

**End-to-End Automation Management
User's Guide and Reference**

Note

Before using this information and the product it supports, read the information in "Notices" on page 375.

Third Edition (March 2006)

This edition applies to Version 2 Release 1 Modification 1 of IBM Tivoli System Automation for Multiplatforms (5724-M00), and to all subsequent releases and modifications until otherwise indicated in new editions.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

A form for readers' comments appears at the back of this publication. If the form has been removed, address your comments to:

IBM Deutschland Entwicklung GmbH
Department 3248
Schoenaicher Strasse 220
D-71032 Boeblingen
Federal Republic of Germany

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 2005, 2006. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures ix

Tables xi

About this guide xiii

Who should read this guide xiii

How to use this guide xiii

Summary of changes for release 2.1.1 xiv

Where to find more information xiv

Conventions used in this guide xiv

 Typeface conventions xiv

 Terminology used in this guide. xv

Related information xvii

Part 1. Introducing end-to-end automation management 1

Chapter 1. What end-to-end automation management can do for you 3

The scope of automated management of resources. 3

The scope of end-to-end automation management of business applications 5

The scope of the operations console 6

Role of an operator 6

Role of an administrator 7

Chapter 2. Components of end-to-end automation management 9

Automation J2EE framework 10

Automation engine 10

Automation manager 11

Automation engine resource adapter 11

First-level automation manager resource adapter 11

Automation adapter 11

Operations console 11

End-to-end automation policy 12

First-level automation domain 12

Automation database 12

Automation Software Development Kit 13

Chapter 3. Operations console modes 15

End-to-end automation mode 15

First-level automation mode 15

Direct access mode 16

Chapter 4. Communication flow between the components. 19

Policy activation and subscription 19

A first-level automation domain sends a resource modified event 21

An operator submits a request against a resource reference 23

The operations console is used in first-level automation mode 24

Chapter 5. Automation concepts. 27

Resources of the end-to-end automation domain 27

 Resource references. 27

 Resource groups. 27

 Choice groups 27

Goal-driven automation 27

How the automation manager is informed about automation goals 28

How the default desired state is determined 29

Understanding relationships. 29

 What is a relationship? 29

 StartAfter relationship 30

 StopAfter relationship 32

 ForcedDownBy relationship 33

How requests become goals 34

Requests processing when relationships exist 35

Request priorities 35

How requests against resource references are processed 36

 User credentials of the end-to-end automation manager 36

 Example scenarios 37

When the end-to-end automation manager will not generate requests 39

 The referenced resource is a monitor resource 39

 The referenced resource is in a transitional state 40

 The referenced resource is in a specific operational state. 40

Requests generated by the end-to-end automation manager are persistent. 40

Canceling obsolete end-to-end automation manager requests on first-level automation resources 41

 Canceling requests on SA for Multiplatforms resources 41

 Canceling requests on SA z/OS resources 42

Part 2. Installation 45

Chapter 6. Planning for installation 47

Packaging 47

 Product CD and WebSphere Application Server

 Upgrade CD 47

 Electronic distribution 48

Product features, DB2 setup options, and user registry options 50

DB2 setup options and user registry options 51

 Using a local or remote DB2 setup. 51

 Using LDAP or DB2 as the user registry. 51

Product requirements 52

 Supported operating systems 52

 Middleware software requirements 53

 Browser requirements 54

| | |
|---|----|
| Hardware requirements | 55 |
| Security concepts | 57 |
| Security considerations | 57 |
| Considerations for choosing between LDAP and DB2 as the user registry | 57 |
| User management | 57 |

Chapter 7. Installing the middleware software 59

| | |
|--|----|
| What the middleware software CDs contain | 59 |
| Installing a DB2 server | 59 |
| DB2 server requirements | 59 |
| DB2 server installation. | 60 |
| Post-installation tasks for remote DB2 setup | 60 |
| Installing a DB2 client | 62 |
| DB2 client requirements | 62 |
| DB2 client installation | 62 |
| Post-installation tasks for remote DB2 setup | 63 |
| Installing WebSphere Application Server. | 64 |
| WebSphere Application Server 6.0.0.0 requirements | 64 |
| Installing WebSphere Application Server 6.0.0.0 | 65 |
| Installing Refresh Pack 2 and the required Interim Fixes | 65 |
| Post-installation tasks | 66 |

Chapter 8. Setting up an LDAP server 67

| | |
|--|----|
| Required LDAP directory tree structure | 67 |
| Required user groups and users | 68 |
| LDAP-related pre-installation tasks | 68 |
| Sample LDAP configuration | 69 |

Chapter 9. Preparing for the installation of the end-to-end automation management component 71

| | |
|--|----|
| Collecting the information you need to provide during installation | 71 |
| Installation directory and Tivoli Common Directory | 71 |
| Installation parameters for DB2. | 72 |
| Installation parameters for WebSphere Application Server | 74 |
| Installation parameters for LDAP | 74 |
| Installation parameters for the operations console | 76 |
| Installation parameters for IBM Tivoli Enterprise Console. | 79 |
| Name of the end-to-end automation domain | 80 |
| What the installation CD contains | 80 |
| Languages supported by IBM Tivoli System Automation | 81 |
| Installation prerequisites | 81 |

Chapter 10. Installing the end-to-end automation management component. . 83

Chapter 11. Verifying the installation . . 99

| | |
|--|----|
| Automation manager | 99 |
| End-to-end automation database | 99 |
| Automation J2EE Framework | 99 |

| | |
|--|-----|
| Verifying that DB2 accepts WebSphere Application Server requests | 99 |
| Automation engine | 100 |
| Operations console | 100 |

Chapter 12. Post-installations tasks 103

| | |
|---|-----|
| Setting up SSL for the operations console | 103 |
| Modifying the LTPA settings | 105 |
| Modifying the HTTP session timeout | 106 |
| Configuring how many users can connect to the automation manager using the operations console | 107 |

Chapter 13. Uninstalling the end-to-end automation management component 109

| | |
|---|-----|
| Launching the graphical uninstallation program on Windows. | 109 |
| Launching the graphical uninstallation program on AIX and Linux | 109 |
| Using the uninstallation program. | 109 |

Chapter 14. Installing service 113

| | |
|---|-----|
| Where to obtain fix packs | 113 |
| Archive naming conventions | 113 |
| Naming conventions of the update installer location | 114 |
| Usage instructions for the platform-specific archives | 114 |
| Windows | 114 |
| AIX. | 114 |
| Linux on IBM x/Series | 115 |
| PPC Linux | 115 |
| Linux on z/Series | 115 |
| Steps for installing a product fix pack | 116 |

Chapter 15. Uninstalling service 117

Chapter 16. Troubleshooting installation problems 119

| | |
|---|-----|
| Installation wizard cannot find WebSphere Application Server on the system. | 119 |
| The installation of the operations console fails | 119 |
| Using the installation log files | 119 |
| Installation log file directories | 119 |
| Installation log files | 120 |
| Procedures for troubleshooting an installation | 123 |
| Gathering information for IBM Support | 124 |

Part 3. First steps 125

Chapter 17. Overview 127

Chapter 18. Starting the sample end-to-end automation domain. 129

Chapter 19. Activating the sample end-to-end automation policy 131

Chapter 20. Creating and activating a new sample policy 133

- Creating a new sample policy 133
- Changing the domain name 134

Chapter 21. Displaying a first-level automation domain on the operations console 137

- Checking the adapter configuration 137
- Starting the adapter 138
- Where to find the first-level automation domain on the operations console 138

Chapter 22. Creating a policy that references actual first-level resources. 139

Part 4. Administering the end-to-end automation management component 141

Chapter 23. Post-installation tasks for administrators 143

- Access roles for end-to-end automation management 143
 - How users are given roles 145
- Creating user groups in Integrated Solutions Console 145
- Assigning access permissions to user groups in Integrated Solutions Console 146
 - Granting user groups access to the pages of Integrated Solutions Console 146
 - Granting user groups access to the operations console of Tivoli System Automation for Multiplatforms 147
- Assigning the user ID of the automation engine to groups in Integrated Solutions Console. 148
- Assigning access roles to user groups in WebSphere Application Server. 149

Chapter 24. Managing users 153

- Managing users and user groups in Integrated Solutions Console 153
 - Creating and authorizing users in Integrated Solutions Console 153
 - Administering users and user groups in Integrated Solutions Console 154
- Managing the user credentials of subcomponents of end-to-end automation management. 156
 - Modifying the default user ID of the automation engine. 156
 - Managing the user ID used by the automation engine to access first-level automation domains . 157
 - Modifying the default user ID used by the automation management server to access DB2 . 157
- Managing JMS authentication 157
 - Modifying the default JMS authentication entry for the automation engine 158

- Modifying the default JMS authentication entry for the operations console 158
- Modifying the default JMS authentication entry for the automation management server. 159

Chapter 25. Creating and modifying policies 161

- What you must know before you define an end-to-end automation policy 162
 - The scope of end-to-end automation policies . 162
 - Identifying cluster-spanning dependencies. . . 164
 - Gathering the required data for defining a policy 166
 - Considerations for referencing SA for Multiplatforms resources 167
 - Restrictions for referencing SA z/OS resources . 167
- Defining an end-to-end automation policy. 168
 - Creating the XML policy file 168
 - Defining the resources of the end-to-end automation domain 171
 - Defining groups 174
 - Defining StartAfter, StopAfter, and ForcedDownBy relationships 177
 - Saving the policy in the policy pool directory . 179
 - Starting the policy checking tool from a command line 179

Chapter 26. Setting up information pages for operators. 181

Chapter 27. Using the command-line interface of the automation engine . . 183

- eezdmn options quick reference 184
- eezdmn options 184
 - start 184
 - shutdown 185
 - monitor 186
 - reconfig 187
 - co 187
 - xd 188
 - ? 188

Chapter 28. Starting and stopping . . 189

- Starting and stopping WebSphere Application Server 189
 - Starting and stopping WebSphere Application Server on Windows 189
 - Starting and stopping WebSphere Application Server on AIX and Linux 190
- Starting and stopping the automation J2EE framework 190
- Starting and stopping the operations console . . 190
 - Starting and stopping the operations console on Windows. 191
 - Starting and stopping the operations console on AIX and Linux 192
- Starting and stopping the automation engine. . . 192

Chapter 29. Configuring the end-to-end automation manager . . . 193

Launching the configuration dialog 193
Configuration dialog pages. 193

Chapter 30. Using Tivoli Enterprise Console with SA for Multiplatforms . . 199

Configuring Tivoli Enterprise Console 199
Checking the Tivoli Event Integration Facility function 200
Enabling Tivoli Enterprise Console event filtering 202
 Activating the default CEI filter 202
 Customizing the default event filter 204

Part 5. Monitoring and managing automated resources 205

Chapter 31. Overview 207

Chapter 32. What you must know about the operations console 209

Configuring your Web browser 209
Logging on 209
 Steps for accessing the operations console 210
Understanding the layout of the operations console 212
What you must know about the topology tree 214
 Navigating the topology tree 215
 Selecting an element in the topology tree 216
 Limiting the scope of the topology tree. 216
 What is displayed in the topology column. 216
 What you can see in the Status column. 217
 What you can see in the Located here column 217
What you must know about the resources section 217
 Section header 218
 View and Search 218
 Resource table views 218
What you must know about the information area 222
What you must know about the Smart refresh bar 223
What you must know about the main menu 223
Using links to quickly jump to a specific element 224

Chapter 33. Monitoring resources . . . 225

State information provided on the operations console 225
 Compound state and operational state 225
 State information provided for domains 226
 State information provided for nodes 231
 State information provided for resources 231
Monitoring tasks 236
 Locating a resource 236
 Switching between resource references and referenced resources 236
 Finding out to which groups a resource belongs 237
 Displaying relationships 237
 Viewing log files 237
 Displaying operator instructions using the info link. 238
 Displaying owner contact information 238

Limiting the scope of the resource table 238
 Displaying only resources that are in an error or warning state 238
 Searching for resources 239
Hiding domains 243
Using non-top-level resources as domain health indicators. 244
Refreshing the operations console 245
Switching to a different end-to-end automation manager 246
 Steps for connecting to a different end-to-end automation manager from the operations console 246
Managing your user credentials for first-level automation domains 247

Chapter 34. Managing resources . . . 249

Working with policies 249
 Activating a policy 249
 Deactivating a policy 250
 Modifying a policy 251
Working with requests 251
 Submitting start requests 251
 Submitting stop requests 252
 Displaying information about an operator request 252
 Displaying request lists 253
 Canceling requests 253
Resetting a resource from an unrecoverable error 254
 Steps for resetting a resource 254
Including a node in automation and excluding a node from automation 255
 Steps for excluding a node from automation 255
 Steps for including a node in automation 255
Working with choice groups 256
 Steps for starting the preferred member of a choice group 256
 Steps for starting a different member of a choice group 257

Part 6. Appendixes 259

Appendix A. Default directories . . . 261

Appendix B. Configuration properties files 263

Appendix C. Automation policy reference 265

Worksheet for gathering data for defining a policy 265
Policy XML reference. 267
 AutomationDomain 267
 AutomationDomainName 267
 AutomationPolicy 268
 ChoiceGroup 268
 Class 269
 Description 270
 DesiredState. 270
 InfoLink 271

| | |
|---|-----|
| Node | 271 |
| Members | 272 |
| Name | 273 |
| Owner. | 273 |
| PolicyAuthor | 273 |
| PolicyDescription | 274 |
| PolicyInformation | 274 |
| PolicyName | 275 |
| PolicyToken | 275 |
| ReferencedResource | 276 |
| Relationship | 276 |
| ResourceGroup | 277 |
| ResourceReference. | 277 |
| Source. | 279 |
| Target | 279 |
| Type | 280 |
| Policy rule reference | 281 |
| Rules for resource groups and choice groups | 281 |
| Relationship rules | 281 |
| Rules that ensure that resource names are | |
| unique | 281 |
| Additional rules | 281 |
| Sample automation policy | 282 |

Appendix D. Additional state-related information about end-to-end automation resources. 287

| | |
|---|-----|
| Values for sa_status_compound | 288 |
| Values for sa_status_desired | 289 |
| Values for sa_status_observed | 289 |
| Values for sa_status_automation | 290 |
| Values for sa_status_control | 290 |

Appendix E. Troubleshooting 293

| | |
|---|-----|
| Where to find the log and trace files. | 293 |
| Where to find the Tivoli Common Directory | 293 |
| Log and trace files of the automation engine | 293 |
| Log and trace files of the automation J2EE | |
| framework and the resource adapters | 294 |
| Trace files of the operations console | 295 |
| How to determine the server port number for | |
| connecting to the operations console | 296 |
| Problems occur when multiple browser windows | |
| are used to connect to the same Integrated | |
| Solutions Console from the same client system | 296 |
| The end-to-end automation domain is not | |
| displayed on the operations console. | 296 |
| A System Automation for Multiplatforms domain | |
| is not displayed in the topology tree. | 296 |
| Critical exceptions in the WebSphere Application | |
| Server log file | 300 |
| Security exception when trying to subscribe to | |
| resources that are hosted on a first-level | |
| automation domain | 301 |

| | |
|--|-----|
| Resolving timeout problems | 301 |
| Watchdog - A mechanism for monitoring the | |
| domain communication states | 302 |
| Database clean-up timeout for automation | |
| domains | 302 |
| Method invocation timeout between the | |
| automation J2EE framework and the automation | |
| adapters | 302 |
| Modifying the environment variables for the | |
| automation J2EE framework | 303 |
| Modifying the time zone settings for the operations | |
| console | 304 |
| Unrecoverable error state displayed for first-level | |
| automation resources is incorrect | 304 |
| WebSphere Application Server cannot be started - | |
| DB2 is used as the user registry | 305 |
| WebSphere Application Server cannot connect to | |
| DB2 | 305 |
| "Unable to set up the event path..." error message | |
| is displayed in Integrated Solutions Console | 306 |
| EEZBus is not started | 306 |
| EEZBus is not started due to a security problem | 306 |
| EEZBus is not started because an internal | |
| database is in an inconsistent state | 307 |

Appendix F. Operations console icons 309

| | |
|--|-----|
| Icons in the topology tree | 309 |
| Resource icons | 310 |
| First-level automation resources | 310 |
| Resource references | 312 |
| Referenced resources | 314 |
| Resource group icons. | 315 |
| Referenced first-level automation groups | 317 |
| Choice groups and first-level automation move | |
| groups | 319 |
| Referenced move groups | 320 |
| Information about requests on the General page | 322 |

Appendix G. Messages 323

| | |
|-------------------------------------|-----|
| How to interpret messages | 323 |
| Messages | 324 |
| Smart refresh messages | 373 |

Notices 375

| | |
|----------------------|-----|
| Trademarks | 376 |
|----------------------|-----|

Glossary 377

Index 381

Figures

| | | |
|-----|---|-----|
| 1. | Components of end-to-end automation management | 10 |
| 2. | Operations console is used for managing first-level automation domains only | 25 |
| 3. | Setup of the end-to-end automation management component | 51 |
| 4. | LDAP directory tree structure | 68 |
| 5. | Log in panel of Integrated Solutions Console | 100 |
| 6. | Welcome panel of Integrated Solutions Console | 101 |
| 7. | Panel for connecting to the operations console | 101 |
| 8. | Common Event Infrastructure Service panel | 199 |
| 9. | Enterprise Applications panel | 199 |
| 10. | Custom properties panel | 200 |
| 11. | Log in panel of Integrated Solutions Console | 210 |
| 12. | Welcome panel of Integrated Solutions Console | 211 |
| 13. | Operations console entry in the navigation tree | 211 |
| 14. | Connect panel | 212 |
| 15. | Main panel of the operations console | 213 |
| 16. | Topology tree and resources section | 215 |
| 17. | Layout of the resources section | 218 |
| 18. | Main menu | 224 |
| 19. | State information on the General page | 232 |
| 20. | Name filters page on the Preferences panel | 242 |
| 21. | Visible automation domains page | 244 |
| 22. | Additional info page for a resource reference | 287 |

Tables

| | |
|--|------|
| 1. End-to-end automation-specific terms | xv |
| 2. Short names used in this guide | xvii |
| 3. Product CD versions | 47 |
| 4. WebSphere Application Server upgrade CD versions | 48 |
| 5. Archives for Windows platforms | 48 |
| 6. Archives for AIX platforms | 49 |
| 7. Archives for Linux on IBM x/Series | 49 |
| 8. Archives for PPC Linux | 49 |
| 9. Archives for Linux on z/Series | 50 |
| 10. Supported operating systems | 52 |
| 11. Disk space requirements on Windows systems | 55 |
| 12. Disk space requirements on AIX and Linux systems | 56 |
| 13. Contents of the WebSphere Application Server 6.0.0 upgrade CD | 65 |
| 14. Installation directory and Tivoli Common Directory | 71 |
| 15. DB2 data for local and remote DB2 setup | 73 |
| 16. WebSphere Application Server installation parameters | 74 |
| 17. Installation parameters for LDAP | 74 |
| 18. Installation parameters for Integrated Solutions Console | 77 |
| 19. Installation parameters for IBM Tivoli Enterprise Console | 80 |
| 20. Name of the end-to-end automation domain | 80 |
| 21. Directories on the product CD | 80 |
| 22. Windows platforms | 114 |
| 23. AIX platforms | 114 |
| 24. Linux on IBM x/Series | 115 |
| 25. PPC Linux | 115 |
| 26. Linux on z/Series | 115 |
| 27. Location of the installation log files | 120 |
| 28. Installation logs in the WebSphere Application Server profile log directory or in the system temporary directory | 122 |
| 29. Installation logs in the system temporary directory | 122 |
| 30. WebSphere Application Server access roles for end-to-end automation management | 144 |
| 31. Mapping of end-to-end automation management roles to groups in WebSphere | 149 |
| 32. Recommendations for referencing SA for Multiplatforms resources in end-to-end automation policies | 167 |
| 33. Steps for defining a new end-to-end automation policy | 168 |
| 34. Command line options for the automation engine | 184 |
| 35. Messages and return codes returned by the automation engine | 186 |
| 36. Valid XPath event selectors | 204 |
| 37. Icons used for the elements of the topology tree | 216 |
| 38. Some flavors of topology tree icons | 216 |
| 39. Icons in the Status column of the topology tree | 217 |
| 40. Compound state icons | 226 |
| 41. Operational state descriptions provided on the General page for a domain | 227 |
| 42. Domain state icons | 229 |
| 43. Communication state | 230 |
| 44. Observed state of a node | 231 |
| 45. Operational state descriptions on the General page for a resource | 232 |
| 46. Operator request icons in the information area | 252 |
| 47. Default directories | 261 |
| 48. Worksheet for defining an end-to-end automation policy | 265 |
| 49. Internal automation states on the Additional info page | 288 |
| 50. Values for sa_status_compound | 288 |
| 51. Values for sa_status_desired | 289 |
| 52. Values for sa_status_observed | 289 |
| 53. Values for sa_status_automation | 290 |
| 54. Values for sa_status_control | 290 |

| | | |
|-----|--|-----|
| 55. | Environment variables of the automation J2EE framework | 301 |
| 56. | Domain icons on the operations console | 309 |
| 57. | Node icons on the operations console | 310 |
| 58. | Icons for first-level automation resources | 310 |
| 59. | Icons for resource references | 312 |
| 60. | Icons for referenced resources | 314 |
| 61. | Resource group icons | 315 |
| 62. | Icons for referenced first-level automation groups | 317 |
| 63. | Icons for choice groups and first-level automation move groups | 319 |
| 64. | Referenced move groups. | 320 |
| 65. | Operator request icons in the information area | 322 |

About this guide

This guide provides information about installing and using the end-to-end automation management component of IBM Tivoli System Automation for Multiplatforms.

After reading this document, you will understand:

- How to plan for end-to-end automation management
- How to install the end-to-end automation management component
- How to administer the end-to-end automation management component
- How to define and activate automation policies for end-to-end automation management
- How to use the graphical user interface of the end-to-end automation management component for monitoring and managing resources
- How to service the end-to-end automation management component

Who should read this guide

This guide is for planners, installers, and administrators who plan to install, configure, and administer the end-to-end automation management component of IBM Tivoli System Automation for Multiplatforms, and for operators who want to monitor and manage resources from the operations console.

How to use this guide

Use the parts of this guide that correspond to the job that you will do:

- Part 1, “Introducing end-to-end automation management,” on page 1 gives you an overview of end-to-end automation management, its goals, the automation concepts, and the functionality provided by the end-to-end automation management component.
- Part 2, “Installation,” on page 45 describes how to plan and prepare for the installation of the end-to-end automation management component, how to install it, and how to customize it to prepare it for use.
- Part 3, “First steps,” on page 125 describes in how you can use sample environment that is configured during installation to learn about end-to-end automation management.
- Part 4, “Administering the end-to-end automation management component,” on page 141 describes how to create policies, manage users, and start and stop the components of end-to-end automation management.
- Part 5, “Monitoring and managing automated resources,” on page 205 describes how to exploit the functionality of end-to-end automation management.
- In the Appendixes you find reference information you may need for using and operating the end-to-end automation management component.

Summary of changes for release 2.1.1

This section provides an overview of changes in the end-to-end automation management component of IBM Tivoli System Automation for Multiplatforms Version 2, Release 1, Modification 1.

What's new for installers:

Installing level 2.1.1 or upgrading to level 2.1.1

To obtain level 2.1.1 of the end-to-end automation management component, you can:

- install level 2.1.1 from scratch
- upgrade from level 2.1 or 2.1.0.1

What's new for operators:

Redesigned operations console

The operations console has been redesigned:

- Resources that are hosted by a domain or node are presented in two different table views. Which view is displayed depends on whether or not you are displaying the results of a search. The table views replace the tree format in which the resources were presented in earlier versions.
- An information bar displays a confirmation message whenever an action was performed against an element on the operations console.

For more information, refer to Chapter 32, "What you must know about the operations console," on page 209.

Where to find more information

The IBM Tivoli System Automation for Multiplatforms home page contains useful up-to-date information, including support links and downloads for maintenance packages.

You find the IBM Tivoli System Automation for Multiplatforms home page at:

www.ibm.com/software/tivoli/products/sys-auto-linux/

Conventions used in this guide

This guide uses several conventions for special terms and actions and operating system commands and paths.

Typeface conventions

This guide uses the following conventions:

- Typically, file names, directories, and commands appear in a different font. For example:
 - File name: `setup.jar`
 - Directory: `/etc/hosts`
 - Command: `startServer server1`
- Variables are either italicized, enclosed in brackets, or both. For example:
 - `http://<hostname.yourco.com>/index.html`
- Frequently, variables are used to indicate a root installation directory:

- Root installation directory of the end-to-end automation management component:
<EEZ_INSTALL_ROOT> or EEZ_INSTALL_ROOT
- WebSphere Application Server root installation directory: <was_root> or was_root
- Runtime root directory of Integrated Solutions Console: <isc_runtime_root> or isc_runtime_root

For more information on such variables, see Appendix A, “Default directories,” on page 261.

- Directories are shown with forward slashes (/), unless operating-system specific information is provided. On Windows systems, you should use backward slashes (\) when typing at a command line, unless otherwise noted.
- Operating-system specific information is provided. For example:
 - **AIX, Linux:** /opt/IBM/tsamp/eez
 - **Windows:** C:\Program Files\IBM\tsamp\eez

Terminology used in this guide

This section describes terms that are specific to end-to-end automation management and that you will frequently encounter in this manual, in other publications related to end-to-end automation management, and on the operations console.

Two different types of terms are introduced in this section:

- The end-to-end automation specific terms that are important for understanding the concepts of end-to-end automation management.
- The short forms of terms that are used in this guide to ensure readability.

End-to-end automation-specific terminology

In the following table, you find the definitions of important terms related to end-to-end automation management. Additional terms are described in Chapter 2, “Components of end-to-end automation management,” on page 9 and in the glossary.

Table 1. End-to-end automation-specific terms

| Term | Description |
|--------------------|--|
| choice group | An end-to-end automation resource group whose members are alternatives. Only one of the members can be active at a time. If the desired state of the choice group is Online, the end-to-end automation manager tries to keep the active resource online but will only start the resource in place if it fails. An operator can start a different member of a choice group from the operations console. |
| direct access mode | An operations console mode in which only resources that are automated by the Base component of IBM Tivoli System Automation for Multiplatforms can be managed and monitored from the console. |

Table 1. End-to-end automation-specific terms (continued)

| Term | Description |
|-----------------------------|---|
| domain health indicators | <p>Resources whose state is used to indicate whether or not a domain is healthy. If the observed state of such a resource differs from its desired state, an error or warning appears on the operations console for the domain by which it is hosted.</p> <p>This makes it possible to monitor resources simply by observing the domains in the topology tree and drilling down to resource level only when a problem is indicated for the domain.</p> <p>By default, a domain's top-level resources are used as domain health indicators. On the operations console you can define that other resources are to be used for this purpose.</p> |
| end-to-end automation mode | An operations console mode in which end-to-end automation management is installed and active. In this mode, resources that are hosted by the end-to-end automation domain and by first-level automation domains can be monitored and managed from the operations console. |
| first-level automation mode | An operations console mode in which only resources that are hosted by first-level automation domains can be monitored and managed from the console. The end-to-end automation management component is installed but end-to-end automation management is not active. |
| monitor resource | <p>A first-level automation resource that has the following characteristics:</p> <ul style="list-style-type: none"> • its current state can be monitored from the operations console • its desired state cannot be changed through start and stop requests |
| resource | <p>Any application, process, or service that is monitored and managed by a first-level or end-to-end automation manager.</p> <p>If not stated otherwise, the term is used to refer to both resources and groups of resources on any automation level and on the specific automation level described in the context in which the term appears.</p> |
| resource group | In end-to-end automation management, a collection of resource references that have the same desired state and are managed and monitored as one unit. The first-level resources referenced by the resource references in a group can be hosted by different first-level domains. Resource groups are defined in the end-to-end automation policy. |
| resource reference | A resource that is managed by the end-to-end automation manager. Resource references are virtual resources that refer to actual resources that are managed by a first-level automation manager. Resource references are defined in the end-to-end automation policy. |
| top-level resource | A resource or resource group that is displayed in the resource table when a domain is first selected. Typically, these are resources that are either not members of a group, or groups that are not nested within other groups. By default, such resources are used as <i>domain health indicators</i> . |

Short names used in this guide

To ensure the readability of this guide, short names are used for some products and for some of the subcomponents of the end-to-end automation management

component of IBM Tivoli System Automation for Multiplatforms. The full names are used whenever the context demands it. For example, the end-to-end automation policy will usually be referred to as policy, however, when it might not become clear from the context whether the term refers to the policy of the end-to-end automation domain or to that of a first-level automation domain, the full term is used.

Table 2. Short names used in this guide

| Term used in this guide | Used for.. |
|--|---|
| automation adapter | end-to-end automation management adapter |
| automation engine | end-to-end automation decision engine |
| automation manager | end-to-end automation manager |
| end-to-end automation management component end-to-end automation management | end-to-end automation management component of IBM Tivoli System Automation for Multiplatforms |
| operations console SA operations console | operations console of IBM Tivoli System Automation for Multiplatforms |
| policy | end-to-end automation policy |
| SA for Multiplatforms | IBM Tivoli System Automation for Multiplatforms |
| SA z/OS | IBM Tivoli System Automation for z/OS |

Related information

Use the following IBM Tivoli System Automation guides to obtain more information about the first-level automation products:

- Tivoli System Automation for Multiplatforms:
 - *IBM Tivoli System Automation for Multiplatforms Base Component User's Guide*, SC33-8210
 - *IBM Tivoli System Automation for Multiplatforms Base Component Reference*, SC33-8212
- Tivoli System Automation for z/OS
 - *System Automation for z/OS V3R1 End-to-End Automation Adapter*, SC33-8271

The latest versions of all WebSphere Application Server publications can be found on the WebSphere Application Server library Web site at

www.ibm.com/software/webservers/appserv/was/library/

DB2 documentation can be found on the IBM DB2 UDB Web site at

www.ibm.com/software/data/db2/udb/support/

You find the link to the PDF manuals in the **Other resources** section on the Web page.

Part 1. Introducing end-to-end automation management

| | |
|---|----|
| Chapter 1. What end-to-end automation management can do for you | 3 |
| The scope of automated management of resources | 3 |
| The scope of end-to-end automation management of business applications | 5 |
| The scope of the operations console | 6 |
| Role of an operator | 6 |
| Role of an administrator | 7 |
| Chapter 2. Components of end-to-end automation management | 9 |
| Automation J2EE framework | 10 |
| Automation engine | 10 |
| Automation manager | 11 |
| Automation engine resource adapter | 11 |
| First-level automation manager resource adapter | 11 |
| Automation adapter | 11 |
| Operations console | 11 |
| End-to-end automation policy | 12 |
| First-level automation domain | 12 |
| Automation database | 12 |
| Automation Software Development Kit | 13 |
| Chapter 3. Operations console modes | 15 |
| End-to-end automation mode | 15 |
| First-level automation mode | 15 |
| Direct access mode | 16 |
| Chapter 4. Communication flow between the components | 19 |
| Policy activation and subscription | 19 |
| A first-level automation domain sends a resource modified event | 21 |
| An operator submits a request against a resource reference | 23 |
| The operations console is used in first-level automation mode | 24 |
| Chapter 5. Automation concepts | 27 |
| Resources of the end-to-end automation domain | 27 |
| Resource references | 27 |
| Resource groups | 27 |
| Choice groups | 27 |
| Goal-driven automation | 27 |
| How the automation manager is informed about automation goals | 28 |
| How the default desired state is determined | 29 |
| Understanding relationships | 29 |
| What is a relationship? | 29 |
| StartAfter relationship | 30 |
| Details on the start behavior of the StartAfter relationship | 30 |
| StopAfter relationship | 32 |
| Details on the stop behavior of the StopAfter relationship | 32 |
| ForcedDownBy relationship | 33 |
| Details on the force down behavior of the ForcedDownBy relationship | 34 |
| How requests become goals | 34 |
| Requests processing when relationships exist | 35 |
| Request priorities | 35 |
| How requests against resource references are processed | 36 |
| User credentials of the end-to-end automation manager | 36 |
| Example scenarios | 37 |
| A policy is activated | 38 |
| An operator issues a request against a resource reference | 38 |
| The state of a referenced resource changes | 39 |
| When the end-to-end automation manager will not generate requests | 39 |
| The referenced resource is a monitor resource | 39 |
| The referenced resource is in a transitional state | 40 |
| The referenced resource is in a specific operational state | 40 |
| Requests generated by the end-to-end automation manager are persistent | 40 |
| Canceling obsolete end-to-end automation manager requests on first-level automation resources | 41 |
| Canceling requests on SA for Multiplatforms resources | 41 |
| Example: The referenced resource is a first-level automation resource group | 41 |
| Example: The referenced resource is a first-level automation resource | 42 |
| Canceling requests on SA z/OS resources | 42 |

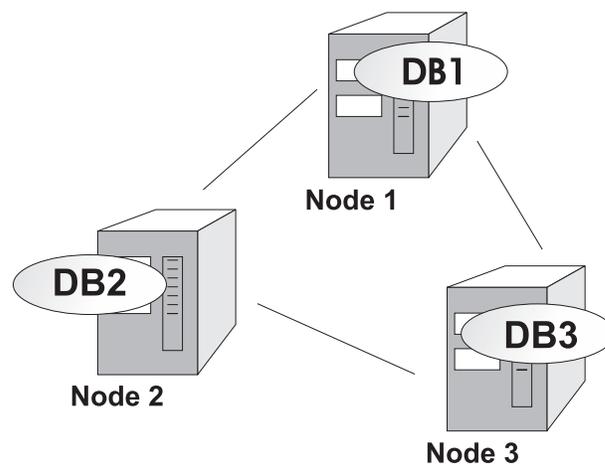
Chapter 1. What end-to-end automation management can do for you

The scope of automated management of resources

Automation means that a certain desired run time behavior of Information Technology (IT) can be described in a formal way and that an automation decision instance, the so-called automation engine, performs tasks on behalf of a human operator.

This is true for many aspects of operations management. The focus of IBM Tivoli System Automation is on automating the availability of IT resources. This is defined as the capability to automatically start and stop IT resources, typically, these are applications. The automation engine acts based on the understanding of operationally related resources and with the knowledge of alternative resource instances that provide the same service in case of outages.

The following figure shows an example in which the databases can run on three different nodes.



When you use SA for Multiplatforms, you no longer need to specify event correlation rules in sophisticated scripts. Such scripts would describe the desired behavior in complex lists such as

```
If (DB3 failed) and (Node 1 running) then (start DB1) else...
```

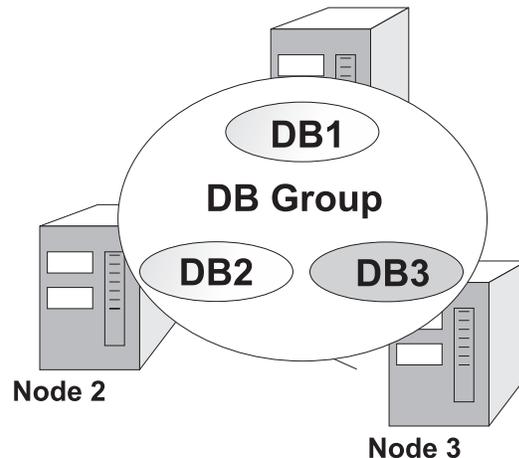
IBM Tivoli System Automation offers a resource management model with a relationship graph and a set of defined abstract resource states as input. The knowledge about how state changes of specific resources are propagated to the related resources is expressed by the semantic of the relationship rather than by exposing those scripting rules.

All required actions are submitted by the automation engine when the desired state and the current situation require an intervention. All you need to describe is the resource topology, namely, the resources, and their relationships and grouping dependencies.

The input specification is done in a so-called automation policy document. Resource groups of different types define the special semantics of the automation behavior of the members inside a group. For example, a group can express that all members must be started and stopped together. Another group type might express that its members are alternatives to each other. Such a group would always allow only one member to run at a time.

Groups also provide aggregated state information about their members. This gives an operator the opportunity to immediately see whether all required and dependent resources are in their desired state. In IBM Tivoli System Automation groups can even be nested, which gives an operator an ever increasing entry point for controlling and monitoring resources.

The following figure shows an example of a so-called move group. The members of move group "DB group" are alternative instances of resource "DB". An instance of resource "DB" is available on each node and the instances are alternatives. For example, if the database on Node 3 fails, Tivoli System Automation chooses one of the alternatives on another node.



You can also define relationships between resources in the policy. Relationships can define:

- sequences for the start and stop behavior of resources
- fault scopes: when one resource fails another resource is forced down
- location constraints: a resource must always or must never run on the same node as another resource

The end-to-end automation management component of SA for Multiplatforms includes a set of products that implement this notion of automation. The technology can be used to describe typical High Availability (HA) scenarios based on HA clustered environments, but can also be used to coordinate the start and stop behavior of heterogeneous distributed applications.

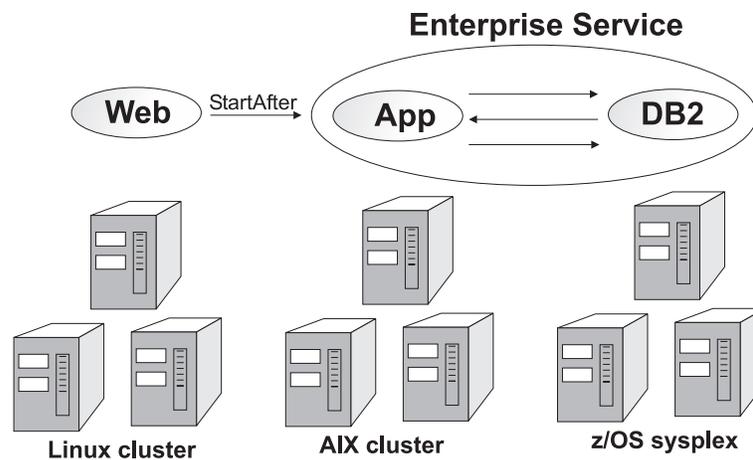
The scope of end-to-end automation management of business applications

This section focuses on the automation aspects of heterogeneous distributed applications with the assumption that many of the resource relationships which are valid in a homogeneous peer node cluster are also of use in heterogeneous environments. For example, the possibility to group IT resources to define a higher level entity is extremely useful to model IT business applications.

Cluster-spanning start and stop ordering is also valid between services on distributed tiers, and the possibility to reflect an overall availability state on a resource that represents the overall business application level is definitely valuable.

The scope of the end-to-end automation management component of SA for Multiplatforms is the automation of operations-related tasks in an environment that consists of multiple server clusters. Each individual server cluster is homogeneous because it is comprised of servers running the same operating system and system software. However, multiple server clusters may each have another operating system environment.

Instead of re-inventing resource management of individual resources at the heterogeneous cluster level, end-to-end automation management makes use of the automation solution that is available on each homogeneous cluster. This functionality is provided, for example, by the other products of the IBM Tivoli System Automation (SA) product family, namely, SA for Multiplatforms and SA z/OS.



In this manual, an automation solution on a homogeneous cluster is called a first-level automation domain. End-to-end automation management does not replace these first-level automation domains but rather builds upon and integrates them.

In the example shown in the figure above the resource Web, which is defined on a Linux cluster, has a startAfter relationship to the group Enterprise Service, which consists of resources that are running on an AIX cluster and on a z/OS sysplex.

In end-to-end automation management, the resources App and DB2, although running on different clusters, can have relationships to each other (which are not further specified in the figure above).

The scope of first-level automation domains is to ensure the high availability of resources as specified in their local (first-level) automation policy. The scope of end-to-end automation is to control the relationships these resources have that span the first-level automation cluster boundary. End-to-end automation does not replace the first-level automation products. Rather, it sends requests to the first-level automation domains in order to accomplish the goals specified in the end-to-end automation policy.

If an operator submits a request to start the resource Web in the example above, end-to-end automation management will first start the resource group Enterprise Service. This is because end-to-end automation sends the requests to start App and DB2 in the correct sequence to the two first-level automation clusters AIX Cluster and z/OS Sysplex. After the resources App and DB2 have been started successfully by the first-level automation product, the group Enterprise Service changes to a Started state, which satisfies the startAfter relationship of the resource Web. End-to-end automation now sends a request to bring Web online on the Linux cluster.

The scope of the operations console

SA for Multiplatforms provides a user front-end, the so-called operations console, that can be used by operators for monitoring and controlling the availability status of all automated resources. The operations console provides this capability on a domain-spanning level. This means that an operator can monitor all automated resources in the enterprise environment from a single console.

This has two major benefits:

- Operators who monitor and manage automated resources that are hosted by clusters of systems spanning different operating systems do not need to have specific knowledge about the particular operating systems.
- Different automation products can be used on different local clusters. An operator does not have to know the different automation concepts or learn how to work with native automation product-specific front-ends (native user interfaces).

To realize these benefits, the automation products must meet the following requirements:

- They must have a common set of resource availability states.
- They must have a common set of operations an operator can perform against the automated resources.

This means that the native user interface may still be required for particular, highly specialized operations and for performing some product-specific monitoring and problem analysis tasks.

Role of an operator

An operator is defined as a person whose responsibility is to ensure the continuous availability of all business-relevant IT resources within a specific enterprise.

An operator must mainly accomplish two major tasks:

- Perform planned maintenance work on IT resources. Resources can be systems, networks, or applications. Maintenance can include applying fixes, replacing defective hardware, and applying (preventive) fixes to applications.

- React to problems. Whenever an IT resource encounters a problem, the operator must be alerted. The operator is in charge of finding the root cause of the problem and resolving it as quickly as possible.

The operations console of SA for Multiplatforms provides a user interface that is designed to support an operator in accomplishing these tasks.

Role of an administrator

The task of an administrator is to define and set up the relationships of IT resources in the data center of the enterprise. In this document it is assumed that administrators are typically not involved in the daily business of keeping the business-relevant IT-resources running. They have a supporting role, they specify automation policies and help operators to resolve severe problems.

Specifying automation policies includes defining automation policies, verifying the correct logic of the policies by running the policy checking tool, and activating the policies from the operations console. These tasks may be performed first on some test systems before the policies are activated on the production systems.

Administrators may also use the operations console to drill down to those applications whose failure is the root cause of a problem.

Chapter 2. Components of end-to-end automation management

This chapter provides an overview of the following components of end-to-end automation management:

- “Automation J2EE framework” on page 10
- “Automation engine” on page 10
- “Automation manager” on page 11
- “Automation engine resource adapter” on page 11
- “First-level automation manager resource adapter” on page 11
- “Automation adapter” on page 11
- “Operations console” on page 11
- “End-to-end automation policy” on page 12
- “First-level automation domain” on page 12
- “Automation database” on page 12
- “Automation Software Development Kit” on page 13

The relationships among the components are illustrated in the following diagram.

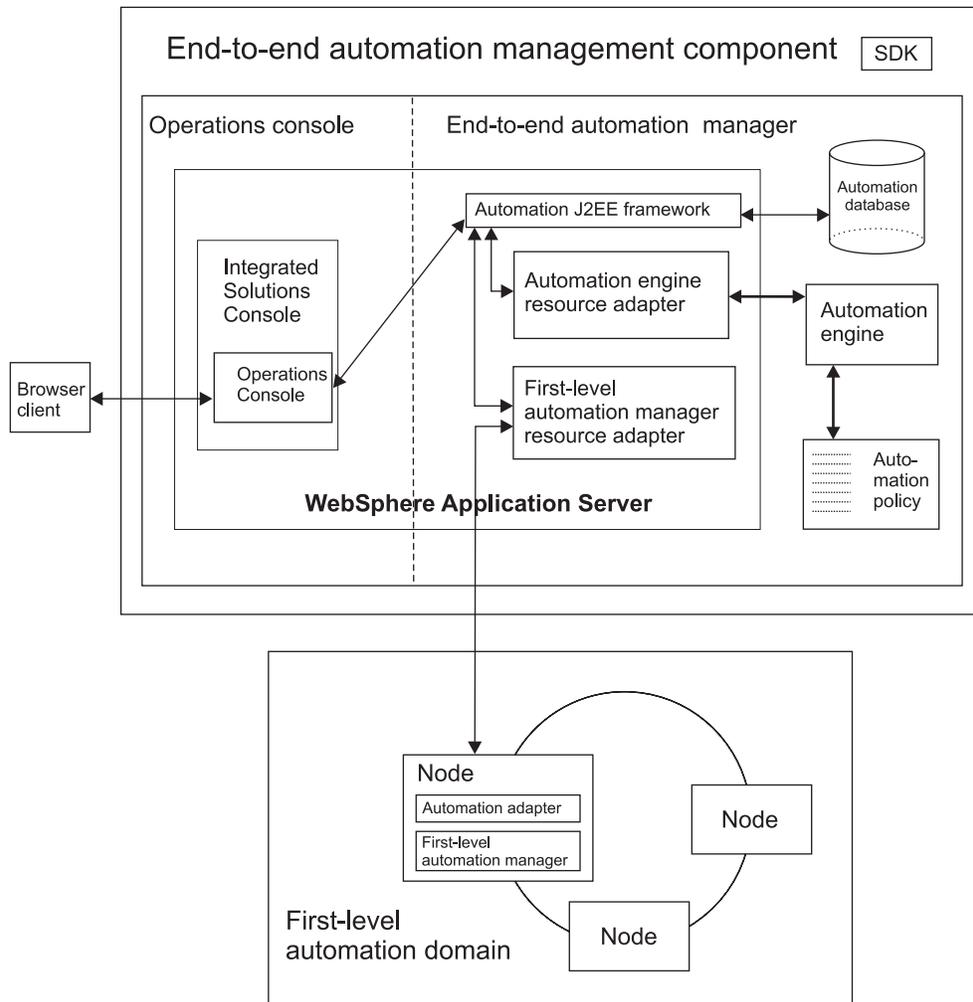


Figure 1. Components of end-to-end automation management

Automation J2EE framework

The automation J2EE framework comprises the components that are deployed within WebSphere Application Server during the installation of the end-to-end automation management component and that act as communication framework between the first-level automation domains and the automation engine and the operations console. Together with the automation database, the framework ensures that required automation domain data and operator preferences are kept in persistent storage.

Automation engine

The automation engine is the decision-making component of the automation manager. It runs as a separate process (daemon or service) on the same system as the WebSphere Application Server where the automation J2EE framework has been installed and is running. The automation engine is notified when the current (observed) state of referenced resources has changed. The automation engine compares the observed state of the resource with its desired state that is defined in the end-to-end automation policy and calculates resulting start or stop requests.

With the help of the automation J2EE framework, the resulting requests are sent to the first-level automation domain that hosts the referenced resource in order to reach the desired state.

The automation engine has to be started by using its command line interface. After startup, the automation engine is displayed as end-to-end automation domain on the operations console. After startup, the automation engine is idling until an end-to-end automation policy is activated.

Automation manager

The term describes the combination of the automation J2EE framework and the automation engine. The end-to-end automation manager's role concerning the management of resource references specified in the end-to-end automation policy can be compared to that of the automation managers that are running on first-level automation domains with respect to the resources managed by them.

Automation engine resource adapter

This resource adapter is a J2EE component that is required by the automation J2EE framework in order to communicate with the automation engine. It is based on the standard J2EE connector architecture. As any other resource adapter, it is deployed and managed using the Administrative Console of WebSphere Application Server.

First-level automation manager resource adapter

This resource adapter is a J2EE component that is required by the automation J2EE framework in order to communicate with the automation adapters that run on the first-level automation domains. It is based on the standard J2EE connector architecture. As any other resource adapter, it is deployed and managed using the Administrative Console of WebSphere Application Server.

The first-level automation manager resource adapter is responsible for all synchronous communication paths to the first-level automation domains. However, the automation engine must always be started in order to receive an event when the state of a resource changes that is hosted by a first-level automation domain.

Automation adapter

An automation adapter process must run on each first-level automation domain. Together with the first-level automation manager resource adapter, the automation adapter ensures normalized communication between the end-to-end automation J2EE framework and the automation manager of the first-level automation domain.

Operations console

The operations console is the Web-based graphical user front-end to the end-to-end automation domain and to the first-level automation domains. Operators use the operations console to monitor the resources that are hosted by the automation domains and to change their states by submitting requests against them.

The operations console consists of the following parts:

- WebSphere Portal Server, which is based on WebSphere Application Server
- Integrated Solutions Console, which runs as an application in WebSphere Portal Server.

Integrated Solutions Console can host multiple application front-ends. Tivoli System Automation is one of these applications.

- The operations console, which is the actual front-end that is used by operator. The operations console runs within Integrated Solutions Console. Operators use a Web browser to contact Integrated Solutions Console and to display the operations console.

For information about the different modes in which you can run the operations console, refer to Chapter 3, “Operations console modes,” on page 15.

End-to-end automation policy

The policy is defined in an XML file. The file contains the definitions of all resource references, groups and relationships which will be managed by the end-to-end automation domain. The document will be read by the end-to-end automation manager at policy activation time. The automation manager will automatically set up the links between the end-to-end automation domain and any available or joining first-level automation domains hosting resources that are referenced by resource references in the currently activate policy.

The end-to-end automation policy describes:

- The aggregation of resource references and of groups of resource references. By gathering resource references in groups and by building group hierarchies, the aggregated state of a complete enterprise application can be monitored easily. In addition, because all members of a group can be started or stopped through a single request, only one request is needed to start or stop all resources that are required by a business application, which may be distributed over multiple first-level automation domains.
- The relationships between resource references, such as which resource must be started before another resource can be activated.
- The desired states of the resource references. The desired state is the automation goal the end-to-end automation manager tries to reach by keeping each defined resource reference in this state.

First-level automation domain

This term is used for an automation back-end hosting resources that are managed by some automation management product, for example, a Linux cluster on which the applications are automated by SA for Multiplatforms. Such a cluster becomes a first-level automation domain when an automation adapter has been installed and configured and is running on one of the nodes of the cluster. Only resources that are managed by a first-level automation domain can be the target of resource references.

Automation database

The automation database is needed by the automation J2EE framework in order to store persistent information about automation domains (the end-to-end automation domain and first-level automation domains) and operator preferences. The database also holds some information about the currently active automation policy. However, the policy itself is not stored in the database. The policy itself is made persistent by specifying it as an XML document and placing it in the policy pool directory which is used by the automation engine.

Automation Software Development Kit

The Automation Software Development Kit defines a set of classes that are used by all other end-to-end automation subcomponents. These classes represent the EEZ common data model and the methods that are needed to access it. The Automation Software Development Kit component is not visible as a running part neither by itself nor within WebSphere Application Server. However, references to the classes may appear in messages in various trace and log files which are written by subcomponents of the end-to-end automation management component.

Chapter 3. Operations console modes

This chapter gives an overview of the three different modes in which the operations console of Tivoli System Automation for Multiplatforms can be used.

End-to-end automation mode

In this mode, end-to-end automation management is active. From the operations console, you can monitor and manage the resources of the end-to-end automation domain and of the first-level automation domains that are connected to the end-to-end automation manager.

Prerequisites for using the operations console in end-to-end automation mode:

- The end-to-end automation management component is installed.
- The end-to-end automation manager is running.
- The automation adapters on the first-level automation domains are configured to send events to the end-to-end automation manager (end-to-end automation mode).
- The automation adapters are running.
- An end-to-end automation policy is active.

This is what you will see on the operations console in end-to-end automation mode:

- In the topology tree, the end-to-end automation domain is displayed.
- First-level domains hosting resources that are referenced in the end-to-end automation policy are displayed as child domains of the end-to-end automation domain.
- First-level domains that are not hosting resources that are referenced in the end-to-end automation policy appear at the same level of the domain hierarchy as the end-to-end automation domain.

This is what you can do on the operations console in end-to-end automation mode:

- You can monitor and manage the resources that are hosted by the end-to-end automation domain and by the first-level automation domains that are connected to the end-to-end automation manager.
- You can activate and deactivate end-to-end automation policies.
- You can perform the full set of tasks described in Part 5, “Monitoring and managing automated resources,” on page 205.

First-level automation mode

In this mode, end-to-end automation management is not active. You can use the operations console for monitoring and managing resources that are managed by SA for Multiplatforms or SA z/OS or both.

Prerequisites for using the operations console in first-level automation mode:

- The end-to-end automation management component is installed.
- The automation engine of the end-to-end automation management component is started in conversion-only mode. In conversion-only mode, the automation

engine is only used for converting events into the required format. No end-to-end automation domain is available and no end-to-end automation is performed.

- The automation adapters on the first-level automation domains are configured to send events to the end-to-end automation manager (end-to-end automation mode).
- The automation adapters are running.

This is what you will see on the operations console in first-level automation mode:

- In the topology tree, all automation domains appear at the same level.

This is what you can do on the operations console in first-level automation mode:

- You can monitor and manage the resources of the first-level automation domains that are connected to the end-to-end automation manager.
- Some functions that are available when end-to-end automation management is active, are not available.

For detailed information about the communication flow that occurs when the operations console is used in first-level automation mode, refer to “The operations console is used in first-level automation mode” on page 24. For information on how you start the automation engine in conversion-only mode, refer to Chapter 27, “Using the command-line interface of the automation engine,” on page 183.

Direct access mode

In this mode, you use the operations console for monitoring and managing resources that are managed by SA for Multiplatforms. On the system on which the operations console of Tivoli SA for Multiplatforms is installed, the end-to-end automation management component must not be installed.

Note: It is not possible to connect a first-level automation domain to both an operations console in direct access mode and to an operations console in one of the other modes at the same time.

The description of how to use the operations console in direct access mode is not within the scope of this guide. For detailed information on installing, configuring, and using the operations console in direct access mode, refer to the *IBM Tivoli System Automation for Multiplatforms Base Component User's Guide*, SC33-8210.

Prerequisites for using the operations console in direct access mode:

- The operations console of SA for Multiplatforms is installed.
- The automation adapters on the SA for Multiplatforms automation domains are configured to send events to the SA for Multiplatforms operations console (direct access mode).
- The automation adapters are running.

This is what you will see on the operations console in direct access mode:

- In the topology tree, you see the automation domains of SA for Multiplatforms.

This is what you can do on the operations console in direct access mode:

- You can monitor and manage the resources that are hosted by the SA for Multiplatforms automation domains.

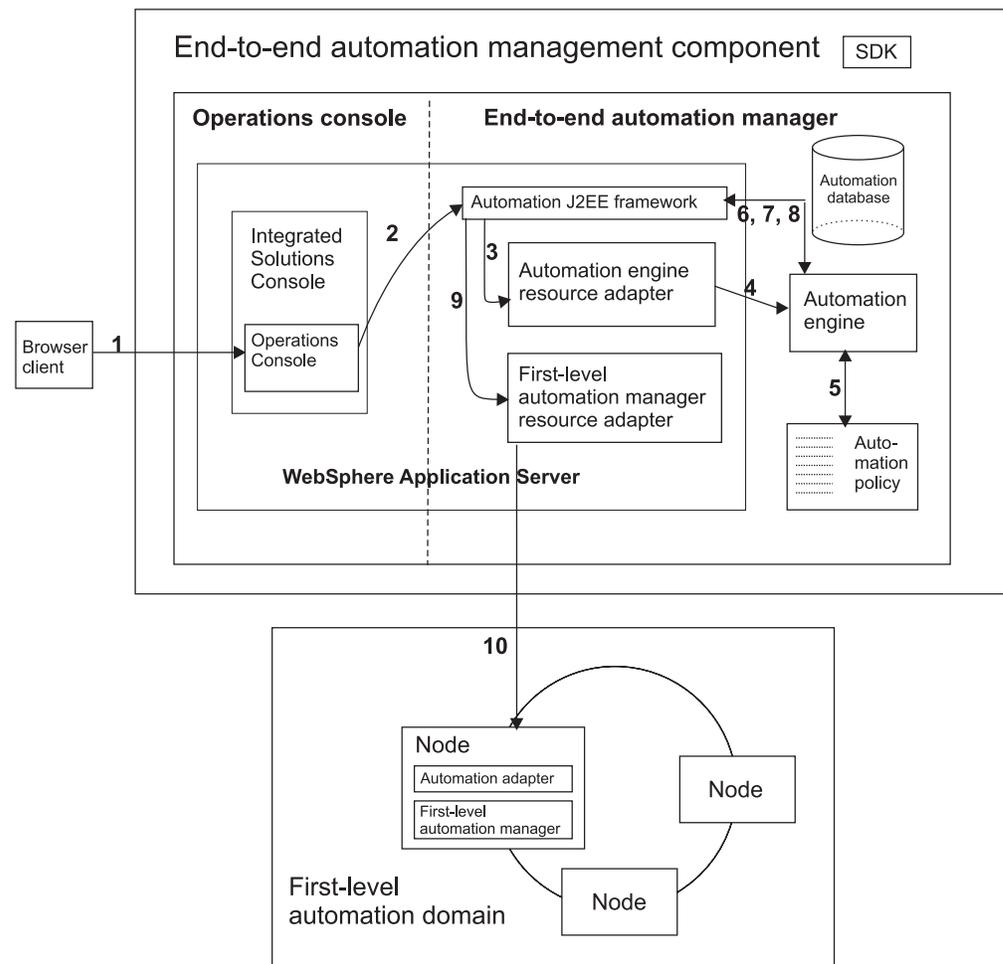
- Some functions that are available when end-to-end automation management is active, are not available.

Chapter 4. Communication flow between the components

The following sections provide an overview of the communication flows that occur between the components involved in end-to-end automation management.

Policy activation and subscription

The following figure shows the communication flow between the components that occurs when a new policy is activated:



This is a description of the scenario shown in the figure above:

1. An operator selects the policy to be activated from the list of policies that are available in the policy pool directory.
2. The name of the policy is passed to the automation J2EE framework with the request for activation.
3. The request is passed to the automation engine resource adapter.

4. The policy activation request is passed to the automation engine.

5. The automation engine loads the policy from the policy pool directory.

6. The automation engine parses the policy XML document and creates all resources, groups, and relationships within its internal storage structure.
At this time, the automation engine has no information about the observed state of any of the defined resource references. It also does not know if the first-level automation domains hosting the referenced resources defined in the policy are currently online. This is why the automation engine now subscribes to the automation J2EE framework to be informed about the state of any first-level automation domain that hosts referenced resources.

7. The automation J2EE framework returns a list of all first-level automation domains that are currently online.
(From then on, the automation engine will be informed of all state changes in the domains it subscribed for, for example, when an automation adapter sends its domain join event at a later time.)

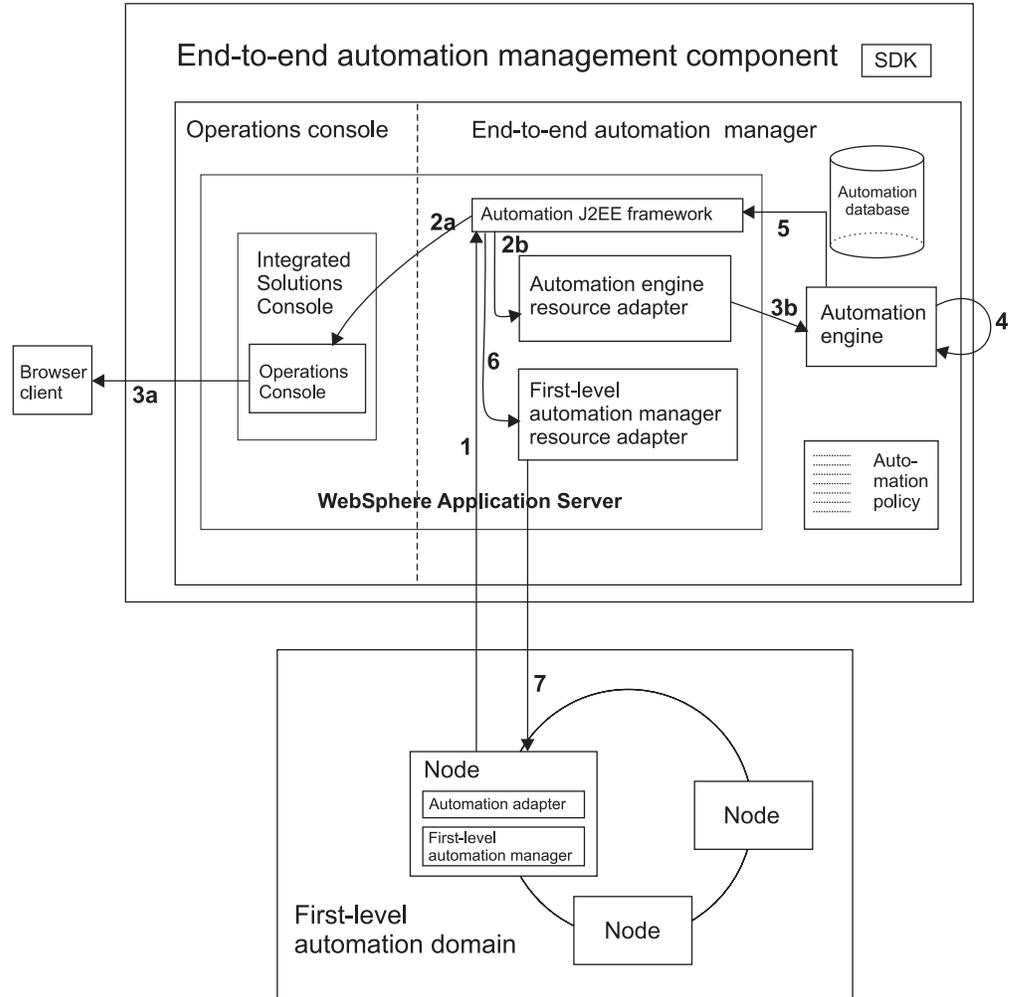
8. The automation engine subscribes to the resources hosted by the first-level automation domains which were returned in step 7. This is done because the automation engine needs to get informed about the current (observed) state of all resources in this first-level automation domain in order to calculate the states and resulting requests for the resource references defined in the automation policy.

9. The subscription for state changes of resources is passed to the first-level automation manager resource adapter.

10. The subscription is passed to the first-level automation domain. From now on, the automation engine will be informed whenever the state of one of the resources it subscribed for changes.

A first-level automation domain sends a resource modified event

The following figure shows the communication flow between the components that occurs when the observed state of a first-level automation resource changes that is referenced in the active policy changes.



This is a description of the scenario shown in the figure above:

1. The observed state of a resource which is referenced by a resource reference in the active end-to-end automation policy changes. In such a case, a so-called state change event is sent to the automation J2EE framework.
-
2. The automation J2EE framework has a list of all subscribers that must be informed when the state of this resource changes. In the scenario shown in the figure above, there are two subscribers for this resource:
 - the end-to-end automation domain has made a subscription (see to “Policy activation and subscription” on page 19)
 - an operator is monitoring this resource from the operations console

Therefore, the automation J2EE framework forwards the state change event to two recipients:

- a. The event is forwarded to the operations console

- b. Via the automation engine resource adapter, it is also forwarded to the automation engine
-

3. The event is forwarded

- a. to the operator monitoring the operations console
 - b. to the automation engine
-

4. The automation engine calculates the new states for the resource reference pointing to this resource and for all groups and related resources. In addition, as a reaction to the new situation, it may generate new requests.

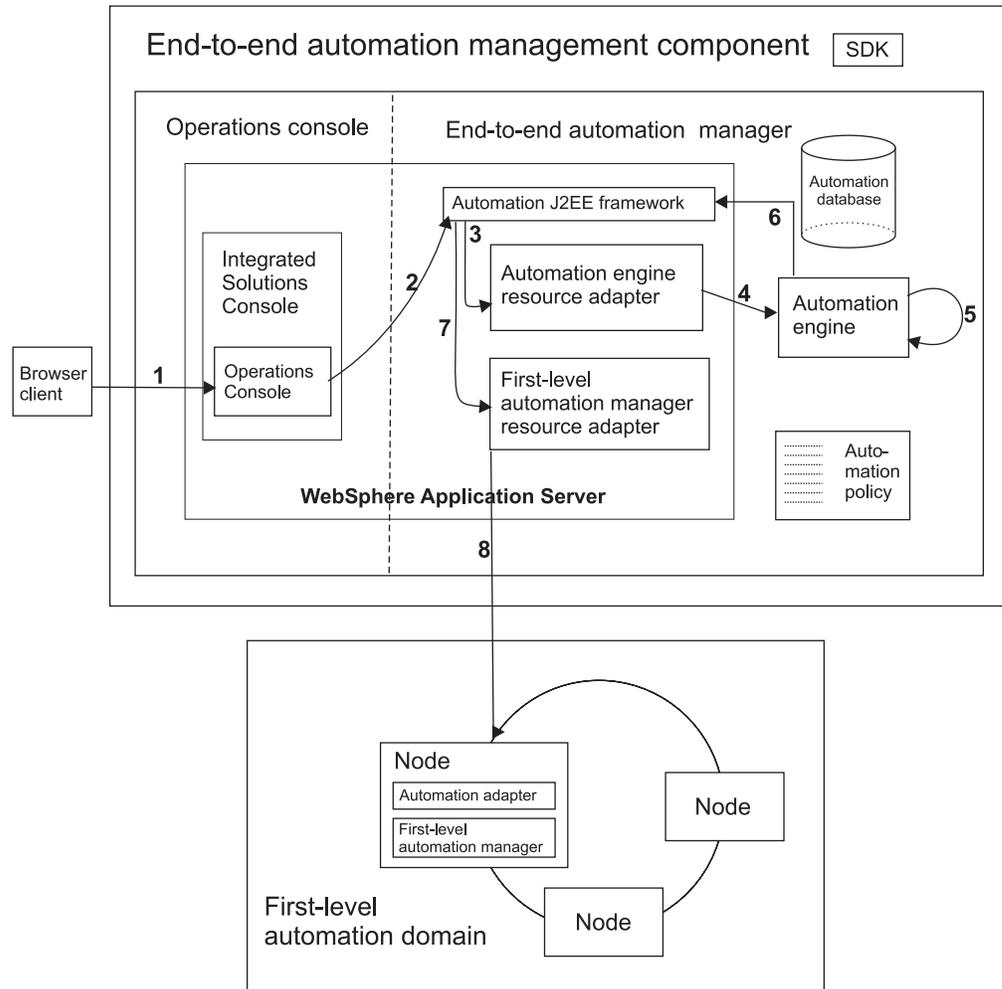
5. Each of the resulting requests is forwarded to the automation J2EE framework. The framework forwards each request to the first-level automation domain that hosts the resource to which the request applies.

6. The request is passed through the first-level automation manager resource adapter.

7. The request is transmitted to the first-level automation domain, which will evaluate the request and react accordingly.

An operator submits a request against a resource reference

The following figure shows the communication flow that occurs when an operator submits a request against a resource reference:



This is a description of the scenario shown in the figure above.

1. An operator submits a request against a resource reference from the operations console.

2. The operations console forwards the request to the automation J2EE framework.

3. The request is passed through the automation engine resource adapter.

4. The request is passed to the automation engine.

5. The automation engine calculates all resulting request which must be issued against referenced resources hosted by first-level automation domains. These calculations take into account all relationships defined in the active end-to-end automation policy.

-
6. All resulting requests against referenced resources are passed to the automation J2EE framework.

 7. The requests are passed through the first-level automation manager resource adapter.

 8. The requests are passed to the first-level automation domains. The first-level automation managers will handle the requests and start or stop the resources depending on the relationships defined in the active first-level automation policy.
-

The operations console is used in first-level automation mode

When you use the operations console in first-level automation mode, in which case you monitor and manage first-level automation domains only, you start the automation engine using the converter option `-co` (**eezdmn -co**). This will start the automation engine in "conversion-only" mode, that is, it will only be used to convert events but no end-to-end automation domain will be available and no end-to-end automation will be performed.

The following figure shows the communication flow that occurs when the automation engine is running in conversion-only mode and the operations console is used for monitoring and managing first-level automation domains only.

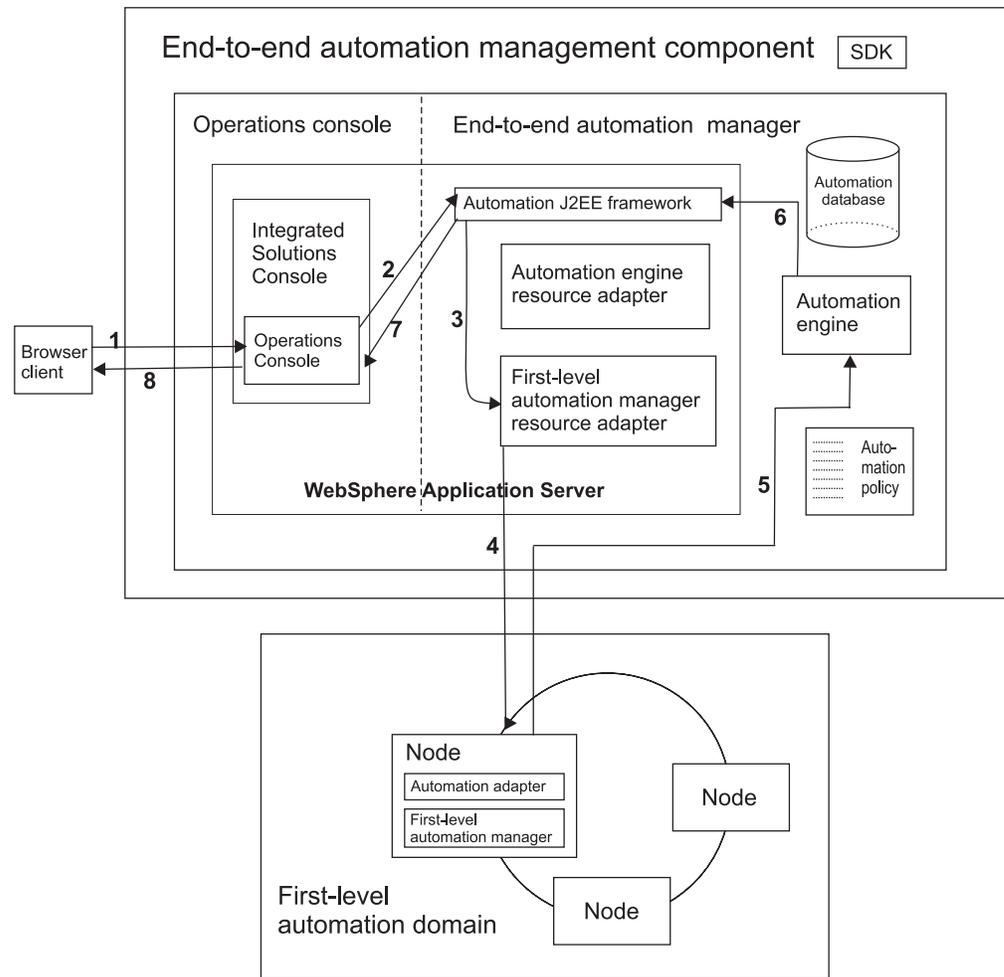


Figure 2. Operations console is used for managing first-level automation domains only

This is a description of the scenario shown in the figure above:

1. The operator opens the resource table for a first-level automation domain on the operation console.

2. The operations console performs a query for resource-related information (states and other information) against the automation J2EE framework. In addition, it also subscribes for this resource in order to be informed about future state changes.

3. The query and the subscription request are passed through the first-level automation manager resource adapter.

4. The query and the subscription request are passed to the first-level automation domain. The query results, that is, the current states of the resources, are returned to the operations console.

5. The observed state of the resource changes. Because the operations console subscribed for such events, a state change event is generated and passed to the automation engine.
-

6. Because the automation engine is running in conversion-only mode, it only translates the EIF event and puts it into the JMS queue that is used by the automation J2EE framework for getting posted about such events.

Note: The automation engine always converts events in this way. This is also true for the other scenarios described in this chapter, where this fact is not mentioned in order to keep the scenarios as simple as possible.

7. The change event is passed to the operations console because it is on the subscriber list.
-

8. The state of the displayed resources is updated accordingly
-

Chapter 5. Automation concepts

Resources of the end-to-end automation domain

The end-to-end automation manager manages the following types of resources:

- Resource references
- Resource groups
- Choice groups

Resource references

End-to-end automation resource references are virtual resources that reference actual resources. The actual resources are hosted by first-level automation domains.

Resource groups

End-to-end automation resource groups are composed of member resource references that are functionally related, share the same automation goal, and will be managed as one unit. Group members can be resource references, choice groups, or other resource groups, thus allowing an arbitrary level of nested groups.

Choice groups

End-to-end automation choice groups have the following characteristics:

- The members of a choice group are configuration alternatives that provide the same functionality (for example, two databases where one is used as production database and the other serves as backup).
- Only one of the members can be online at a time.
- The members can be either resource references or resource groups.
- One member of the choice group is defined as the *preferred member*. When the desired state of the choice group is online, the preferred member is kept online by the automation manager. The other members are kept offline.
- When a member other than the preferred member is to be brought online, an operator must change the preferred member.

Goal-driven automation

End-to-end automation is goal-driven. This means:

- The automation manager knows the automation goal for each resource it manages. The automation goal is the so-called desired state of the resource. Possible desired states for a resource are Online or Offline. The end-to-end automation manager pursues the automation goal by trying to keep the resource in its desired state.
- The automation manager is aware of relationships between resources that are defined in the end-to-end automation policy. It ensures that the relationships are fulfilled before a resource is started or stopped, that is, it ensures that any other resources that must be started or stopped first are actually started or stopped first.
- The automation manager pursues the automation goals not by issuing start or stop commands, but rather by submitting requests to the first-level automation managers that ask that the automation goal of the resource be changed. This ensures that a resource is only started or stopped when the first-level

automation manager has determined that any relationships defined for the resource in the first-level automation policy are fulfilled and no higher priority requests exist.

To ensure that each resource is kept in its desired state, the automation manager keeps track of various states for each resource. The following list gives a short overview of the states the automation manager knows for a resource and that are also displayed on the operations console:

Desired state

The desired state is the automation goal the automation manager pursues. Possible desired states are Online and Offline. When the desired state is online, the automation manager tries to keep the resource online. When the desired state is offline, the automation manager tries to keep the resource offline.

Compound state

The compound state indicates whether the resource or resource group works as desired or whether problems have occurred. It provides a traffic-light-like indicator informing operators when they need to react to a situation.

Operational state

The operational state provides additional information about the compound state.

Observed state

The observed state describes the current state of the actual first-level automation resource as reported by the first-level automation manager.

For a description of all states that are displayed in the operations console, refer to “State information provided on the operations console” on page 225.

How the automation manager is informed about automation goals

The automation manager is informed about the automation goal for a specific resource in the following ways:

- The default desired state for a resource is defined in the end-to-end automation policy.
- At runtime, the desired state is influenced by operator actions (start and stop requests) and by a resource’s relationships (StartAfter, StopAfter, and ForcedDownBy relationships):
 - Operators can change the desired state of a resource at runtime by submitting a start or stop request. If such a start or stop request can be fulfilled, the desired state of the resource changes to the new value. The new automation goal remains valid until the request is canceled or overruled by another request.
 - When the automation goal of a resource changes and the resource has StartAfter or StopAfter relationships, the desired states of the resources that are involved in the relationship change as well (if they are not in the requested desired state already). In such a case, the change of the desired state also persists until the original request is canceled or overruled by a higher priority request.
 - A ForcedDownBy relationship will result in a transient change of the automation goal when another resource is forced down.

How the default desired state is determined

The default desired state of any resource of the end-to-end automation domain (resource reference, resource group, and choice group) depends on the definition in the policy. The default desired state is the automation goal the automation manager will pursue if no other requests against the resource exist. The XML tag for defining the desired state in the XML policy is optional, this means that the default desired state can but need not necessarily be defined for each resource.

This is how the default desired state of a resource is determined:

- When the desired state of a resource reference is not defined in the policy and the resource reference is not a member of a resource group or choice group, the default (Online) is used.
- All members of a resource group have the same default desired state. The desired state of a resource group takes precedence over the desired state defined in the policy for any of its members. When the desired state is defined in the policy for a member of the group, it will be ignored even if it differs from the desired state of the group.
- When the desired state of a resource group is not defined in the policy, the default (Online) will be used.
- The default desired state of the members of a choice group depends on the default desired state of the choice group:
 - If the default desired state of the choice group is online, which is also the default that is used when the desired state is not defined in the policy, the automation manager will try to keep the so-called preferred member online and the other members offline.
 - If the default desired state is offline, all members will be kept offline.

Understanding relationships

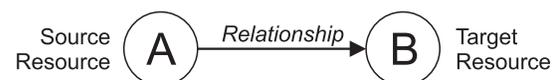
The end-to-end automation manager is aware of relationships between resources. Relationships are defined in the end-to-end automation policy. In end-to-end automation management, there are three types of relationships:

- StartAfter relationships
- StopAfter relationships
- ForcedDownBy relationships

What is a relationship?

Relationships can exist between two resource references, a resource reference and a group, and between two groups. The resources involved in a relationship can be hosted by different domains.

A relationship exists between a source resource and a target resource.



As the arrow in the figure above indicates, relationships always have a direction: In a StartAfter relationship, for example, target resource B would be started before source resource A.

By using combinations of managed relationships, complex automation scenarios can be defined. This is shown in this figure:



The arrows between the resources in the figure could, for example, represent the following three relationship definitions in the policy:

1. A StartAfter B
2. B StopAfter A
3. B StartAfter C

The source or target of a relationship can be resource references or groups of the end-to-end automation domain.

Whenever the automation goal of a resource is changed, for example, by a start or stop request, the automation manager checks whether StartAfter or StopAfter relationships are defined for the resource and, if this is the case, ensures that the relationships are fulfilled.

StartAfter relationship

The StartAfter relationship ensures that the source resource is only started when the target resource is online.

The StartAfter relationship provides the following behavior scheme:



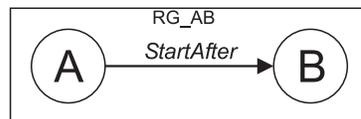
This StartAfter relationship defines the start sequence for resources A and B:

- When source resource A has to be started, then the target resource B is started first.
- After resource B has become online, resource A is started.

Details on the start behavior of the StartAfter relationship

The start behavior is controlled through the observed state of the target resource. At the time when the observed state of resource B has become online, resource A is started. Here are some examples for the start behavior that results from StartAfter relationships:

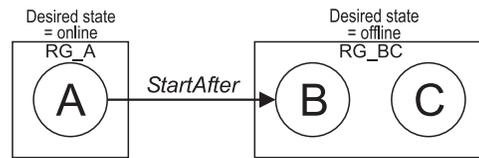
- In the example shown in the following figure, resource A and resource B are members of the same resource group:



When the desired state of their resource group is set to online, for example by a start request, both members A and B are started. Due to the StartAfter relationship from A to B, resource B is started first. Once the observed state of resource B is online, resource A is started.

- In the example shown in the following figure, resource A is a member of resource group RG_A, and resource B is a member of resource group RG_BC, and a StartAfter relationship is defined between A and B. Then the start

behavior of the StartAfter relationship is triggered when the desired state of RG_A is set to online, for example, by a start request.



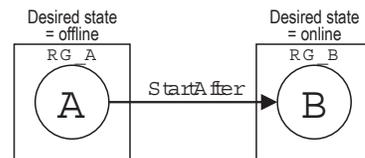
Due to the start sequence defined by the StartAfter relationship, resource B has to be started first. However, because RG_BC's desired state is set to offline, the following conflict exists:

RG_BC wants resource B to be offline whereas the StartAfter relationship forces B to be started. The end-to-end automation manager resolves this conflict in such a way that the online request is always more important than the offline request. Therefore resource B is started even though other possible group members of RG_BC will not be started since the desired state of their group is offline. After resource B is online, the end-to-end automation manager will try to start resource A. Resource C is not started.

When the desired state of RG_A is changed to offline in this scenario, resources A and B are stopped simultaneously. The reason for this behavior is that resource B was started due to the start request against resource group RG_A, which had been passed on to resource B due to the StartAfter relationship.

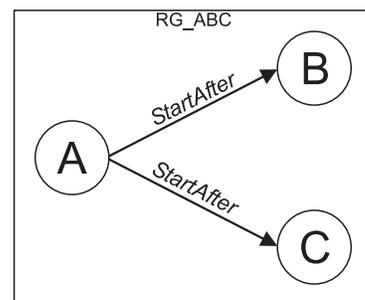
When the desired state of RG_A is set to offline, the start request for resource B is removed and the desired state of RG_B, which is offline, causes resource B to be stopped.

- The StartAfter relationship only acts in the forward direction of the relationship. In this example, resource A and resource B are members of different resource groups (A belongs to RG_A and B belongs to RG_B). In this case, setting the desired state of RG_B to online does not result in any action on resource A because resource B has no forward relationship to resource A.



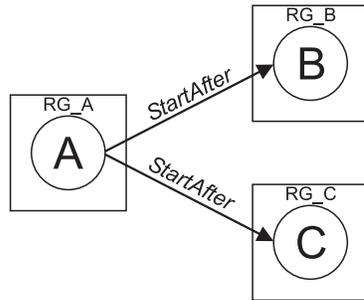
When RG_A's desired state is set to online, resource A can be started right away since resource B is already online.

- In this example, resource A has a StartAfter relationship to resource B and resource C.



In this case, starting A requires that both resources B and C are online before the end-to-end automation manager can start resource A. If A, B, and C are members of the resource group RG_ABC, setting the desired state of RG_ABC to online causes that resources B and C are started in parallel first. When the observed state of both resources is online, then resource A is started.

- In this example, resource A is a member of resource group RG_A, resource B is a member of resource group RG_B, and resource C is a member of resource group RG_C.



A has a StartAfter relationship to both B and C. Setting RG_A's desired state to online causes that due to the StartAfter relationship resource C and resource B are started. After both resources B and C are online, A is started.

StopAfter relationship

The StopAfter relationship ensures that the source resource can only be stopped when the target resource is offline.

The StopAfter relationship provides the following behavior scheme:



Resource A will not be stopped unless the target resource B has been brought offline before.

Details on the stop behavior of the StopAfter relationship

The stop behavior is controlled via the observed state of the target resource. At the time when the observed state of resource B has become offline, resource A is stopped. Here are some examples for the stop behavior that results from StopAfter relationships:

- This is an example of a simple StopAfter relationship. Source resource A cannot be stopped while target resource B is in observed state online.



When the desired state of resource A is set to offline, the automation manager stops B first. Once B is offline, A will be stopped.

- In this example, source resource A and target resource B are members of the same resource group.



When the desired state of resource group RG_AB is set to Online, both members A and B are started. Since the StopAfter relationship does not define a start sequence, resources A and B can be started simultaneously. Setting their resource group's desired state to offline causes that all members are stopped. Due to the

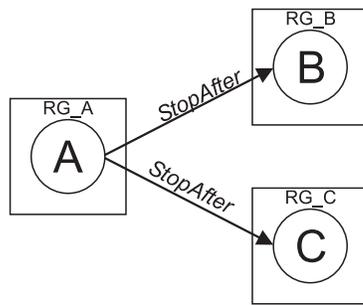
relationship from A to B, resource B is stopped first. When the observed state of resource B is offline, resource A is stopped.

- In this example, resources A and B are members of different resource groups (A belongs to RG_A, and B belongs to RG_B). RG_B has the desired state offline.



As long as the desired state of RG_B remains Offline, you can start and stop RG_A without any dependency to resource group RG_B. If you set the desired state of RG_B to online and the desired state of RG_A to offline, source resource A cannot stop as long as target resource B is Online. If the desired of RG_A is offline, you can start or stop RG_B without any dependency to resource A.

- In this example, resource A is a member of resource group RG_A, resource B is a member of resource group RG_B, and resource C is a member of resource group RG_C. A has a StopAfter relationship to both B and C.



If the desired state of RG_A is online and you want to stop it, RG_A cannot be stopped as long as the desired state of both RG_B and RG_C is online. Only when both RG_B and RG_C have a desired state of offline, resource A can be stopped.

ForcedDownBy relationship

Use the ForcedDownBy relationship to ensure that the source resource is brought down if the target resource comes offline.

The ForcedDownBy relationship provides the following behavior scheme:



Resource A is forced offline when either the target resource goes offline. The stop of resources A and B can happen in parallel. The force down of resource A will be triggered when resource B enters any of the regular down states (Offline) after having previously been in an Online state or when resource B fails while it is offline.

Note: After Resource A has stopped, its desired state will change to the current desired state again. For example, if Resource A has the desired state Online and is forced down because Resource B fails, the following happens:

1. Resource A is brought offline.

- When the observed state of Resource A has changed to Offline, its desired state changes to Online again and Resource A will be started.

Details on the force down behavior of the ForcedDownBy relationship

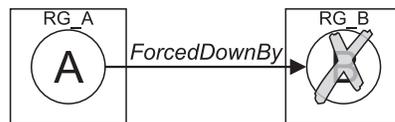
The basic principle of the ForcedDownBy relationship is that source resource A must be forced Offline when target resource B goes offline or fails. Here are some examples that illustrate the behavior when a ForcedDownBy relationship is defined:

- In this example, resource A has a ForcedDownBy relationship to resource B.



Both resources are online. In case resource B goes offline, resource A will be forced down.

- In this example, resource A is member of resource group RG_A, and resource B is member of resource group RG_B, and A has a ForcedDownBy relationship with resource B. The force down behavior of the ForcedDownBy relationship is triggered by a failure of resource B. Due to the ForcedDownBy relationship, resource A will be stopped as well. This will happen even though the desired state of RG_A is Online. However, because the desired state of RG_A is still online, resource A will be restarted by the end-to-end automation manager. To achieve the behavior that resource A remains offline as long as resource B is offline, add an additional StartAfter relationship between resource A and resource B.



How requests become goals

In end-to-end automation management, operators start and stop resources by submitting requests.

A request asks that one specific resource should be moved to a specific desired state (its automation goal). Using requests instead of commands ensures that the priority of requests is honored and that any relationships that have been defined for the resource are fulfilled before a resource is started or stopped.

Here is a simplified example that describes what happens when an operator submits a start request against a resource reference:

- The end-to-end automation manager checks whether a request has been submitted against the resource reference that has a higher priority than the current request. If this is not the case, the operator request wins and the desired state of the resource reference is set to online.
- The end-to-end automation manager checks whether StartAfter relationships are defined for the resource reference in the automation policy. When no such relationship exists, the automation manager sends a start request against the referenced resource to the first-level automation manager.
- The first-level automation manager checks whether requests against the resource exist that have a higher priority than the current request. If this is not the case, the first-level automation manager checks whether relationships have been

defined for the resource in the first-level automation policy that must be fulfilled before the resource can be started. When no such relationship is defined there, the first-level automation manager initiates the start of the resource.

This means that what happens after a start or stop request is submitted depends on the following conditions:

- whether the resource has StartAfter or StopAfter relationships.
- whether other higher priority requests exist for the resource itself or for a resource to which it has a relationship.

Requests processing when relationships exist

When a start or stop request is submitted against an end-to-end automation resource, the automation managers involved ensure that any relationships defined for the resource are fulfilled before the source resource is started or stopped. To achieve this goal, automation managers use two types of requests, namely, genuine requests and votes. Votes are a special type of request that have the following characteristics:

- Votes are internal requests that an automation manager generates against the target resource of a relationship.

To ensure that a relationship of a resource reference is fulfilled when a request is submitted against the source resource, the end-to-end automation manager will generate both a vote and a request:

- A vote is generated against the target resource reference.
- If the vote wins, that is, if no higher priority request against the target resource reference exists, the automation manager will generate a request against the referenced resource and forward it to the first-level automation manager.
- When a vote wins, the desired state of the target resource is changed accordingly. The new desired state persists until it is either overruled by a higher priority request or the request against the source resource is canceled.
- When the request against the resource is canceled, the votes that were generated against the target resources of a relationship are canceled as well.
- Operator requests can be canceled by any other operator from the operations console. Votes that were generated due to an operator request cannot be canceled directly. They are canceled automatically when the request against the source resource is canceled.

Request priorities

Requests that are submitted against a resource are kept in the resource's request list. Whether a request to change the desired state of a resource is successful, that is, if the request wins, depends on the priority of the requests that are already in the resource's request list. A request will only win if it has a higher priority than any of the other requests or votes in the list. The priority of a request depends on:

- whether the request was submitted by an operator or by an automation manager
- whether it is a start or stop request

In end-to-end automation management, the following priority rules apply when the request list contains conflicting requests:

- Requests have a higher priority than votes that were generated by the automation manager of the same automation domain.

- Requests generated by the end-to-end automation manager against a first-level automation resource have a lower priority than votes generated against the same resource by the first-level automation manager.
- Operator requests have a higher priority than requests generated by an automation manager.

Note: When an operator submits a request against a resource reference, resource group, or choice group, the request that is forwarded to the first-level automation manager is generated by the end-to-end automation manager. As requests that are generated by an automation manager have a lower priority than requests that are submitted by an operator, such a request will not win when the request list contains an operator request that was submitted directly against the first-level automation resource.

- Requests submitted by different operators have the same priority.
- Requests generated by any automation manager against the same resource have the same priority.
- Requests generated by the same automation manager replace each other.
- When conflicting requests have the same priority, the start request wins.

Request lists can be viewed on the operations console (refer to “Displaying request lists” on page 253).

How requests against resource references are processed

This chapter describes how requests against resource references are processed by the end-to-end automation manager.

As described above, resource references are virtual resources that are hosted by the end-to-end automation engine. Resource references point to actual resources that are hosted by first-level automation domains. The actual resources that are referenced by a resource reference are called referenced resources.

Requests against referenced resources are evaluated by the end-to-end automation manager and result from the following scenarios:

- An operator issues a request against a resource reference. After evaluating the request, the end-to-end automation manager forwards this request to one or more referenced resources.
- A state change event of a referenced resource causes the end-to-end automation manager to react by generating requests against one or more referenced resources.
- An operator activates an end-to-end automation policy. The end-to-end automation manager creates requests against all referenced resources to ensure that the desired state of the resource references defined in this policy is fulfilled.

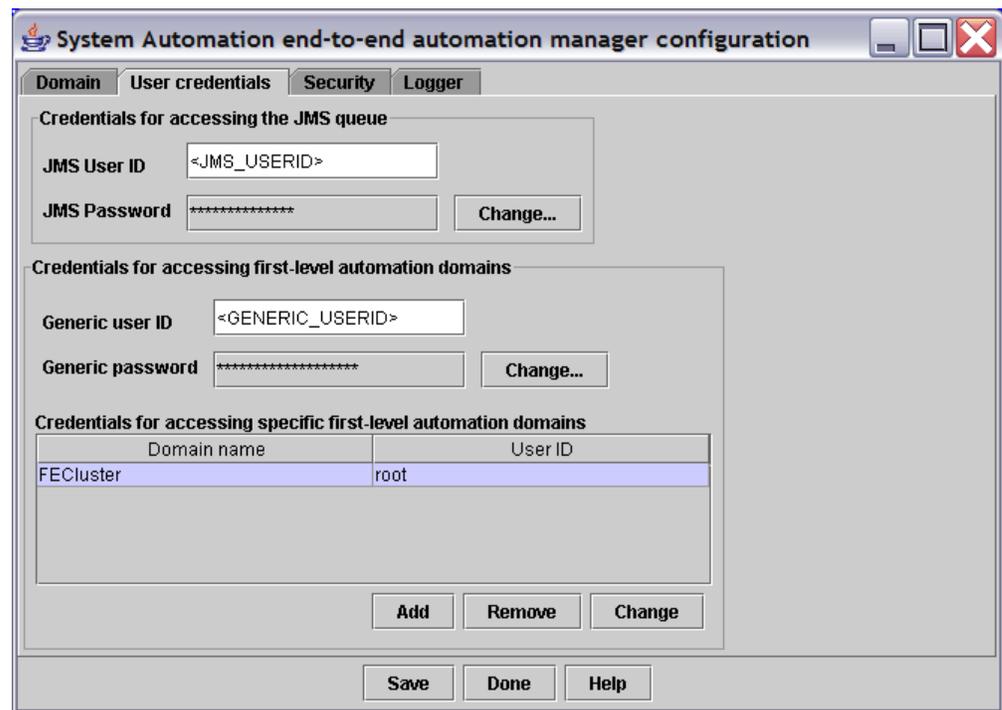
User credentials of the end-to-end automation manager

When the end-to-end automation manager issues requests against referenced resources, it must authenticate itself to the first-level automation domains that host the referenced resources. For authentication, the end-to-end automation manager uses the user credentials (user ID and password) that are specified on the User credentials page of the configuration dialog.

The user credentials are needed because the automation manager is a stand-alone process that must be able react to exceptional situations even if no operator is logged in.

If the referenced resource that is targeted by the request is hosted by a first-level automation domain for which specific user credentials have been specified, the automation manager uses these credentials for authentication. If no specific user credentials for the domain are specified in the configuration dialog, the automation manager uses the generic credentials that must be specified in the configuration dialog.

This is an example of how the user credentials for the automation engine are specified in the configuration dialog:



On the User credentials page shown above, specific credentials are only defined for the first-level automation domain FECluster. When the end-to-end automation manager issues requests against referenced resources that are hosted by FECluster, it uses the user ID root and the corresponding password.

When it issues requests against referenced resources that are hosted by other first-level automation domains, it uses the user ID and the password specified in the fields **Generic user ID** and **Generic password**.

Example scenarios

In the scenarios described in the following sections, it is assumed that the end-to-end automation policy contains the following specifications:



|
Resource A on
FEPLEX

|
Resource B on
MYDOMAIN

A policy is activated

When the policy containing the definitions above is activated, the automation engine first subscribes for the referenced resources Resource A, which is hosted by domain FEPLEX, and Resource B, which is hosted by domain MYDOMAIN (see also “Policy activation and subscription” on page 19). To make the subscriptions, the automation engine uses the user credentials specified in the configuration dialog (refer to Chapter 29, “Configuring the end-to-end automation manager,” on page 193).

After receiving the subscriptions, the automation managers on both first-level automation domains create a so-called initial resource event for each referenced resource and send them to the end-to-end automation manager. The initial resource events inform the end-to-end automation manager of the current observed state of Resource A and Resource B.

After receiving and processing these events, the end-to-end automation manager sets the states of both resource references (ResourceReference A and ResourceReference B) accordingly. Depending on which desired state is defined for the resource references in the end-to-end automation policy, the end-to-end automation manager generates requests and sends them to the referenced resources.

Note:

- After receiving the initial event for a resource, the end-to-end automation manager always generates a request against the referenced resource and sends it to the first-level automation domain. This is done even if the current observed state of the referenced resource already matches the desired state of the resource reference in the end-to-end automation policy. This ensures the desired state from the end-to-end automation policy is known on the first-level automation domain.
- The end-to-end automation manager writes a message to the domain log file that contains the user ID of the operator who activated the policy from the operations console.

An operator issues a request against a resource reference

An operator can issue requests against resource references from the operations console (see “An operator submits a request against a resource reference” on page 23 for a description of the complete flow). This request is passed to the end-to-end automation manager with the operator’s user ID. The end-to-end manager writes a message to the log file of the end-to-end automation domain. This message contains the user ID of the operator who issued this request from the operations console.

Subsequently, the end-to-end automation engine calculates the resulting actions. Assume that the operator with the user ID Charles issued a start request against ResourceReference A. The end-to-end automation manager will evaluate the new desired states of all resource references defined in the automation policy. In this particular case, also assume that ResourceReference B currently is in an offline state. As a startAfter relationship between ResourceReference A and ResourceReference B is defined in the policy, the first resulting action is to ensure that ResourceReference B is started, this results in an Online request against Resource B. The automation engine generates an Online request against Resource

B. This Online request is forwarded to the first-level automation domain MYDOMAIN with the credentials specified in the configuration dialog for this domain (in this case, with the user ID bob).

The request can now be viewed on the referenced resource Resource B. The request that has been added by the end-to-end automation manager has the source E2EMGR and the user ID that is specified for this domain in the configuration dialog (bob).

Subsequently, the end-to-end automation engine waits for the request to be processed by the first-level automation domain MYDOMAIN. After the end-to-end automation manager receives the resource status change event that informs it of the fact that Resource B has become online, the end-to-end automation engine generates the Online request against Resource A, which is hosted by FEPLEX, authenticating itself with the user ID root. This request can now be viewed on the referenced resource Resource A. The source of this request is E2EMGR. On Resource Reference A, the end-to-end operator request issued by Charles can also be viewed. On this level, however, the request source is OPERATOR, and the user ID is Charles.

To sum up: When an operator submits a request against a resource reference from the operations console, this may result in the generation of requests against more than one referenced resource. These resulting requests are issued by the end-to-end automation manager using the credentials from configuration dialog. The user ID of the operator who submits or cancels a request against a resource reference is logged in the log file of the end-to-end automation domain. It can also be viewed when the resource reference is selected.

The state of a referenced resource changes

Whenever the state of a referenced resource changes, the end-to-end automation manager is informed of the state change through an event. The state of the resource reference is updated accordingly. In some cases, the automation engine of the automation manager will create requests against this referenced resource or other referenced resources because of the state change. As described in the scenarios above, the end-to-end automation manager will use the user credentials specified in the configuration dialog when it issues the requests against the referenced resources.

When the end-to-end automation manager will not generate requests

The previous sections described the situations in which the end-to-end automation manager generates requests against referenced resources that are hosted by first-level automation domains. The following sections describe in which situations the end-to-end automation manager will not generate requests.

The referenced resource is a monitor resource

In some situations, a first-level automation manager is not able to handle requests against specific resources.

When the end-to-end automation manager or the operations console subscribes for events for such a resource, the initial resource event contains the information that the particular resource is a so-called monitor resource.

The end-to-end automation manager will never generate requests against such resources. Whenever a state change event is received from these resources, the specific state of the resource reference is only updated.

However, a state change of a monitor resource can still cause some other resource references to be started or stopped by a request that is generated by the automation manager. This happens if the resource reference referencing the monitor resource is a member of some relationship.

The referenced resource is in a transitional state

The end-to-end automation engine does not generate requests if the referenced resource is in a so-called transitional state. Transitional states are, for example, the states Starting or Stopping. The end-to-end manager waits until the transition is completed before generating a request.

The referenced resource is in a specific operational state

Some operational states of referenced resources also cause the end-to-end automation engine not to create requests. In general, it can be said that whenever the referenced resource is in a state where it cannot accept requests, the end-to-end automation engine will not create one.

In any state change event from a referenced resource, the first-level automation manager not only sends the current observed state but also the current operational state. If the operational state already indicates an error, the end-to-end automation manager assumes that the first-level automation manager already handles the current state of this referenced resource. The first-level automation manager already reacts to the particular situation. Therefore, it would not make sense for the end-to-end automation manager to also create a new request which might request the same operations as the first-level automation manager is already trying to perform.

The following list contains the operational state descriptions that will cause the end-to-end automation manager not to create requests:

- Warning: Waiting for initial state info
- Warning: Online/Offline request pending
- Warning: The communication has been interrupted
- Error: The resource has an unrecoverable problem
- Error: The hosting node is gone
- Error: The resource has been excluded from automation
- Error: The resource cannot be started/stopped because the online/offline request did not win at this moment
- Error: The resource reference references a resource that does not exist
- Error: The resource cannot be started/stopped because of unfulfilled dependencies
- Error: Unable to contact the referenced resource
- Error: The referenced resource is in an error state

Requests generated by the end-to-end automation manager are persistent

The end-to-end automation manager never cancels previously generated requests against referenced resources. If the desired state of a resource reference changes, for example, from Online to Offline, the end-to-end automation manager does not cancel the Online request against the referenced resource but generates an Offline request and sends it to the referenced resource.

The first-level automation domain handles this request by overwriting the previous request and processing the new request.

If the end-to-end automation manager fails and is restarted, the policy that was active at the time of failure is automatically activated again. The end-to-end automation manager again subscribes for the referenced resources and sends default requests to the referenced resources.

Canceling obsolete end-to-end automation manager requests on first-level automation resources

When an administrator deactivates the currently active end-to-end automation policy or activates a new one, the desired states of the resource references from the old policy that were propagated to the referenced first-level automation resources are retained as automation requests. This has the advantage that the referenced resources do not have to be restarted when the desired state in the old and new policy is identical.

However, the new policy may not contain references to the relevant first-level automation resources at all. In such a case, some of the requests that are retained in a first-level automation domain may be obsolete. The following sections describe how you can identify and delete such obsolete requests.

Canceling requests on SA for Multiplatforms resources

Perform the following steps to find and remove requests that were issued by the end-to-end automation manager against resources or resource groups hosted by SA for Multiplatforms:

1. To obtain a list of all resource groups against which a request has been issued, enter the following command:

```
lsrgreq -L
```

2. In the list, identify all resource groups with a request from source Automation
-

3. Cancel these request with the following command:

```
regreq -o cancel -s Automation <GROUPNAME>
```

4. To obtain a list of all group members against which a request has been issued, enter the following command:

```
lsrgreq -L -m
```

5. In the list, identify all resources with a request from source Automation
-

6. Cancel these request with the following command:

```
rgmbrreq -o cancel -s Automation <MEMBERNAME>
```

Example: The referenced resource is a first-level automation resource group

To list all requests against resource groups in a first-level automation domain, issue the following command:

lsrgreq -L

The following list is generated:

```
Displaying Resource Group request information:
All request information
ResourceGroup Priority Action Source NodeList ActiveStatus UserID ...
my_rg          high   start Automation {} Active e2e
```

The Active request is a relict from the old end-to-end automation policy. To remove the remaining request, enter the following command:

```
rgreq -o cancel -S Automation my_rg
```

Example: The referenced resource is a first-level automation resource

To list all requests that were issued directly against first-level automation resources, enter the following command:

lsrgreq -L -m

The following list is generated:

```
Displaying Member Resource request information:
All request information

Member Resource 1:
  Class:Resource:Node[ManagedResource] = IBM.Application:my_resource
  Priority                               = High
  Action                                 = start
  Source                                 = Automation
  ActiveStatus                           = Active
  UserID                                 = e2e
  Comments                               = 20050503142734+0200 |
```

The Active request is a relict from the old automation policy. To remove the obsolete request, enter the following command:

```
rgmbrreq -o cancel -S Automation IBM.Application:my_resource
```

Note: When the referenced resource is a SA for Multiplatforms fixed resource, the node name must be appended:

```
rgmbrreq -o cancel -S Automation IBM.Application:my_resource:node1
```

When the request has been removed, the observed state of my_resource changes from Online to Offline as defined in the first-level automation policy.

Canceling requests on SA z/OS resources

This is an example of a REXX script which can be used for the following purposes:

- Find all requests which have been issued by the end-to-end automation manager
- Cancel the requests that were found

```
/**/
Address NetVAsis,
'PIPE (NAME INGVOTE)',
  ' NETV INGVOTE,OUTMODE=LINE',
  ' DROP FIRST 3 LINES',
  ' DROP LAST 1 LINE',
  ' SEP',
  ' CASEI COLLECT BREAK BEFORE 27.5 /Req :/',
```

```

    '| CASEI LOC 27.12 /Org : E2EMGR/',
    '| EDIT 1.25 1 SKIPTO /:/ WORD 2.1 NW',
    '| FWDLINE 2 SKIPTO /:/ UPTO /( / 2.* NW',
    '| STEM data.'
Do i = 1 To data.0
  Parse Var data.i name type system . 27 request source .
  resource = Strip(name/'type'/'system','T','/')
  say,
  'INGSET KILL' resource' REQUEST='request 'SOURCE='source 'VERIFY=NO'
End i

```


Part 2. Installation

| | |
|--|-----|
| Chapter 6. Planning for installation. | 47 |
| Packaging | 47 |
| Product CD and WebSphere Application Server | |
| Upgrade CD | 47 |
| Product CD | 47 |
| WebSphere Application Server upgrade CD | 48 |
| Electronic distribution | 48 |
| Archives | 48 |
| Product features, DB2 setup options, and user registry options | 50 |
| DB2 setup options and user registry options | 51 |
| Using a local or remote DB2 setup. | 51 |
| Using LDAP or DB2 as the user registry. | 51 |
| Product requirements | 52 |
| Supported operating systems | 52 |
| Middleware software requirements | 53 |
| Software prerequisites for a local DB2 setup | 54 |
| Software prerequisites for a remote DB2 setup | 54 |
| Software prerequisites: LDAP is used as the user registry | 54 |
| Software prerequisites: DB2 is used as the user registry | 54 |
| Browser requirements | 54 |
| Hardware requirements | 55 |
| Memory | 55 |
| TCP/IP connectivity | 55 |
| Disk space requirements | 55 |
| Security concepts | 57 |
| Security considerations | 57 |
| Considerations for choosing between LDAP and DB2 as the user registry | 57 |
| User management | 57 |
| Chapter 7. Installing the middleware software | 59 |
| What the middleware software CDs contain | 59 |
| Installing a DB2 server | 59 |
| DB2 server requirements | 59 |
| DB2 server installation. | 60 |
| Post-installation tasks for remote DB2 setup | 60 |
| Creating the automation manager database and the database tables | 61 |
| Creating the operations console database | 61 |
| Installing a DB2 client | 62 |
| DB2 client requirements | 62 |
| DB2 client installation | 62 |
| Post-installation tasks for remote DB2 setup | 63 |
| Cataloging a TCP node | 63 |
| Cataloging the operations console database. | 63 |
| Cataloging the automation manager database | 64 |
| Installing WebSphere Application Server. | 64 |
| WebSphere Application Server 6.0.0.0 requirements | 64 |
| Installing WebSphere Application Server 6.0.0.0 | 65 |
| Installing Refresh Pack 2 and the required Interim Fixes | 65 |
| Installing Refresh Pack 2 | 65 |
| Installing the required Interim Fixes | 65 |
| Post-installation tasks | 66 |
| Chapter 8. Setting up an LDAP server. | 67 |
| Required LDAP directory tree structure | 67 |
| Required user groups and users | 68 |
| LDAP-related pre-installation tasks | 68 |
| Sample LDAP configuration | 69 |
| Chapter 9. Preparing for the installation of the end-to-end automation management component | 71 |
| Collecting the information you need to provide during installation | 71 |
| Installation directory and Tivoli Common Directory | 71 |
| Installation parameters for DB2. | 72 |
| Installation parameters for WebSphere Application Server | 74 |
| Installation parameters for LDAP | 74 |
| Installation parameters for the operations console | 76 |
| Installation parameters for IBM Tivoli Enterprise Console. | 79 |
| Name of the end-to-end automation domain | 80 |
| What the installation CD contains | 80 |
| Languages supported by IBM Tivoli System Automation | 81 |
| Installation prerequisites | 81 |
| Chapter 10. Installing the end-to-end automation management component | 83 |
| Chapter 11. Verifying the installation | 99 |
| Automation manager | 99 |
| End-to-end automation database | 99 |
| Automation J2EE Framework | 99 |
| Verifying that DB2 accepts WebSphere Application Server requests | 99 |
| Automation engine | 100 |
| Operations console | 100 |
| Chapter 12. Post-installations tasks | 103 |
| Setting up SSL for the operations console | 103 |
| Modifying the LTPA settings | 105 |
| Modifying the HTTP session timeout | 106 |
| Configuring how many users can connect to the automation manager using the operations console | 107 |
| Chapter 13. Uninstalling the end-to-end automation management component | 109 |
| Launching the graphical uninstallation program on Windows. | 109 |
| Launching the graphical uninstallation program on AIX and Linux | 109 |
| Using the uninstallation program. | 109 |

| | |
|--|---------|
| Chapter 14. Installing service | 113 |
| Where to obtain fix packs | 113 |
| Archive naming conventions | 113 |
| Naming conventions of the update installer location | 114 |
| Usage instructions for the platform-specific archives | 114 |
| Windows | 114 |
| AIX | 114 |
| Linux on IBM x/Series | 115 |
| PPC Linux | 115 |
| Linux on z/Series | 115 |
| Steps for installing a product fix pack | 116 |
| Chapter 15. Uninstalling service | 117 |
| Chapter 16. Troubleshooting installation problems | 119 |
| Installation wizard cannot find WebSphere Application Server on the system | 119 |
| The installation of the operations console fails | 119 |
| Using the installation log files | 119 |
| Installation log file directories | 119 |
| Installation log files | 120 |
| Pre-installation phase log files | 121 |
| Installation phase log files | 121 |
| Procedures for troubleshooting an installation | 123 |
| Gathering information for IBM Support | 124 |

Chapter 6. Planning for installation

This chapter contains the information you need for preparing the installation of the end-to-end automation management component.

Packaging

The end-to-end automation management component can be ordered from IBM as media pack or downloaded from an IBM software distribution download site.

Product CD and WebSphere Application Server Upgrade CD

When you order the end-to-end automation management component on CD, you receive the following CDs:

- One product CD for each operating system on which the product can be installed. You use the product CD to install the end-to-end automation management component.
- One WebSphere Application Server Upgrade CD for each operating system on which the product can be installed. You use the upgrade CD to bring WebSphere Application Server to the version level required for the end-to-end automation management component.

For information about the middleware software CDs that are shipped with the end-to-end automation management component, refer to “What the middleware software CDs contain” on page 59.

Product CD

The following table lists the versions of the product CDs that are available for the end-to-end automation management component. To install the product, you use the installation wizard file listed in the right column of the table.

Table 3. Product CD versions

| Operating system | Product CD label | Installation wizard file |
|------------------------|---|-------------------------------------|
| Windows | IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component for Windows | EEZ2110E2EWindows/Windows/setup.exe |
| AIX | IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component for AIX | EEZ2110E2EAIX/AIX/setup |
| Linux for IBM x-Series | IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component for Linux xSeries | EEZ2110E2EI386/i386/setup |
| Linux PPC | IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component for Linux PPC | EEZ2110E2EPPC/ppc/setup |
| Linux for IBM z-Series | IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component for Linux zSeries | EEZ2110E2ES390/s390/setup |

WebSphere Application Server upgrade CD

The following table lists the available versions of the WebSphere Application Server upgrade CDs.

Table 4. WebSphere Application Server upgrade CD versions

| Operating system | CD label |
|------------------------|--|
| Windows | IBM Tivoli System Automation Multiplatform V2.1.1 WAS 6.0 upgrade for Windows |
| AIX | IBM Tivoli System Automation Multiplatform V2.1.1 WAS 6.0 upgrade for AIX |
| Linux for IBM x/Series | IBM Tivoli System Automation Multiplatform V2.1.1 WAS 6.0 upgrade for Linux xSeries |
| Linux PPC | IBM Tivoli System Automation Multiplatform V2.1.1 WAS 6.0 upgrade for Linux PPC |
| Linux for IBM z/Series | IBM Tivoli System Automation Multiplatform V2.1.1 WAS 6.0 upgrade for Linux zSeries |

Electronic distribution

You can also obtain the end-to-end automation management component through electronic distribution. In this case, you can download the deliverables from a URL you receive after purchasing the product.

The following tables lists the archives that you need to download for each platform to install the WebSphere Application Server upgrade, which is required for running the end-to-end automation management component, and the product itself.

Archives

The following tables list the archives that you need for installing the WebSphere Application Server upgrade and the product itself.

Windows:

Table 5. Archives for Windows platforms

| Archive name | Description |
|--------------|--|
| C892VML.exe | This is the archive you use to install the product. The archive is self-extracting. After extraction, the directory structure is identical to that on the corresponding CD. |
| C8931ML.exe | The self-extracting archive contains the files that you need for installing the WebSphere Application Server upgrade. After extraction, the directory structure is identical to that on the corresponding CD. |

AIX:

Table 6. Archives for AIX platforms

| Archive name | Description |
|--------------|--|
| C892WML.bin | This is the archive you use to install the product. The archive is self-extracting. After extraction, the directory structure is identical to that on the corresponding CD. |
| C8932ML.bin | The self-extracting archive contains the files that you need for installing the WebSphere Application Server upgrade. After extraction, the directory structure is identical to that on the corresponding CD. |

Linux on IBM x/Series:

Table 7. Archives for Linux on IBM x/Series

| Archive name | Description |
|--------------|--|
| C892XML.tar | This is the archive you use to install the product. For extracting the archive, GNU tar 1.13 or later is required. Use the tar -xf command to extract the files to a temporary directory. After extraction, the directory structure is identical to that on the corresponding CD. |
| C8933ML.tar | The archive contains the files that you need for installing the WebSphere Application Server upgrade. After extraction, the directory structure is identical to that on the corresponding CD. |

PPC Linux:

Table 8. Archives for PPC Linux

| Archive name | Description |
|--------------|--|
| C892YML.tar | This is the archive you use to install the product. For extracting the archive, GNU tar 1.13 or later is required. Use the tar -xf command to extract the files to a temporary directory. After extraction, the directory structure is identical to that on the corresponding CD. |
| C8935ML.tar | The archive contains the files that you need for installing the WebSphere Application Server upgrade. After extraction, the directory structure is identical to that on the corresponding CD. |

Linux on z/Series:

Table 9. Archives for Linux on z/Series

| Archive name | Description |
|--------------|---|
| C892ZML.tar | <p>This is the archive you use to install the product.</p> <p>For extracting the archive, GNU tar 1.13 or later is required.</p> <p>Use the tar -xf command to extract the files to a temporary directory.</p> <p>After extraction, the directory structure is identical to that on the corresponding CD.</p> |
| C8934ML.tar | <p>The archive contains the files that you need for installing the WebSphere Application Server upgrade.</p> <p>After extraction, the directory structure is identical to that on the corresponding CD.</p> |

Product features, DB2 setup options, and user registry options

The two major subcomponents of the end-to-end automation management component, namely, the end-to-end automation manager and the operations console are installed on the same node and run in the same WebSphere Application Server environment:

- During installation, the automation J2EE framework and the resource adapters are installed to an existing server in WebSphere Application Server. Typically, the name of this server is server1.
- The installation of the operations console creates a new server in WebSphere Application Server. The name of this server is ISC_Portal.

The following figure depicts the setup of the end-to-end automation management component.

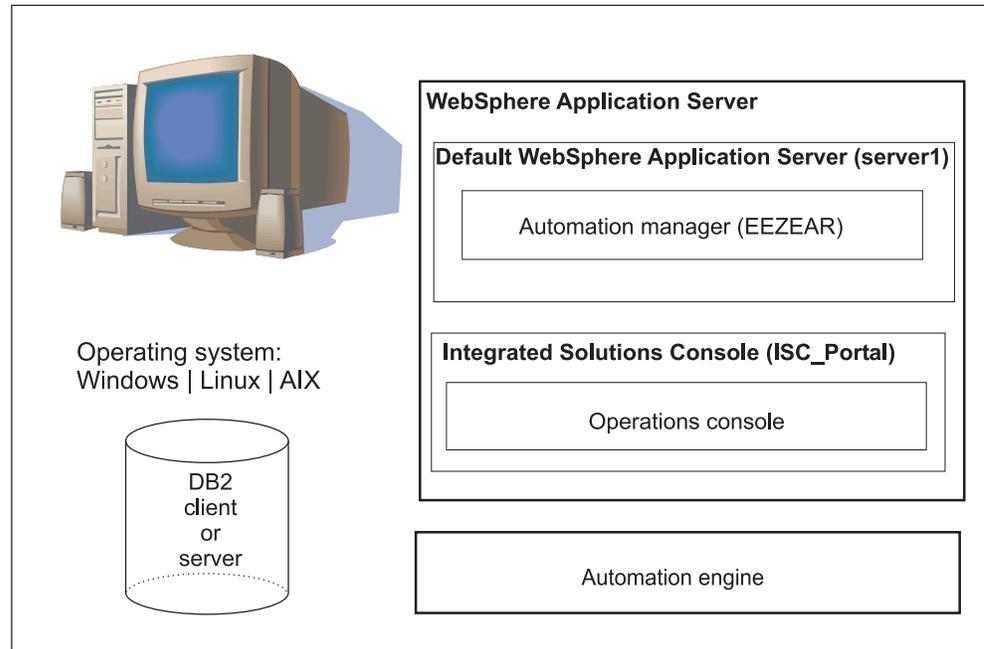


Figure 3. Setup of the end-to-end automation management component

For more information about the components of the end-to-end automation management component, refer to Chapter 2, “Components of end-to-end automation management,” on page 9.

DB2 setup options and user registry options

When planning for the installation of the end-to-end automation management component, you must decide:

- whether you want to use a local or remote DB2 setup for the two DB2 databases that are required for the end-to-end automation management component.
- whether you want to use DB2 or LDAP as the user registry.

Using a local or remote DB2 setup

In a local DB2 setup, the DB2 server is installed and runs on the same node on which the end-to-end automation management component is installed.

In a remote DB2 setup, the DB2 server is installed and runs on a node other than that on which the end-to-end automation management component is installed. In this case, you need to install a DB2 client on the end-to-end automation management node.

Using LDAP or DB2 as the user registry

When using LDAP, the user and group information needed by the operations console for authentication is stored in an LDAP directory.

When using DB2 as the user registry, the user and group information needed by the operations console for authentication is stored in a DB2 database.

For more information, refer to “Security concepts” on page 57.

Product requirements

The following sections list the software and hardware requirements for the end-to-end management component.

Supported operating systems

The following table lists the operating systems that are supported by the end-to-end automation management component:

Table 10. Supported operating systems

| Operating system | IBM x/Series ¹ | IBM i/Series | IBM p/Series | IBM z/Series |
|---|---------------------------|----------------|----------------|----------------|
| Windows Server 2003 Standard Edition (32 bit) | X | | | |
| Windows Server 2003 Enterprise Edition (32 bit) | X | | | |
| AIX 5.2 (32 bit) (AIX 5L Version 5.2) ML 5 | | | X | |
| AIX 5.2 (64 bit) (AIX 5L Version 5.2) ML 5 | | | X | |
| AIX 5.3 (32 bit) (AIX 5L Version 5.3) ML 2 ⁹ | | | X | |
| AIX 5.3 (64 bit) (AIX 5L Version 5.3) ML 2 ⁹ | | | X | |
| SUSE SLES 8 SP 3 (31 bit/32 bit ^{2, 3}) | X | | | X |
| SUSE SLES 8 SP 3 (64 bit ⁴) | | X | X | |
| SUSE SLES 9 (32 bit ³) | X | | | |
| SUSE SLES 9 (64 bit ⁴) | | X | X | X ⁸ |
| Red Hat RHEL 3.0 AS QU 3 (31 bit/32 bit ^{2, 3}) | X ⁵ | | | X ⁷ |
| Red Hat RHEL 3.0 AS QU 3 (64 bit ⁴) | | X ⁶ | X ⁶ | |

Table 10. Supported operating systems (continued)

| Operating system | IBM x/Series ¹ | IBM i/Series | IBM p/Series | IBM z/Series |
|--|---------------------------|--------------|--------------|----------------|
| Red Hat RHEL 4.0 AS (32 bit ³) | X | | | |
| Red Hat RHEL 4.0 AS (64 bit ⁴) | | X | X | X ⁸ |

Notes:

1. IBM x/Series systems with IA32, EM64T, or AMD64 architecture. Any other systems with IA32, EM64T, or AMD64 architecture are also supported. Systems with IA64 architecture are not supported.
2. The following Linux kernel architectures are supported for running with 31 bit:
 - s390 on IBM z/Series
3. The following Linux kernel architectures are supported for running with 32 bit:
 - x86 on IBM x/Series
4. The following Linux kernel architectures are supported for running with 64 bit:
 - ppc64 on IBM i/Series and IBM p/Series
 - s390x on IBM z/Series is supported for some distributions
5. Red Hat RHEL 3.0 AS QU3 with x86 kernel architecture requires the following packages:
 - compat-gcc-7.3-2.96.122
 - compat-libstdc++-7.3-2.96.122
 - compat-libstdc++-devel-7.3-2.96.122
 - compat-glibc-7.x-2.2.4.32.5
 - compat-gcc-c++-7.3-2.96.122
 - compat-db-4.0.14-5
 - rpm-build-4.2.1-4.2
6. Red Hat RHEL 3.0 AS QU3 with ppc64 kernel architecture requires the following package:
 - rpm-build-4.2.1-4.2
7. Red Hat RHEL 3.0 AS QU3 with s390 kernel architecture requires the following packages:
 - compat-libstdc++-7.2.2-2.95.3.77
 - compat-db-4.0.14-5
 - rpm-build-4.2.1-4.2
 - compat-pwdb-0.62-3
8. SUSE SLES 9 and RHEL 4.0 AS on s390x kernel require IBM DB2 UDB Version 8.2 Run-Time Client with Fix Pack 10 running as 31 bit application. This precludes IBM DB2 UDB Version 8.2 server from running on the same system, that is, remote DB2 setup is required.
9. APAR IY65979 must be installed.

Middleware software requirements

Before you can install the end-to-end automation management component, the following two software prerequisites must be manually installed on the system on which the end-to-end automation management component will run:

- A DB2 server for a local DB2 setup or a DB2 client for a remote DB2 setup
- IBM WebSphere Application Server 6.0.2 (with particular Interim Fixes)

The automation manager and the operations console share the same WebSphere Application Server environment (see Figure 3 on page 51).

Software prerequisites for a local DB2 setup

Before you can install the automation manager and the operations console, the following software prerequisite must be manually installed on the system on which the end-to-end automation management component will run:

- IBM DB2 UDB Version 8.2.3 (equivalent to Version 8.1.10) Enterprise Server Edition

You can either directly install this level of DB2 UDB or install a lower level of IBM DB2 UDB Version 8 and apply Fix Pack 10.

Software prerequisites for a remote DB2 setup

Before you can install the automation manager and the operations console, the following software prerequisites must be manually installed:

- An IBM DB2 UDB server Version 8.2.3 (equivalent to Version 8.1.10) must be installed on a system other than that on which the end-to-end automation management component will run.

You can either directly install this level of DB2 UDB or install a lower level of IBM DB2 UDB Version 8 and apply Fix Pack 10.

- Additionally, a DB2 client must be installed on the system on which the end-to-end automation management component will run. The following DB2 clients are supported:

- IBM UDB DB2 Run-Time Client Version 8.2.3 (equivalent to Version 8.1.10; available for all operating systems)

You can either directly install this level of DB2 UDB or install a lower level of IBM DB2 UDB Version 8 and apply Fix Pack 10.

- IBM UDB DB2 Run-Time Client Lite Version 8.2.3 (equivalent to Version 8.1.10; only available for Windows)

You can either directly install this level of DB2 UDB or install a lower level of IBM DB2 UDB Version 8 and apply Fix Pack 10.

Software prerequisites: LDAP is used as the user registry

When using LDAP as the user registry, one of the following LDAP servers must be running on an arbitrary system and must be reachable from the system on which the end-to-end automation management component will run:

- IBM Tivoli Directory Server Version 5.2
- IBM Directory Server Version 5.1

More information on IBM Tivoli Directory Server can be found at

www.ibm.com/software/tivoli/products/directory-server/

Note: IBM Tivoli Directory Server is not contained on the CDs that are shipped with the end-to-end automation management component. If you want to use LDAP, you must provide one of the LDAP servers listed above.

Software prerequisites: DB2 is used as the user registry

When using DB2 as the user registry, you do not need to install additional software manually.

Browser requirements

The operations console is the user interface of the end-to-end automation management component. It is displayed in a Web browser that connects to the WebSphere Application Server on which the operations console is running. The Web browser may run on an arbitrary system.

The following Web browsers are supported:

- Microsoft Internet Explorer 6.x
- Mozilla 1.7
- Netscape 7

For information on how the Web browser must be configured, refer to “Configuring your Web browser” on page 209.

Hardware requirements

The following sections describe the hardware requirements for the end-to-end automation management component. The description does not cover the hardware requirements that need to be satisfied for installing and running the required middleware software. For more information on hardware requirements for the required middleware software, refer to Chapter 7, “Installing the middleware software,” on page 59.

Memory

Memory: 1.5 GB is required on the server on which the end-to-end automation management component is installed (for WebSphere Application Server, Integrated Solutions Console, and the end-to-end automation management component)

TCP/IP connectivity

The end-to-end automation management component consists of various components that may run on one or several systems. For example:

- When you use a local DB2 setup, the automation manager, the operations console, and the DB2 server run on the same system (single-node setup).
- When you use a remote DB2 setup, the end-to-end automation management component and the DB2 server run on separate systems (multi-node setup).

Be sure that TCP/IP connections can be established between the following components:

- the WebSphere Application Server that is running the automation J2EE framework, the resource adapters, and the operations console
- the DB2 server
- the LDAP server (if LDAP security setup is used)

Disk space requirements

Disk space requirements on Windows systems: The following table lists the disk space requirements for the end-to-end automation management component on Windows systems. Note that the table does not include the space required for the installation of the middleware software.

Table 11. Disk space requirements on Windows systems

| Description | Default directory | Disk space |
|---|--------------------------------|------------|
| End-to-end automation management component installation directory | C:\Program Files\IBM\tsamp\eez | 70 MB |

Table 11. Disk space requirements on Windows systems (continued)

| Description | Default directory | Disk space |
|--|---|------------|
| Automation manager and operations console deployed in WebSphere Application Server | C:\Program Files\IBM\WebSphere\AppServer | 60 MB |
| Operations console installation directory | C:\Program Files\IBM\ISC | 600 MB |
| DB2 database | C:\DB2 | 120 MB |
| Installation log and response files | The value of the system variable %TEMP% Typically, this is: C:\Documents and Settings\Administrator\Local Settings\Temp | 75 MB |
| Temporary disk space needed for installation | The value of the system variable %TEMP% Typically, this is: C:\Documents and Settings\Administrator\Local Settings\Temp | 100 MB |
| Configuration file directory and policy pool directory of the end-to-end automation management component | C:\Program Files\IBM\tsamp\eez\cfg C:\Program Files\IBM\tsamp\eez\policyPool | 1 MB |
| Tivoli Common Directory | C:\Program Files\IBM\tivoli\common\eez | 250 MB |
| Installer registry | C:\Windows\vpd.properties | 10 KB |

Disk space requirements on AIX and Linux systems: The following table lists the disk space requirements for the end-to-end automation management component on AIX and Linux systems. Note that the table does not include the space required for the installation of the middleware software.

Table 12. Disk space requirements on AIX and Linux systems

| Description | Default directory | Disk space |
|--|--|------------|
| Installation directory of the end-to-end automation management component | /opt/IBM/tsamp/eez | 70 MB |
| Automation manager and operations console deployed in WebSphere Application Server | AIX: /usr/IBM/WebSphere/AppServer Linux: /opt/IBM/WebSphere/AppServer | 60 MB |
| Operations console installation directory | /opt/IBM/ISC | 500 MB |
| DB2 database | ~db2inst1 | 120 MB |
| Installation log and response files | /tmp | 75 MB |
| Temporary disk space needed for the installation | /tmp | 100 MB |
| Configuration file directory and policy pool directory of the end-to-end automation management component | /etc/opt/IBM/tsamp/eez/cfg /etc/opt/IBM/tsamp/eez/policyPool | 1 MB |
| Tivoli Common Directory | /var/ibm/tivoli/common/eez | 250 MB |
| Installer registry | ~root/vpd.properties | 10 KB |

Security concepts

The following sections describe the security concepts for end-to-end automation management.

Security considerations

If you are using a local DB2 setup, the end-to-end automation management component and the external components needed for running it, namely, WebSphere Application Server, Integrated Solutions Console, and DB2, are installed on the same system and you do not have to secure the connections between these components over SSL.

However, external connections will be established between the components listed below. If possible, these connections should be secured with SSL. This is recommended when the external components are running in different security domains, separated by firewalls.

External connections will be established between the following components (the port numbers given in brackets are default values):

- The connection between the automation engine and the automation adapters (port 2001).
- The connection between the automation adapters and the automation engine (port 2002). Note that SSL is not supported for this connection.
- When LDAP security is used: The connection between WebSphere Application Server and the LDAP server (port 389; for connections secured over SSL, port 636 is used)
- The connection between the Web browser in which the operations console is displayed and Integrated Solutions Console (HTTP port 8421, HTTPS port 8422)

Considerations for choosing between LDAP and DB2 as the user registry

WebSphere Application Server can be configured to use either an LDAP directory or a DB2 database for storing user information and for authenticating the users of the end-to-end automation management component. Regardless of which type of configuration you choose, you can use the user management function of Integrated Solutions Console to manage the user IDs and passwords for end-to-end automation management.

The following considerations may help you to decide which type of user registry to choose:

- If you are already using an LDAP server for authentication or are planning to set up an LDAP server anyway, you should consider using this LDAP server for end-to-end automation management as well. Note that only LDAP servers with write support can be used for end-to-end automation management. Read-only LDAP servers are not supported.
- If you are currently not using an LDAP server, you can use DB2 for storing user information and for authenticating users.

User management

For information on how user IDs and passwords for end-to-end automation management are managed, refer to Chapter 24, “Managing users,” on page 153.

Chapter 7. Installing the middleware software

Depending on the setup type you choose, middleware software has to be installed on one or more systems before the end-to-end automation management component can be installed.

- For information on possible setup types, refer to “DB2 setup options and user registry options” on page 51
- For information on the required middleware software for each system, refer to “Middleware software requirements” on page 53

What the middleware software CDs contain

The middleware software CDs that are shipped with the end-to-end automation management product CDs contain the following software products:

- IBM DB2 UDB Version 8.2 Enterprise Server Edition (DB2 server)
- IBM UDB DB2 Run-Time Client Version 8.2 (DB2 client)
- IBM UDB DB2 Run-Time Client Lite Version 8.2 (DB2 client)
- IBM WebSphere Application Server Base Version 6.0 (WAS)

Notes:

1. In addition to a WebSphere Application Server 6.0 CD for every supported platform, a WAS 6.0.0 upgrade CD is available for every supported platform. The upgrade CD is needed for bringing WebSphere Application Server to the required product level.
2. Please note that the IBM Tivoli Directory Server is not contained on the middleware software CDs.

Installing a DB2 server

DB2 server requirements

Check which requirements need to be met for installing and running a DB2 server. The information can be found in the following publications:

- IBM DB2 Universal Database - Quick Beginnings for DB2 Servers - Version 8.2 (GC09-4836)
- IBM DB2 Universal Database - Release Notes - Version 8

The latest versions of these publications can be found on the IBM DB2 UDB Web site at

www.ibm.com/software/data/db2/udb/support/

You find the link to the PDF manuals in the **Other resources** section on the Web page.

In addition, check for the latest system requirements at

www.ibm.com/software/data/db2/udb/sysreqs.html

The DB2 release notes can also be found on the CD labeled *IBM DB2 Universal Database Enterprise Server Edition Version 8.2* for your platform. Make sure that all

requirements for installing and running a DB2 server are met. Otherwise, the end-to-end automation management component may not install or work properly.

DB2 server installation

You can use the DB2 Setup wizard to install the DB2 server. You find the DB2 Setup wizard on the CD labeled *IBM DB2 Universal Database Enterprise Server Edition Version 8.2* for your platform.

The typical installation of a single-partition database environment is recommended.

On a Windows system, the DB2 administration server user, the fenced user, and the instance owner user must be local users.

You need to create a DB2 instance as a prerequisite for installing the end-to-end automation management component. If you are asked whether the DB2 instance is to be installed with 32-bit or 64-bit instance word size, choose 32-bit.

Make sure that the DB2 server has the required version level (refer to “Middleware software requirements” on page 53).

When you install the end-to-end automation management component, you will need the following information:

- The host name of the system where the DB2 server is installed.
- The port number of the DB2 instance
The port number is displayed on the summary panel of the DB2 Setup wizard. The summary panel appears immediately before the wizard copies the program files.
- The directory to which the DB2 server is installed if a local DB2 setup is used
- The name and password of the instance owner user or of a different user who is authorized to drop and create databases and database tables, and to select, insert, delete, and update rows in tables.

Note: Chapter 7. “Controlling database access” in *IBM DB2 Universal Database Administration Guide: Implementation, Version 8.2 (SC09-4820)* explains the security concepts of the DB2 server and describes how a user ID can be authorized to perform certain tasks.

Post-installation tasks for remote DB2 setup

The following tasks must be completed on the DB2 server system:

- Identify the DB2 instance that will hold the databases. Identify the instance owner user ID and password. Identify the host name of the DB2 server system. Identify the TCP/IP port of the DB2 instance.
- Create the automation manager database (for information on how to do this, see below).
- Create the automation manager tables in the database (for information on how to do this, see below).

Note: If the database has already been created and tables already exist, you must drop the existing tables before creating the tables.

- Create the operations console database (for information on how to do this, see below).

The CD labeled *IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component* for your platform contains scripts for creating the required databases and tables.

Creating the automation manager database and the database tables

On Windows: Perform the following steps if your DB2 server runs under Windows:

1. Log in with a user ID that has SYSADM privileges on the DB2 instance.
2. On the CD *IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component for Windows*, change the directory to DDL\Script
3. Run the following batch file from this directory:

```
db2_create_automgr_db.bat <db_name> <instance_owner> <instance_pwd>
```

where

- <db_name> is the desired name of the automation manager database (Example: EAUTODB)
- <instance_owner> is the instance owner user ID of the DB2 instance (Example: db2admin)
- <instance_pwd> is the password of the instance owner user ID

On AIX and Linux: Perform the following steps if your DB2 server runs under Linux or AIX:

1. Log in as root.
2. On the CD labeled *IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component* for your operating system, change the directory to DDL/Script.
3. Run the following shell script from this directory:

```
db2_create_automgr_db.sh <db_name> <instance_owner> <instance_pwd>
```

where

- <db_name> is the desired name of the automation manager database (Example: EAUTODB)
- <instance_owner> is the instance owner user ID of the DB2 instance (Example: db2inst1)
- <instance_pwd> is the password of the instance owner user ID

Creating the operations console database

On Windows: Perform the following steps if your DB2 server runs under Windows:

1. Log in with a user ID that has SYSADM privileges on the DB2 instance.
2. On the CD labeled *IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component for Windows*, change the directory to DDL\Script
3. Run the following batch file from this directory:

```
db2_create_opcons_db.bat <db_name>
```

where <db_name> is the desired name of the operations console database (Example: OPCONDB)

On AIX and Linux: Perform the following steps if your DB2 server runs under Linux or AIX:

1. Log in as root.
2. On the CD labeled *IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component* for your operating system, change the directory to DDL/Script.
3. Run the following shell script from this directory:
`db2_create_opcons_db.sh <instance_owner> <db_name>`

where

- <instance_owner> is the instance owner user ID of the DB2 instance (Example: db2inst1)
- <db_name> is the desired name of the operations console database (Example: OPCONDB)

Installing a DB2 client

The following sections only apply when you are using a remote DB2 setup.

DB2 client requirements

Check which requirements need to be met for installing and running a DB2 client. The information can be found in the following publications:

- IBM DB2 Universal Database - Quick Beginnings for DB2 Clients - Version 8.2 (GC09-4832)
- IBM DB2 Universal Database - Release Notes - Version 8

The latest versions of these publications can be found on the IBM DB2 UDB Web site at

www.ibm.com/software/data/db2/udb/support/

You find the link to the PDF manuals in the **Other resources** section on the Web page.

In addition, check for the latest system requirements at

www.ibm.com/software/data/db2/udb/sysreqs.html

The DB2 release notes can also be found on the CD labeled *IBM DB2 Run-time Client Version 8.2* for your platform. Make sure that all requirements for installing and running a DB2 client are met. Otherwise, the DB2 client or the end-to-end automation management component may not work properly.

DB2 client installation

You can use the DB2 Setup wizard to install the DB2 client. You find the DB2 Setup wizard on the CD labeled *IBM DB2 Run-Time Client Version 8.2* for your platform.

If possible, use a 31-bit or 32-bit DB2 client.

Make sure that the DB2 client has the required version level.

When you install the end-to-end automation management component, you will need the following information:

- The directory into which the DB2 client is installed.

Post-installation tasks for remote DB2 setup

The following tasks must be performed on the DB2 client system:

- Catalog a TCP node for the DB2 server system. This is described in “Cataloging a TCP node.”
- Catalog the operations console database. This is described in “Cataloging the operations console database.”

This step generates a database alias.

- Catalog the automation manager database. This is described in “Cataloging the automation manager database” on page 64.

This step generates a database alias.

Cataloging a TCP node

Log in to the system as a user with sufficient database privileges:

- **Windows:** Log in as a user with SYSADM privileges on the DB2 client. Open a DB2 command line prompt.
- **AIX/ Linux:** Log in to DB2 using the name of the DB2 instance or your database user name. If you are logged in as root, enter the following command to change to the DB2 instance owner:

```
su - <INSTANCE_OWNER>
```

Issue the following command:

```
db2 CATALOG TCPIP NODE <NODE_NAME> REMOTE <DB2_SERVER_HOST>  
SERVER <DB2_SERVER_INSTANCE_PORT>
```

where

- <NODE_NAME> is an arbitrary name for the node
- <DB2_SERVER_HOST> is the (fully qualified) host name or the IP address of the DB2 server system
- <DB2_SERVER_INSTANCE_PORT> is the number of the port on which the selected DB2 server instance is listening

Cataloging the operations console database

Log in to the system as a user with sufficient database privileges:

- **Windows:** Log in as a user with SYSADM privileges on the DB2 client. Open a DB2 command line prompt.
- **AIX/ Linux:** Log in to DB2 using the name of the DB2 instance or your database user name. If you are logged in as root, enter the following command to change to the DB2 instance owner:

```
su - <INSTANCE_OWNER>
```

Run the following command:

```
db2 CATALOG DB <DB_NAME> AS <DB_ALIAS> AT NODE <NODE_NAME>
```

where

- <DB_NAME> is the name of the operations console database on the DB2 server
- <DB_ALIAS> is an arbitrary alias for the operations console database, which is later on specified in the installation wizard of the end-to-end management component
- <NODE_NAME> is the arbitrary name for the node that you specified when you cataloged the TCP node

Cataloging the automation manager database

Log in to the system as a user with sufficient database privileges:

- **Windows:** Log in as a user with SYSADM privileges on the DB2 client. Open a DB2 command line prompt.
- **AIX/ Linux:** Log in to DB2 using the name of the DB2 instance or your database user name. If you are logged in as root, enter the following command to change to the DB2 instance owner:

```
su - <INSTANCE_OWNER>
```

Run the following command:

```
db2 CATALOG DB <DB_NAME> AS <DB_ALIAS> AT NODE <NODE_NAME>
```

where

- <DB_NAME> is the name of the automation manager database on the DB2 server
- <DB_ALIAS> is an arbitrary alias for the automation manager database, which is later on specified in the installation wizard of the end-to-end management component
- <NODE_NAME> is the arbitrary name for the node that you specified when you cataloged the TCP node

Installing WebSphere Application Server

Note that installing WebSphere Application Server from the IBM WebSphere Application Server Base Version 6.0 CD does not result in the required version level. To obtain the required version level, additional service from the WAS 6.0.0 upgrade CD must be applied in the following sequence:

1. WebSphere Application Server 6.0 Refresh Pack 2, to obtain WebSphere Application Server 6.0.2
2. Particular Interim Fixes

WebSphere Application Server 6.0.0.0 requirements

Check which requirements need to be met for installing and running WebSphere Application Server Base. The information can be found in the following publications:

- The ReadMe file on the product CD labeled *IBM WebSphere Application Server, Version 6*
- The "Getting started" topics in the Information center for IBM WebSphere Application Server, Version 6

The latest versions of all WebSphere Application Server publications can be found on the WebSphere Application Server library Web site at

www.ibm.com/software/webservers/appserv/was/library/

In addition, check for the latest system requirements at

www.ibm.com/software/webservers/appserv/was/requirements/

An IBM WebSphere Application Server, Version 6, *Getting started* document is available on the product CD for your platform, where it is also referred to as *Installation Guide*. Make sure that all requirements for installing and running

WebSphere Application Server are met. Otherwise, the end-to-end automation management component may not work properly.

Installing WebSphere Application Server 6.0.0.0

You can use the WebSphere Application Server installation wizard to install WebSphere Application Server. The WebSphere Application Server installation wizard can be started from the WebSphere Application Server LaunchPad. You find the LaunchPad on the CD labeled *IBM WebSphere Application Server Version 6.0* for your platform.

The typical installation is recommended.

Installing Refresh Pack 2 and the required Interim Fixes

The following table gives an overview of what you will find on the WebSphere Application Server 6.0.0 upgrade CD. For details, refer to the sections below.

Table 13. Contents of the WebSphere Application Server 6.0.0 upgrade CD

| Directory | Contents |
|-----------|--|
| Upgrade | An archive containing Refresh Pack 2 and the UpdateInstaller for installing Refresh Pack 2 |
| Fixes | The archive that contains the subdirectories with the required Interim Fixes |

Installing Refresh Pack 2

In the directory Upgrade on the WebSphere Application Server 6.0.0 upgrade CD, you find an archive that contains Refresh Pack 2 and the UpdateInstaller for installing it.

The archive names have the following syntax:

6.0-WAS-<platform>-RP<refresh_pack_number>.<archive_type>

where

- <platform> represents the platform on which the end-to-end automation management component is installed
- <refresh_pack_number> represents the number of the refresh pack
- represents the platform-specific file extension of the archive

The documentation for the UpdateInstaller is available in the archive in directory updateinstaller\docs.

For more information about Refresh Pack 2, refer to

<http://www.ibm.com/support/docview.wss?rs=180&uid=swg24009813>.

On the Web page, read the document *Readme for multiplatforms* to obtain information about how to install the Refresh Pack.

Installing the required Interim Fixes

In the directory Fixes on the WAS 6.0.0 upgrade CD, you find the required Interim Fixes.

The required Interim Fixes are located in directories that are prefixed with a number:

```
Fixes
01_PKxxxxx
<NAMExxxxx>.pak
...
nn_PKyyyyy
<NAMEyyyyy>.pak
```

The Interim Fixes are the files with the extension .pak.

When you install the Interim Fixes, be sure to observe the following rules:

- Install all Interim Fixes that are located in the directory Fixes.
- Install the Interim Fixes in exactly the sequence indicated by the number prefix of the directories in the directory Fixes.
- Install exactly these Interim Fixes, do not leave out any Interim Fixes and do not install any additional Interim Fixes unless you are explicitly advised to do so by Tivoli System Automation product support.

Post-installation tasks

Remove the default WebSphere Application Server profile using the `wasprofile` command and create a new default profile using the profile creation wizard.

When creating the new default profile, it is recommended that you accept the default port settings.

Refer to Chapter 10. "Configuring the product after installation" in the manual *WebSphereApplication Server, Version 6 - Installing your application serving environment* for more information on the `wasprofile` command and the profile creation wizard.

Under Windows, you can choose to run WebSphere Application Server as a service. When doing so, make sure that the user ID used to run the service has a valid DB2 environment setup.

For more information on setting up the DB2 environment, refer to Chapter 1. "Before creating a database" -> "Preparing to create a database" in *IBM DB2 Universal Database Administration Guide: Implementation, Version 8.2 (SC09-4820)*.

Chapter 8. Setting up an LDAP server

Note that the IBM Tivoli Directory Server is not contained on the middleware software CDs that are shipped with the end-to-end automation management component. If you want to use LDAP as the user registry for the end-to-end automation management component, you have to obtain a supported version of the product separately.

When LDAP is used as the user registry, user authentication and user membership in groups is performed using an LDAP server.

Before you can install the end-to-end automation management component, you need to create at least one group and one user for the end-to-end automation management component. The group that is required is `iscadmins` or an equivalent. Members of this group have administrative authority within the operations console. At least one operations console administrator user must be created within the group (refer to “Required user groups and users” on page 68 for more information).

Required LDAP directory tree structure

The LDAP directory tree that is used by the operations console must meet the following requirements:

- One distinguished name (DN) suffix needs to be defined. It must contain the users and groups mentioned above.
- The operations console must be able to add, delete, modify, and search entries under this suffix. In fact, new entries will be added below the suffix during the installation of the end-to-end automation management component.
- Directly below the suffix, one entry is needed for containing user entries. The relative distinguished name (RDN) of this entry is the LDAP user suffix.
- Directly below the suffix, one entry is needed for containing group entries. The relative distinguished name (RDN) of this entry is the LDAP group suffix.
- The object class of the user entries must allow for a password specification so that an LDAP bind request can be performed against user entries. An example of such an object class is `inetOrgPerson`, which is defined in IETF RFC 2798.
- The object class of the group entries must allow for specifying which user entries are members of a group. An example of such an object class is `groupOfUniqueNames`, which is defined in IETF RFC 2256.

The following figure shows the required structure. For an example of an LDAP configuration, refer to “Sample LDAP configuration” on page 69.

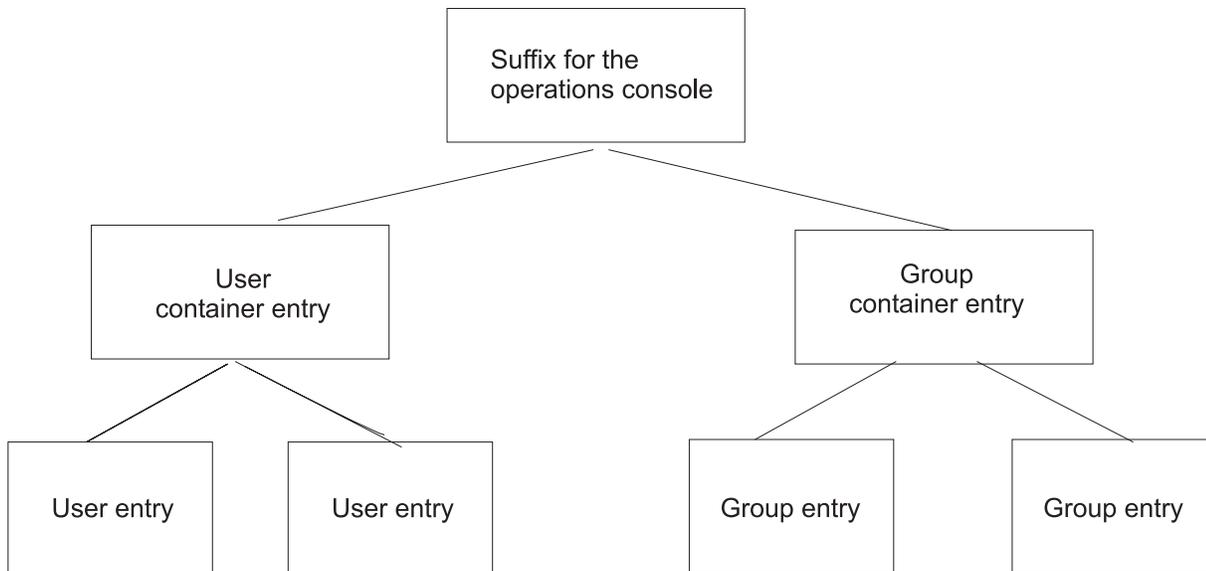


Figure 4. LDAP directory tree structure

Required user groups and users

When using LDAP as the user registry, the following rules apply:

- A system user is required for operating the automation engine. In this guide, this user is referred to as `iscadmin`.

For this user, you need to create a group (`iscadmins` or an equivalent) and a user (`iscadmin` or equivalent) in the LDAP user registry before installing the end-to-end automation management component. You will need to enter this user ID on the LDAP configuration panels in the installation wizard during the installation of the end-to-end automation management component.

- A number of additional groups are required for the end-to-end automation management component. You can either create these groups in the LDAP user registry before installing the end-to-end automation management component or create them in Integrated Solutions Console after the installation is complete.

Regardless of which approach you use for creating the groups, you need to map the groups to the end-to-end automation management-specific access roles in WebSphere Application Server after installation.

For more information on the required user groups and the related access roles, refer to Chapter 23, “Post-installation tasks for administrators,” on page 143.

- You can create new users or authorize existing users for end-to-end automation management. You can create new users in the LDAP user registry or in Integrated Solutions Console. You authorize the user by assigning them to the end-to-end automation-specific user groups, either in the LDAP user registry or in Integrated Solutions Console.

For an example of an LDAP configuration, refer to “Sample LDAP configuration” on page 69.

LDAP-related pre-installation tasks

Before you install the end-to-end automation management component, perform the following task on the LDAP server:

1. Install and configure the LDAP server as described in the installation instructions provided with the LDAP server.
-
2. You must create at least one group (iscadmins or an equivalent) and create at least one user (iscadmin or an equivalent) in that group.
Observe the following guidelines when creating the user ID and password:
 - The user ID and password must be unique
 - The length must be 3 to 60 characters
 - Valid characters are a-z, A-Z, period (.), hyphen (-), and underscore (_)
No other characters are permitted. For example, diacritics, such as the umlaut, and double-byte characters are not permitted.
-

Sample LDAP configuration

The following LDIF creates a user ID and a group in the LDAP server.

Before the LDIF can be loaded, the suffix must exist. You can create it using the following command:

```
ldapcfg -s o=<your_organization_suffix>
```

for example:

```
ldapcfg -s o=ibm.com
```

Use the following LDAP server command to load the LDIF file into the LDAP server:

```
ldapmodify -h <ldap_server_hostname> -p 389  
-D cn=<ldap_administrator> -w ldapadm -f e2e.ldif
```

for example:

```
ldapmodify -h <ldap_server_hostname> -p 389 -D cn=ldapdb2 -w ldapadm -f e2e.ldif
```

This is a sample LDIF file:

```
dn: o=ibm.com  
objectclass: top  
objectclass: organization  
o: ibm  
o: ibm.com  
  
dn: ou=users,o=ibm.com  
objectclass: top  
objectclass: organizationalUnit  
ou: users  
  
dn: ou=groups,o=ibm.com  
objectclass: top  
objectclass: organizationalUnit  
ou: groups  
  
dn: uid=iscadmin,ou=users,o=ibm.com  
objectclass: top  
objectclass: person  
objectclass: organizationalPerson  
objectclass: inetOrgPerson  
uid: iscadmin  
userpassword: iscadmin  
sn: iscadmin
```

```
cn: iscadmin
dn: cn=iscadmins,ou=groups,o=ibm.com
objectclass: top
objectclass: groupOfUniqueNames
ou: iscadmins
cn: iscadmins
uniqueMember: uid=iscadmin, ou=users, o=ibm.com
```

Chapter 9. Preparing for the installation of the end-to-end automation management component

Collecting the information you need to provide during installation

The installation of the end-to-end automation management component is wizard-driven. The wizard guides you through the installation and prompts you for installation and configuration parameters. The following tables list the parameters you need to specify on the installation wizard panels in the order in which they must be specified.

Installation directory and Tivoli Common Directory

The parameters listed in the following table must always be specified.

Table 14. Installation directory and Tivoli Common Directory

| Parameter | Description | Default |
|-----------------------------|--|--|
| Installation directory name | <p>The directory to which the installable features are installed.</p> <p>In this guide, this directory is referred to as EEZ_INSTALL_ROOT.</p> | <p>Windows:</p> <p>C:\Program Files\IBM\tsamp\eez</p> <p>AIX, Linux:</p> <p>/opt/IBM/tsamp/eez</p> |
| Tivoli Common Directory | <p>The Tivoli directory for storing serviceability information.</p> <p>During installation, you are only prompted for input when no Tivoli Common Directory is found on the system.</p> <p>A directory name eeZ is created in the Tivoli Common Directory for storing product-specific data.</p> <p>In this guide, this directory is referred to as Tivoli_Common_Directory.</p> | <p>Windows:</p> <p>C:\Program Files\IBM\tivoli\common</p> <p>AIX, Linux:</p> <p>/var/ibm/tivoli/common</p> |

Table 14. Installation directory and Tivoli Common Directory (continued)

| Parameter | Description | Default |
|--|--|---|
| <p>Operations console installation directory</p> | <p>The directory into which the operations console is installed.</p> <p>In this guide, this directory is referred to as <code>isc_runtime_root</code>.</p> <p>The directory also contains the product registry for Integrated Solutions Console (<code>product.reg</code>) and the properties file (<code>isc.properties</code>). Do not modify these files. They are used by the operations console at runtime.</p> <p>When specifying a directory other than the default, observe the following restrictions:</p> <ul style="list-style-type: none"> • If you use an existing directory, the directory cannot contain any of the following files and directories: <ul style="list-style-type: none"> – The files <code>product.reg</code> and <code>isc.properties</code> – The directory <code>_uninst</code> or a file named <code>_uninst</code> – The directory <code>\AppServer</code> or a file named <code>AppServer</code> • The length of the installation path must be 32 characters or less. • On AIX and Linux systems, the space character and the following special characters are not allowed in the installation path: <code>- \ : = ~ * ! ? () ' # " + < > [] \$ { }</code> • On Windows systems, the space character and the following special characters are not allowed in the installation path: <code>/ % * ! ? () ' # " + < > [] \$ { }</code> <p>The only exception for the colon character is immediately after the drive letter. For example, <code>C:\ISC</code> is allowed, but <code>C:\ISC:DIR</code> is not allowed.</p> • Use only ISO-8859-1 characters in the installation path. | <p>Windows:</p> <p><code>C:\Program Files\IBM\ISC\</code></p> <p>AIX, Linux:</p> <p><code>/opt/IBM/ISC</code></p> |

Installation parameters for DB2

The parameters listed in the following table must always be specified.

Table 15. DB2 data for local and remote DB2 setup

| Parameter | Description | Default |
|----------------------------------|---|---|
| DB2 UDB directory | <p>The installation location of the DB2 client directory.</p> <p>If you are using a local DB2 setup, you use the DB2 client that is part of the DB2 server installation. In this case, you need to specify the DB2 server directory.</p> | The location is detected on your system and displayed as default directory. |
| DB2 instance host name | The host name of the DB2 instance in which the automation manager and operations console databases are located. | <p>Local DB2 setup: The fully qualified host name of your system is displayed as default value. Do not change it to localhost!</p> <p>Remote DB2 setup: No default value is provided.</p> |
| DB2 instance port number | The port number of the DB2 instance in which the automation manager and operations console databases are located. | 50000 |
| Database instance owner name | <p>The instance owner user ID of the DB2 instance in which the automation manager and operations console databases are located.</p> <p>In a local DB2 setup, this user ID will be used for creating the databases and tables.</p> <p>In a remote DB2 setup, the user ID will be used for creating tables.</p> <p>The user ID will be used by WebSphere Application Server to connect to the automation manager and operations console databases and to select, insert, delete, and update rows in tables.</p> | <p>Windows: db2admin</p> <p>AIX, Linux: db2inst1</p> |
| Database instance owner password | The password for the instance owner user ID of the DB2 instance in which the automation manager and operations console databases are located. | No default value is provided |
| Automation manager database | <p>Automation manager database for use by WebSphere Application Server.</p> <p>In a local DB2 setup, a database with this name will be created in the DB2 instance related to the specified instance owner.</p> <p>In a remote DB2 setup, a database with this name must already exist in the remote DB2 instance.</p> | EAUTODB |

Table 15. DB2 data for local and remote DB2 setup (continued)

| Parameter | Description | Default |
|-----------------------------|---|---------|
| Operations console database | <p>Operations console database for use by WebSphere Application Server.</p> <p>In a local DB2 setup, a database with this name will be created in the DB2 instance related to the specified instance owner.</p> <p>In a remote DB2 setup, a database with this name must already exist in the remote DB2 instance and an alias must be defined to the local system.</p> | OPCONDB |

Installation parameters for WebSphere Application Server

The parameters listed in the following table must always be specified.

Table 16. WebSphere Application Server installation parameters

| Parameter | Description | Default |
|--|---|--|
| WebSphere Application Server directory | The installation location of WebSphere Application Server. There must be exactly one installation of WebSphere Application Server on your system. | The location is detected on your system and displayed as default directory. |
| WebSphere Application Server profile | The WebSphere Application Server profile to be used for the automation manager and the operations console. | All existing profiles are detected on your system and displayed in a single-choice list. |
| WebSphere Application Server name | The server to be used for the automation manager. | The server name is detected on your system and displayed as default value. |

Installation parameters for LDAP

The parameters listed in the following table only need to be specified when you are using LDAP as the user registry.

Table 17. Installation parameters for LDAP

| Parameter | Description | Default |
|-----------------------|--|------------------------------|
| LDAP server host name | The fully-qualified host name of the LDAP server that will be used by the operations console. | No default value is provided |
| LDAP server port | <p>The port of the LDAP server that will be used by the operations console. The port number depends on whether SSL is used:</p> <ul style="list-style-type: none"> The default non-SSL port is 389 The default SSL port is 636 <p>Note: On the relevant installation wizard panel, port number 389 is always displayed as default value. If you will be using SSL, you must manually change the port number to the appropriate value.</p> | 389 (non-SSL) |

Table 17. Installation parameters for LDAP (continued)

| Parameter | Description | Default |
|------------------------|---|------------------------------|
| LDAP admin bind DN | <p>The user ID of the LDAP directory administrator in the distinguished name format.</p> <p>The operations console uses this ID to bind to the LDAP to retrieve user attributes, to create new users and groups in the LDAP, and to update user attributes.</p> <p>This ID is not required to be the LDAP admin DN, but rather an ID with sufficient authority for the use cases mentioned above.</p> <p>Note: Make sure to type the value in lower case, regardless of the case used in the distinguished name (DN).</p> <p>Example: cn=ldapdb2</p> | No default value is provided |
| LDAP admin password | The password of the LDAP directory administrator. | No default value is provided |
| LDAP suffix | <p>The LDAP directory DN suffix which the operations console uses for storing user and group data.</p> <p>Example: o=ibm.com</p> | No default value is provided |
| Console admin bind DN | <p>The user ID for the LDAP bind authentication in the distinguished name format.</p> <p>This user ID is used by WebSphere Application Server to bind to the LDAP server to retrieve user attributes required for authentication.</p> <p>Note: Make sure to type the value in lower case, regardless of the case used in the distinguished name (DN).</p> <p>Example: uid=iscadmin,ou=users,o=ibm.com</p> | No default value is provided |
| Console admin password | The password for the console admin bind DN. | |
| Use SSL for LDAP | <p>Indicates whether Secure Sockets Layer (SSL) communication is enabled for the LDAP server.</p> <p>Whether SSL is enabled or not determines which port number must be used for the LDAP port.</p> | False |

Table 17. Installation parameters for LDAP (continued)

| Parameter | Description | Default |
|--------------------------------|---|--------------------|
| LDAP user prefix | The attribute name used for building the relative distinguished name (RDN) of user entries in the LDAP directory tree. The DN of user entries begins with this prefix. If the LDAP user object class is inetOrgPerson, this value is uid. | uid |
| LDAP user suffix | The RDN of the user container entry in the LDAP directory tree. | ou=users |
| LDAP user object class | The object class used by the LDAP server to store user entries. | inetOrgPerson |
| Console admin group short name | The operations console administrator group name. The ISC administrator user is a member of this group. | iscadmins |
| LDAP group prefix | The attribute name used for building the relative distinguished name (RDN) of group entries in the LDAP directory tree. The DN of group entries begins with this prefix. If the LDAP group object class is groupOfUniqueNames, this value is cn. | cn |
| LDAP group suffix | The RDN of the group container entry in the LDAP directory tree. | ou=groups |
| LDAP group object class | The object class used by the LDAP server to store group entries. | groupOfUniqueNames |
| LDAP group member attribute | The property that specifies the attribute name of the membership attribute of group entries in the LDAP directory tree. If the LDAP group object class is groupOfUniqueNames, this attribute name is uniqueMember. | uniqueMember |

Installation parameters for the operations console

The parameters listed in the following table must always be specified.

Table 18. Installation parameters for Integrated Solutions Console

| Parameter | Description | Default |
|---------------------------|--|--|
| ISC administrator user | <p>The user ID of the operations console administrator.</p> <p>During installation, the administrator is given access to all console modules. The user ID is added to the operations console administrator group.</p> <p>After the installation, the administrator can change the password and add other user IDs to the operations console administrator group.</p> <p>The user ID must comply with the following conditions:</p> <ul style="list-style-type: none"> • The user ID must be unique. • The length is 3 to 20 characters. • A valid user ID may contain only the characters a-z, A-Z, period (.), hyphen (-), underscore (_), and double-byte character set (DBCS) characters. <p>No other characters are permitted in this field. For example, diacritics, such as the umlaut, are not permitted.</p> <p>If security with a DB2 database user registry is enabled, this user is created during the installation.</p> <p>If security with an LDAP user registry is enabled, this user must already exist before beginning the installation.</p> | iscadmin |
| Password | <p>The password of the operations console administrator.</p> <p>The password must comply with the following conditions:</p> <ul style="list-style-type: none"> • The length is 5 to 60 characters. • A valid password may contain only the characters a-z, A-Z, period (.), hyphen (-), and underscore (_). <p>No other characters are permitted in this field. For example, DBCS characters and diacritics, such as the umlaut, are not permitted.</p> <p>If security with a DB2 database user registry is enabled, the password is for a user which is created during installation. For this reason, the password must be confirmed.</p> <p>If security with an LDAP user registry is enabled, this password is for a user which must already exist before beginning the installation.</p> | No default value is provided |
| Fully qualified host name | The fully qualified host name of the system where the operations console will be installed | The fully qualified host name is detected on your system and displayed as default value. |

Table 18. Installation parameters for Integrated Solutions Console (continued)

| Parameter | Description | Default |
|--|--|---------|
| HTTP Port | <p>The number of the HTTP port that the operations console will use.</p> <p>Select a port that is not being used by another process on the system.</p> <p>After the operations console is installed, you must include this port number in the URL for opening the console.</p> <p>The URL is composed of the protocol name, plus the fully-qualified host name, plus the port, plus <code>ibm/console</code>.</p> <p>This is an example of a full URL as it is needed for connecting to the operations console: <code>http://myhost.com:8421/ibm/console</code></p> | 8421 |
| HTTPS Port | <p>The port that the operations console will use for secure HTTP transport (HTTPS).</p> <p>This value must not conflict with existing port assignments on the system.</p> <p>To enable HTTPS, you also must perform the procedure described as post-installation task in “Setting up SSL for the operations console” on page 103 after the operations console is installed.</p> | 8422 |
| Bootstrap/RMI Port | <p>The address for the bootstrap function and the port number for the Java Remote Method Invocation (RMI) connector on the operations console server.</p> <p>This value must not conflict with existing port assignments on the system.</p> | 8424 |
| SOAP Port | <p>The address for the Simple Object Access Protocol (SOAP) connector on the operations console server.</p> <p>This value must not conflict with existing port assignments on the system.</p> | 8425 |
| Admin HTTP Port | <p>The HTTP Administrative Console port on the operations console server.</p> <p>This value must not conflict with existing port assignments on the system.</p> | 8431 |
| Admin HTTPS Port | <p>The HTTPS Administrative Console secure port on the operations console server.</p> <p>This value must not conflict with existing port assignments on the system.</p> | 8432 |
| SAS SSL ServerAuth Listener Address Port | <p>The SAS SSL ServerAuth Listener Address port on the operations console server. This value must not conflict with existing port assignments on the system.</p> | 8439 |
| CSIV2 SSL ServerAuth Listener Address Port | <p>The CSIV2 SSL ServerAuth Listener Address port on the operations console server.</p> <p>This value must not conflict with existing port assignments on the system.</p> | 8440 |

Table 18. Installation parameters for Integrated Solutions Console (continued)

| Parameter | Description | Default |
|---|---|---------|
| CSIV2 SSL MutualAuth Listener Address Port | The CSIV2 SSL MutualAuth Listener Address port on the operations console server. This value must not conflict with existing port assignments on the system. | 8441 |
| Console Help Port | The port that the help system (based on Eclipse technology) will use to receive requests for help files. This value must not conflict with existing port assignments on the system. | 8423 |
| Register ISC server and ISC Help server as system service | To automatically restart the operations console and the console help server each time the system is restarted, these services can be registered as system services | Enabled |
| Console Service ID | To automatically restart the console server each time the operating system is restarted, specify this parameter. For Linux, the string must be 1 to 4 characters. The installation program checks the length of the value you specify. For AIX, the length must be 1 or more characters. For Linux and AIX, the operating system file /etc/inittab is edited directly to include the value you specify. The line that is added to the file has the following format: <pre>service_ID:23:boot:isc_runtime_root/PortalServer/ bin/startISC.sh ISC_Portal ISCUSER ISCPASS</pre> For Windows systems only: Set the value to a unique string. Valid characters are a-z, A-Z, and 0-9. The string must be 1 or more characters and the value is used to add a service to the operating system. If this parameter is specified, you must also specify the Console Help Service ID parameter. | CS01 |
| Console Help Service ID | For Windows systems only: Set the value to a unique string. Valid characters are a-z, A-Z, and 0-9. The string must be 1 or more characters. The value is used to add a service to the operating system. If this parameter is specified, you also must specify the Console Service ID parameter. For AIX and Linux: The Console Help Service is started as part of the service defined by the Console Service ID parameter. The Console Help Service ID parameter is not shown. | HS01 |

Installation parameters for IBM Tivoli Enterprise Console

Optionally, you can use Tivoli Enterprise Console for monitoring end-to-end automation management events. The parameters listed in the following table are only required if you will be utilizing Tivoli Enterprise Console for end-to-end automation management.

Table 19. Installation parameters for IBM Tivoli Enterprise Console

| Parameter | Description | Default |
|------------------------|---|-----------|
| TEC host name | The name of the host where the Tivoli Enterprise Console server is installed. | localhost |
| TEC server port number | The port number for the Tivoli Enterprise Console server | 5529 |

Note: When you TEC server runs on AIX or Linux, the TEC server port number must be set to 0. For more information, refer to Chapter 30, “Using Tivoli Enterprise Console with SA for Multiplatforms,” on page 199.

Name of the end-to-end automation domain

Table 20. Name of the end-to-end automation domain

| Parameter | Description | Default |
|------------------------|---|-------------|
| Automation domain name | The name of the end-to-end automation domain represented by this instance of the automation engine. The domain name must be unique and may not be used for any other automation domain. | FriendlyE2E |

What the installation CD contains

The end-to-end automation management component can be ordered from IBM as media pack or downloaded as an Electronic Software Distribution (ESD) image from an IBM software distribution download site.

There are multiple CDs for each supported platform.

This is what the CD labeled *IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component for <operating_system_name>* contains:

- The files for launching the installation wizard
- The readme file
- Directories containing the files required to install components that are embedded into the end-to-end automation management installation. These are:

Table 21. Directories on the product CD

| Directory | Content |
|-------------------------|---|
| README | For example, copyright notices and license agreements |
| license | License key |
| DDL | Scripts for creating DB2 databases and tables when remote DB2 setup is used |
| <PLATFORM> ¹ | Product installer and files needed for installing the product |

Note:

1. <PLATFORM> is one of the following:
 - AIX
 - PPC (Linux for PPC)
 - Windows
 - i386 (Linux for x/Series)
 - S390 (Linux for z/Series)

Languages supported by IBM Tivoli System Automation

This section is only of interest for you if you want to use IBM Tivoli System Automation for Multiplatforms in a language other than English as shown in the following tables.

The following encodings are supported for the Linux distribution:

| Language | UTF-8 | ISO-8859-1 | EUC/GBK | Euro | GB18030/BIG5 |
|-------------------------|-------------|----------------------------|----------------------------|------------|----------------------|
| German | de_DE.UTF-8 | de_DE, de_DE.ISO-8859-1 | | de_DE@euro | |
| Spanish | es_ES.UTF-8 | es_ES, es_ES.ISO-8859-1 | | es_ES@euro | |
| French | fr_FR.UTF-8 | fr_FR, fr_FR.ISO-8859-1 | | fr_FR@euro | |
| Italian | it_IT.UTF-8 | it_IT, it_IT.ISO-8859-1 | | it_IT@euro | |
| Japanese | ja_JP.UTF-8 | | ja_JP.eucJP | | |
| Korean | ko_KR.UTF-8 | | ko_KR.eucKR | | |
| Portugese/ Brazilian | pt_BR.UTF-8 | pt_BR | | | |
| Simplified Chinese | zh_CN.UTF-8 | | zh_CN.GBK, zh_CN.GB2312 | | zh_CN.GB18030 |
| Traditional Chinese | zh_TW.UTF-8 | | | | zh_TW.Big5, zh_TW |

The following encodings are supported on the AIX distribution:

| Language | UTF-8 | ISO-8859-1 | EUC/GBK | SJIS/GB18030/BIG5 |
|---------------------|-------|------------|---------|-------------------|
| German | DE_DE | de_DE | | |
| Spanish | ES_ES | es_ES | | |
| French | FR_FR | fr_FR | | |
| Italian | IT_IT | it_IT | | |
| Japanese | JA_JP | | ja_JP | Ja_JP |
| Korean | KO_KR | | ko_KR | |
| Portugese/Brazilian | PT_BR | pt_BR | | |
| Simplified Chinese | ZH_CN | | zh_CN | Zh_CN |
| Traditional Chinese | ZH_TW | | zh_TW | Zh_TW |

Installation prerequisites

The following prerequisites must be satisfied before you can start the installation wizard for the end-to-end automation management component:

- WebSphere Application Server must be installed as described in “Installing WebSphere Application Server” on page 64. There must be no other WebSphere Application Server product installation on the same system. Security must be disabled in WebSphere Application Server.
- A DB2 server must be installed as described in “Installing a DB2 server” on page 59. The DB2 server instance must be running and accepting client connections.

If you are using a local DB2 setup, the DB2 server instance should be empty or should at least not contain neither the automation manager database nor the operations console database from an earlier attempt to install a Tivoli System Automation for Multiplatforms component. Additionally, no other TCP/IP node must have been cataloged with a database alias that refers to an already existing automation manager or operations console database.

- If remote DB2 setup is used, a DB2 client must be installed on the same system as WebSphere Application Server as described in “Installing a DB2 client” on page 62.
- If LDAP is used as the user registry, an LDAP server must be configured as described in Chapter 8, “Setting up an LDAP server,” on page 67. The LDAP server must be running and accepting client connections.
- The user ID that is used to run the installer for the end-to-end automation management component must be able to run DB2 client code. This means that the DB2 environment must be made available for the user ID.

For more information on setting up the DB2 environment, refer to Chapter 1. “Before creating a database” -> “Preparing to create a database” in *IBM DB2 Universal Database Administration Guide: Implementation, Version 8.2 (SC09-4820)*.

On Linux and AIX systems, the DB2 environment must be prepared in a way such that the DB2 environment is automatically available in every sub-shell that is opened. This is usually accomplished by extending a startup shell script. The user ID that is used to run the installed end-to-end automation management component must have the same DB2 environment setup.

Attention:

On Linux and AIX systems, this DB2 environment setup is not automatically done by a DB2 server or DB2 client installation and must be performed manually. If this is not done, the installation of the end-to-end automation management component will fail.

- The user ID that is used to run the installer for the end-to-end automation management component must have administrator authority.
On Linux and AIX, this user ID is typically “root”.
- When installing the end-to-end automation management component to an AIX or Linux system, you must ensure that an XWindows session is available for displaying the graphical installation wizard panels.

Chapter 10. Installing the end-to-end automation management component

This section describes how to install the end-to-end automation management component. For the installation, you use a graphical installation program, the so-called installation wizard. The required steps are described below.

On the panels of the installation wizard, enter the data you have collected using the lists in section “Collecting the information you need to provide during installation” on page 71.

Notes:

1. Although the panels in this section show a Windows installation, the panels that are displayed for other operating systems have a similar appearance. Make sure to conform to the conventions of your platform when specifying directory locations, files names and so on.
2. In this section, only those panels are depicted on which user action is required.
3. The installation comprises two phases:
 - a. In the pre-installation phase, you specify the installation parameters. This may take up to half an hour.
 - b. The installation phase, which begins when you click the **Install** button on the last pre-installation panel, may take up to two hours to complete (depending on processor speed).

To install the end-to-end automation management component, perform these steps:

1. Make sure that all installation prerequisites are met (refer to “Installation prerequisites” on page 81).

-
2. Insert the following CD in the CD drive:

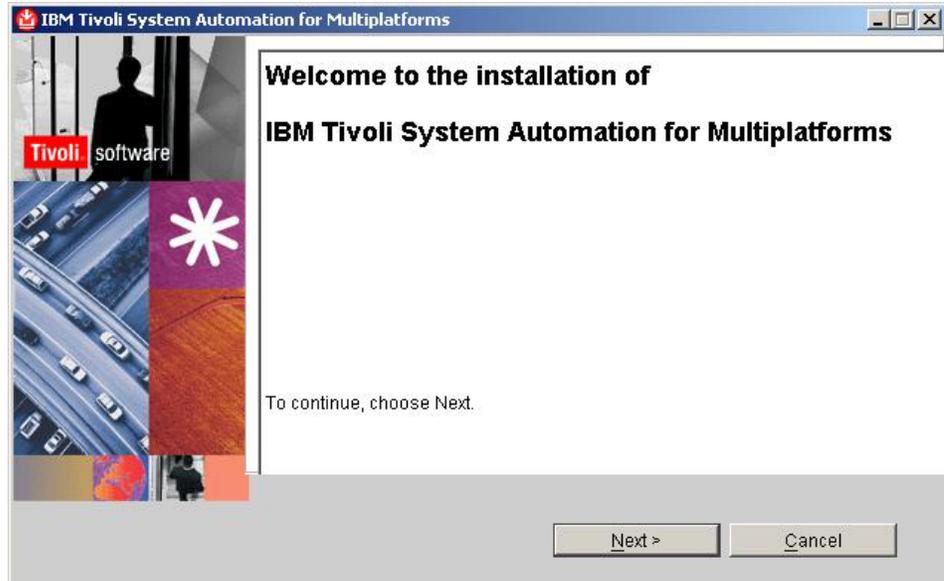
IBM Tivoli System Automation Multiplatform V2.1.1 End-to-end component for <operating_system_name>

There are multiple CDs. Be sure to use the one for your platform.

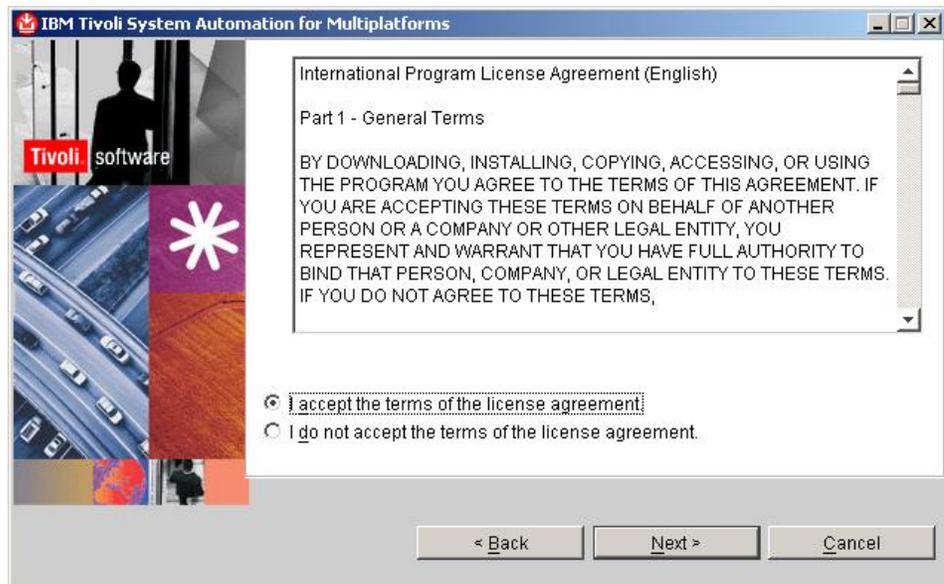
-
3. Launch the installation wizard by starting the following program from the <operating_system> subdirectory on the CD:
 - **Windows:** setup.exe
 - **AIX, Linux:** setup

When the wizard is launched successfully, the Welcome panel appears.

-
4. On the Welcome panel, click **Next** to display the License agreement panel.

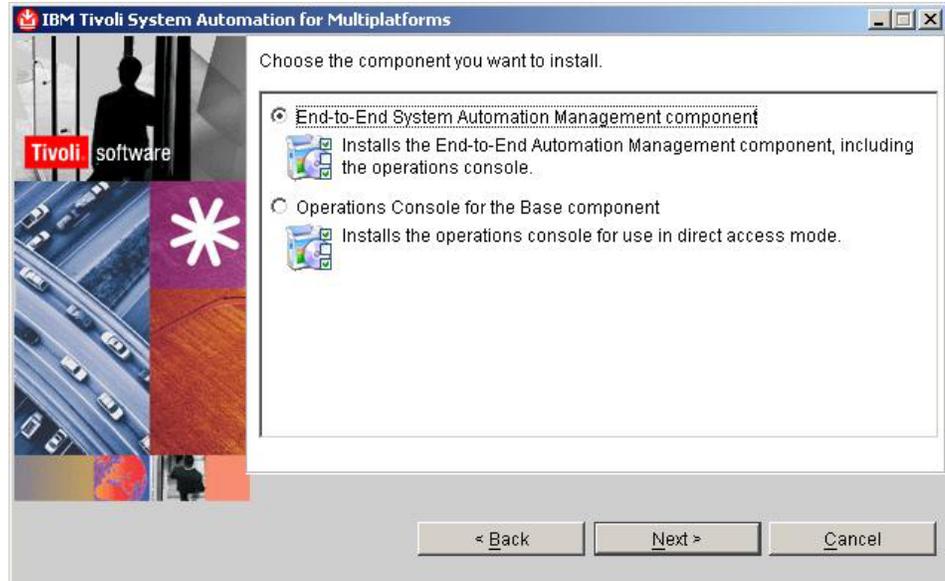


5. Select **I accept the terms of the license agreement** and click **Next**.

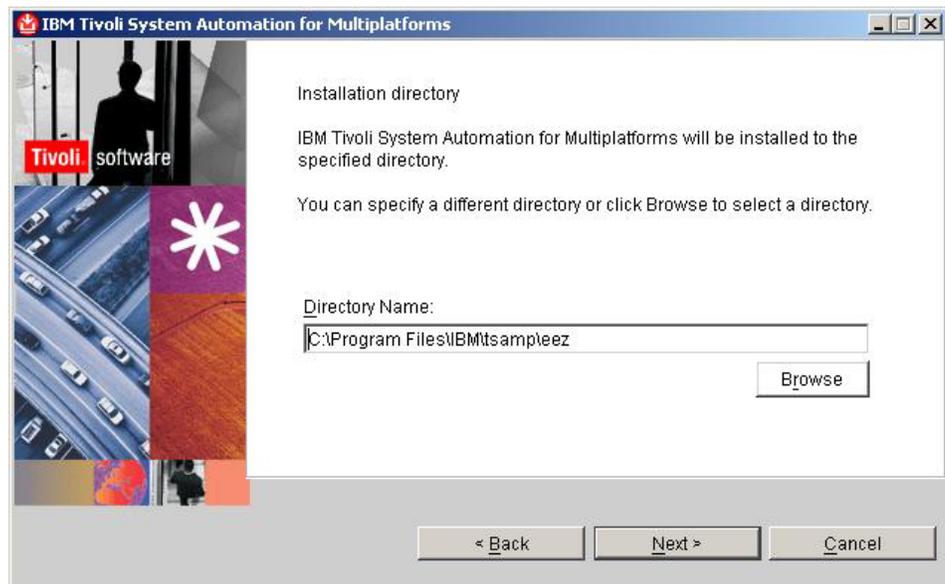


Note: After you click **Next**, the installation wizard checks whether the subdirectories **ISC**, **EIF**, and **ISC** are available in the directory in which the installation wizard program is located. If the subdirectories are not found, a panel is displayed that prompts you for the fully qualified path to the directory in which the subdirectories can be found.

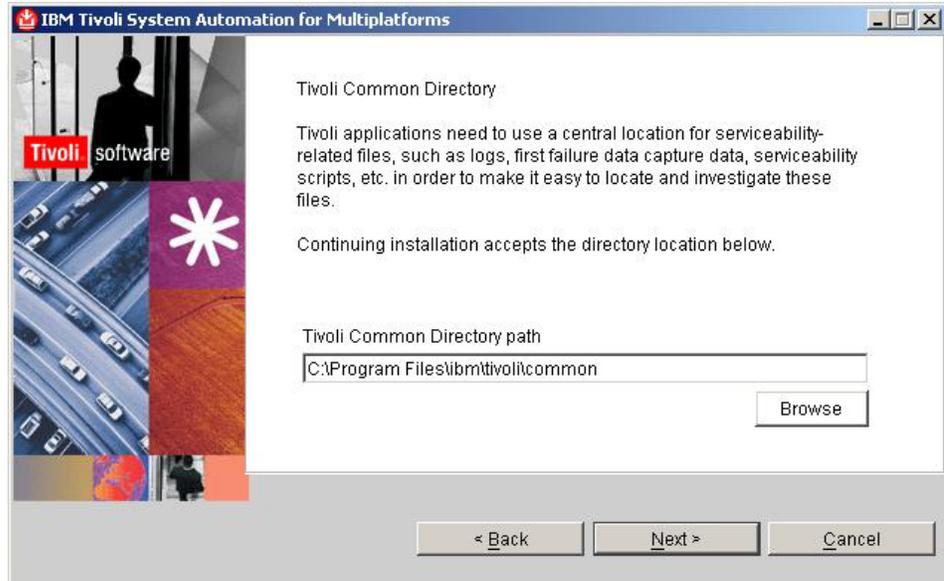
6. Select **End-to-End System Automation Management component** and click **Next**.



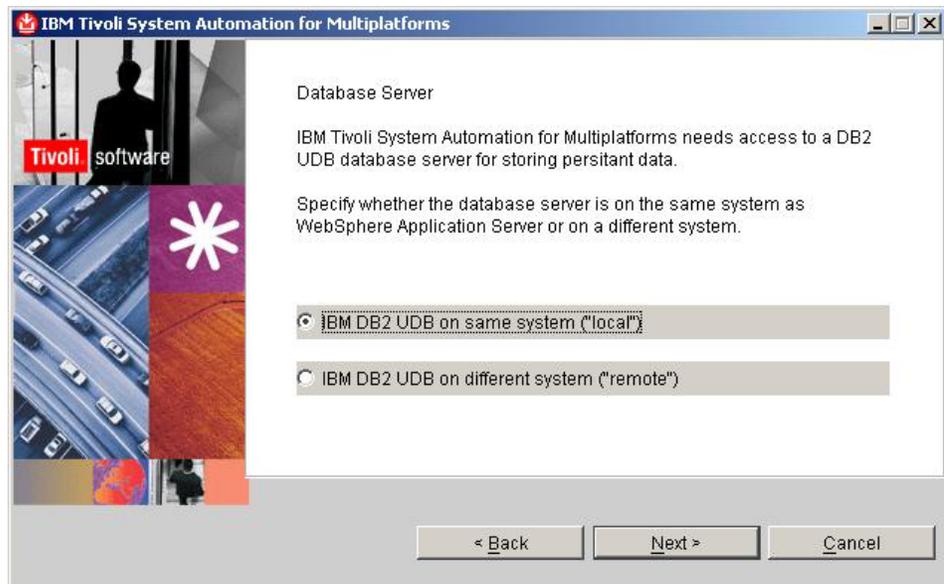
7. Specify the directory where you want to install the end-to-end automation management component or accept the default location.
Click **Next**.



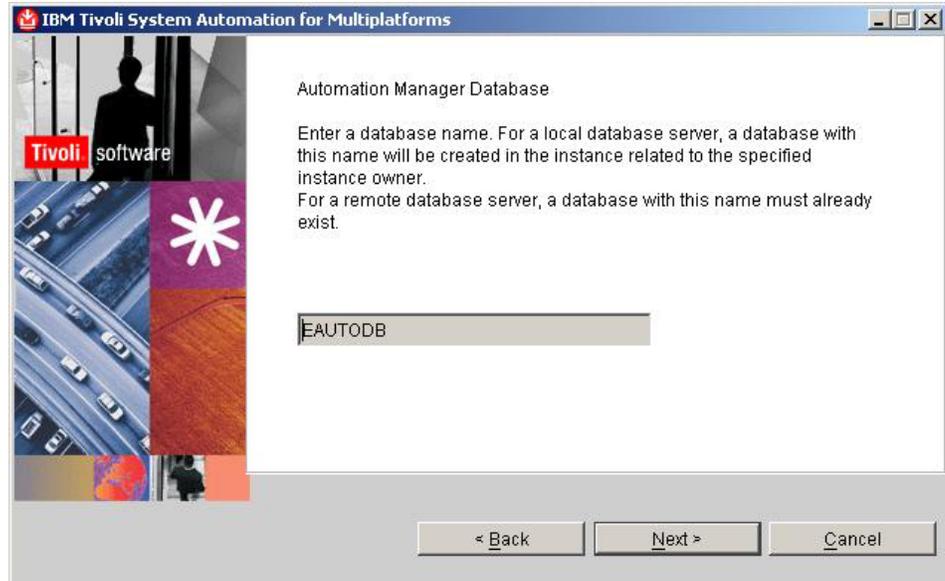
8. If the installation program did not detect a Tivoli Common Directory on your system, accept the default location or specify the directory to which the Tivoli log files are to be written. Click **Next**.



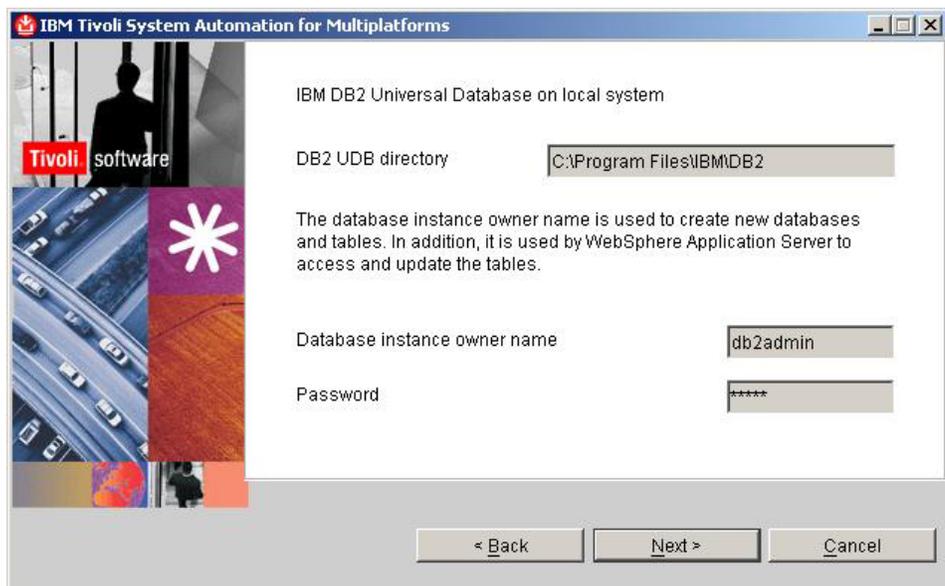
-
9. Select the DB2 setup type you are using and click **Next**.



-
10. Your actions on this panel depend on the type of DB2 setup you are using:
- **Local DB2 setup:** Specify the automation manager database name you want to use or accept the default name.
Click **Next** and proceed with step 11 on page 87.
 - **Remote DB2 setup:** Specify the name of the database you created for the automation manager.
Click **Next** and proceed with step 13 on page 88.

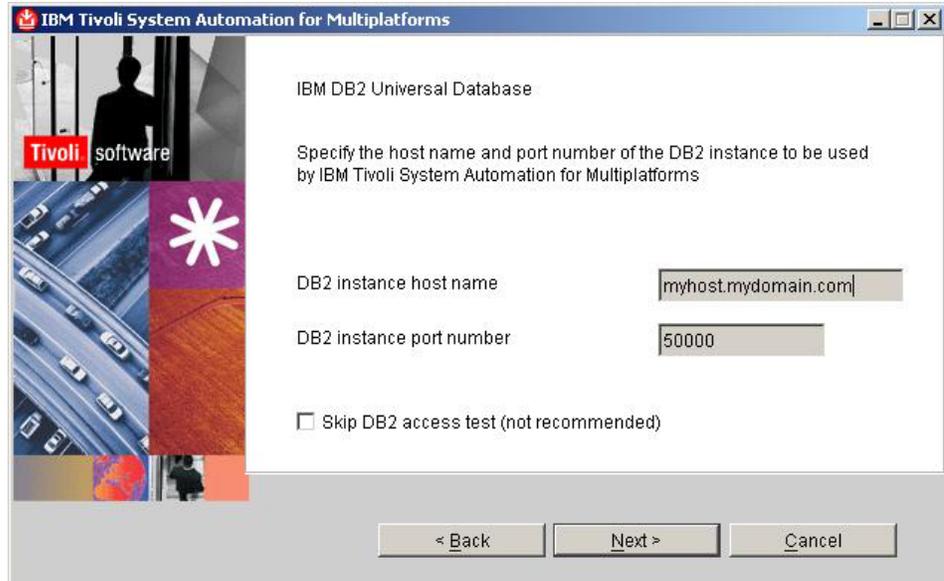


-
11. **This panel only appears when you are using a local DB2 setup.**
Specify the name and password of the DB2 instance owner and click **Next**.

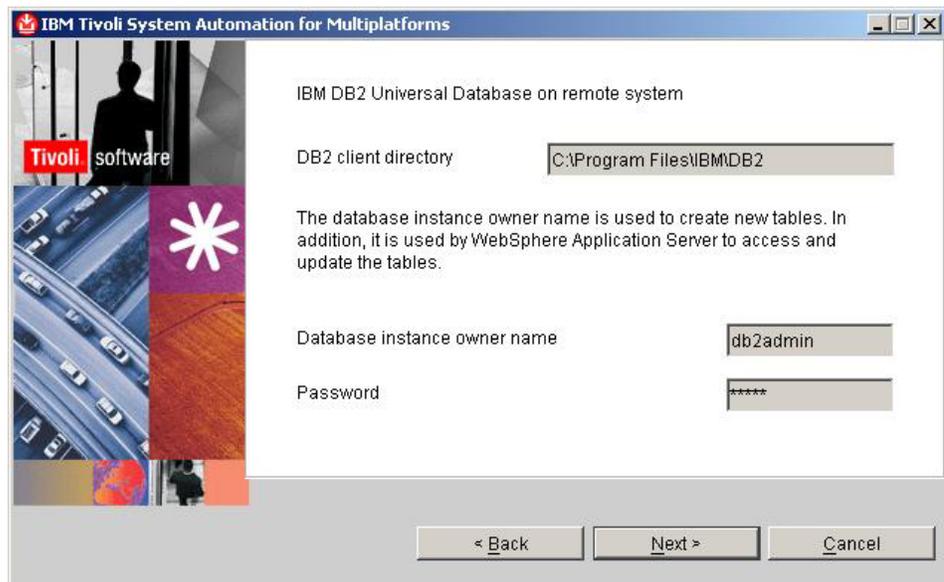


-
12. **This panel only appears when you are using a local DB2 setup.**
The fully qualified host name of your system was detected by the installation program. The value is displayed in the field **DB2 instance host name**. Do not change the value to localhost!
In the field **DB2 instance port number**, specify the port number you want to use or accept the default value.
Click **Next** and proceed with step 15 on page 89.

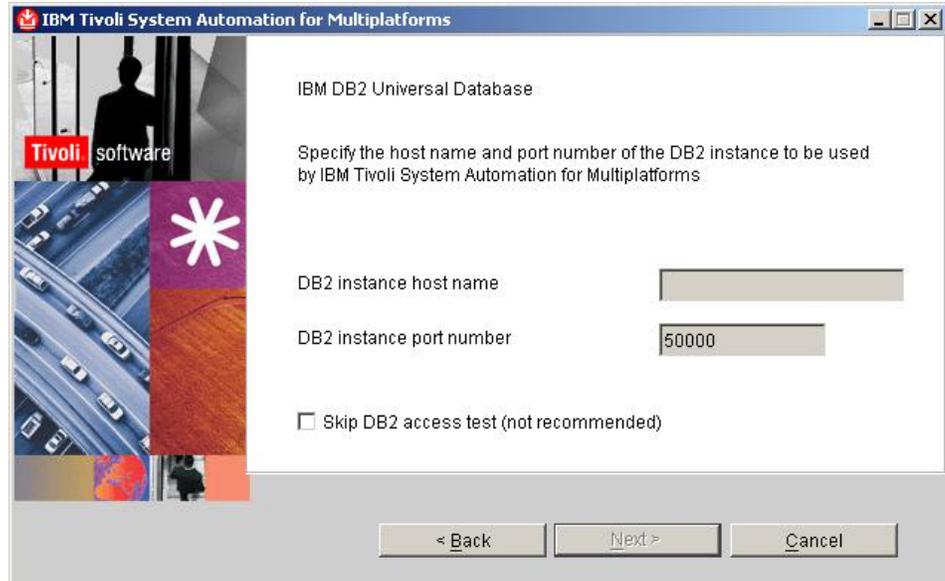
Note: After you click **Next**, the installation program checks whether the database can be accessed with the values you specified on the panel. If you want to skip the check, select the check box on the panel.



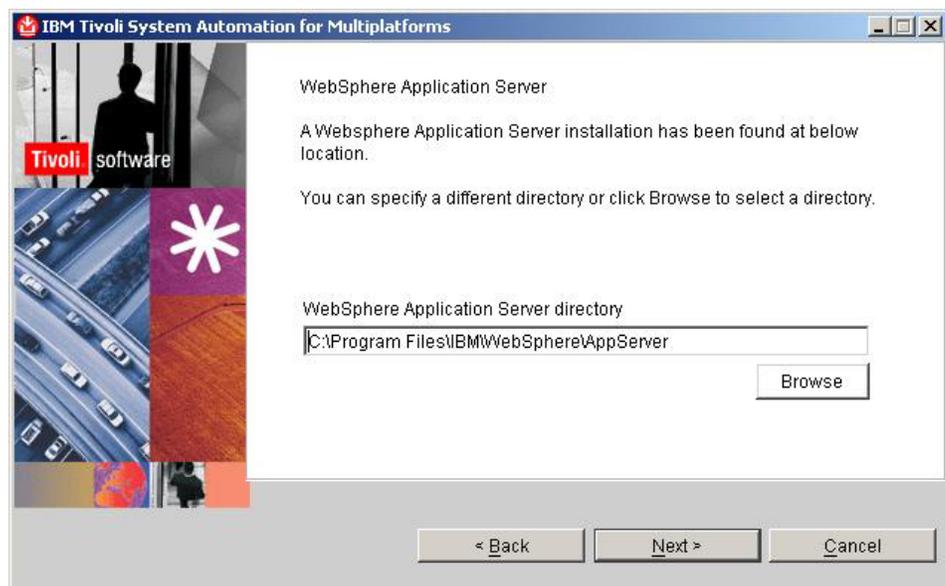
13. **This panel only appears when you are using a remote DB2 setup.** Specify the name and password of the database instance owner and click **Next**.



14. **This panel only appears when you are using a remote DB2 setup.** Specify the fully qualified host name of the system where the DB2 server is installed and the appropriate port number. Click **Next**.
- Note:** After you click **Next**, the installation program checks whether DB2 can be accessed with the values you specified on the panel. If you want to skip the check, select the check box on the panel.



15. The installation directory of WebSphere Application Server is detected on your system and displayed. Click **Next**.



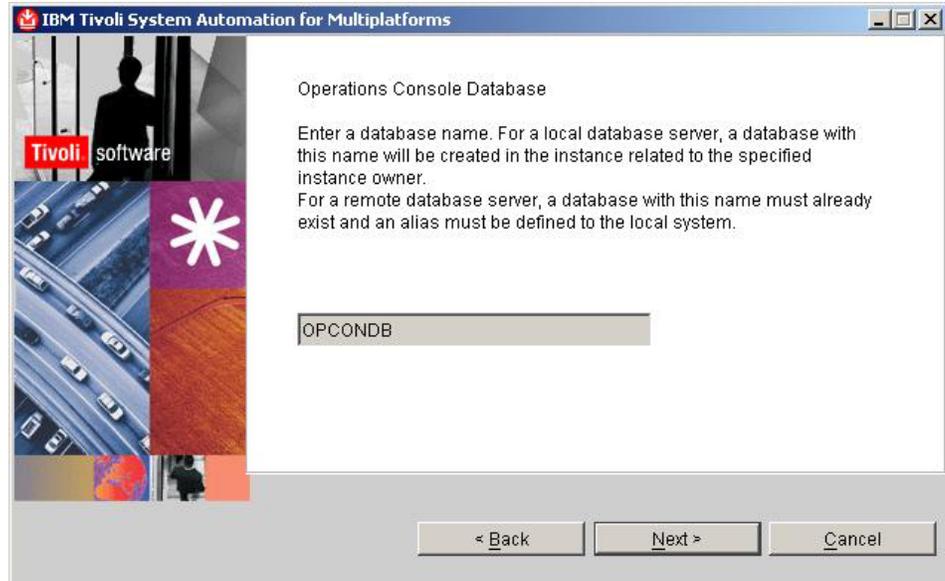
16. The existing WebSphere Application Server profiles are detected on your system and displayed. Select the profile you want to use and click **Next**.



-
17. Specify the WebSphere Application Server name you want to use or accept the default name. Click **Next**.

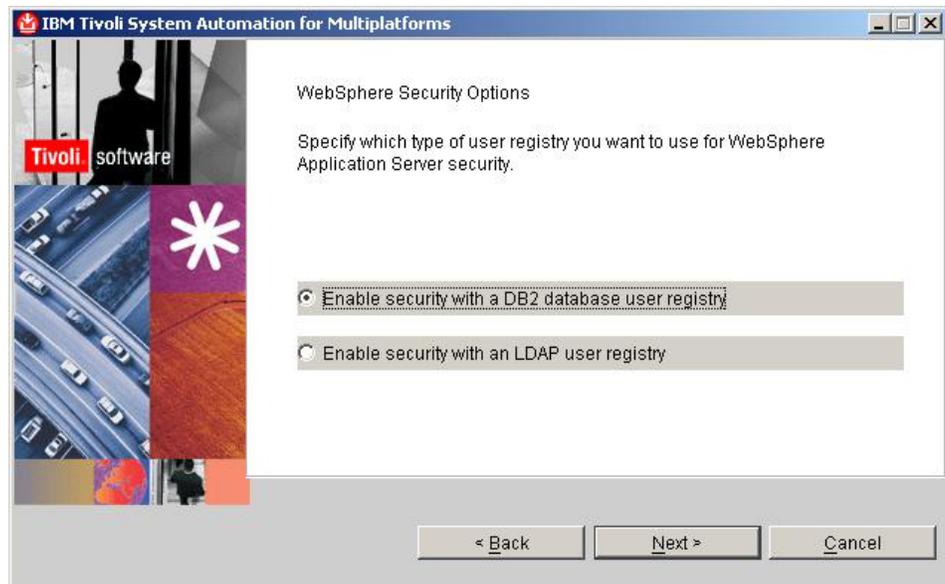


-
18. Your actions on this panel depend on the type of DB2 setup you are using:
- **Local DB2 setup:** Specify the operations console database name you want to use or accept the default name and click **Next**.
 - **Remote DB2 setup:** Specify the name of the database alias you cataloged for the operations console and click **Next**.



19. Select the type of user registry you are using:

- You are using a DB2 database as the user registry:
Select the appropriate radio button, click **Next**, and proceed with step 20.
- You are using an LDAP directory as the user registry:
Select the appropriate radio button, click **Next**, and proceed with step 21 on page 92.



20. **This panel only appears if you are using a DB2 database as the user registry.**

Choose a user ID and password for the Integrated Solutions Console administrator. If you specify an operations console installation directory other than the default, note that the length of the installation path must be 32

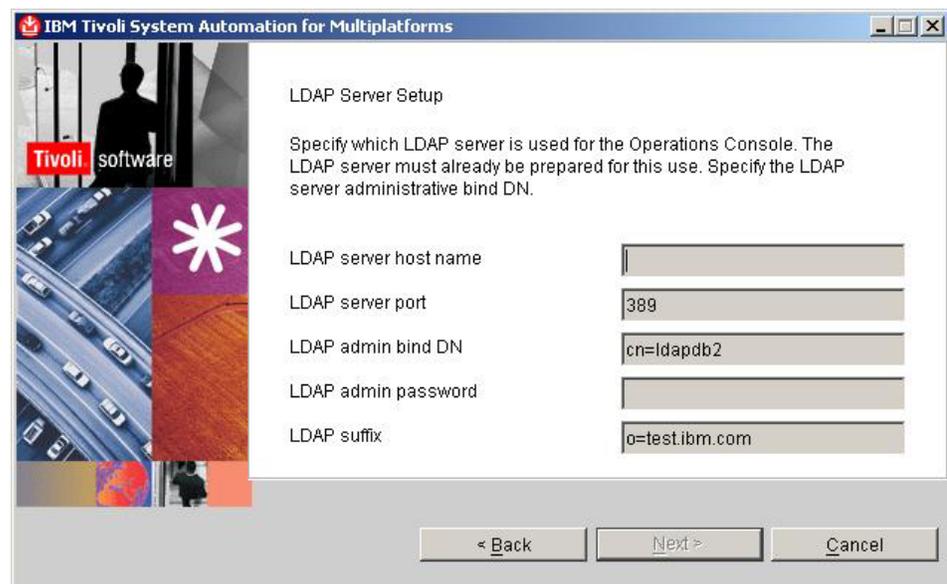
characters or less (for further restrictions, refer to Table 14 on page 72). Click **Next**, and proceed with step 26 on page 94.



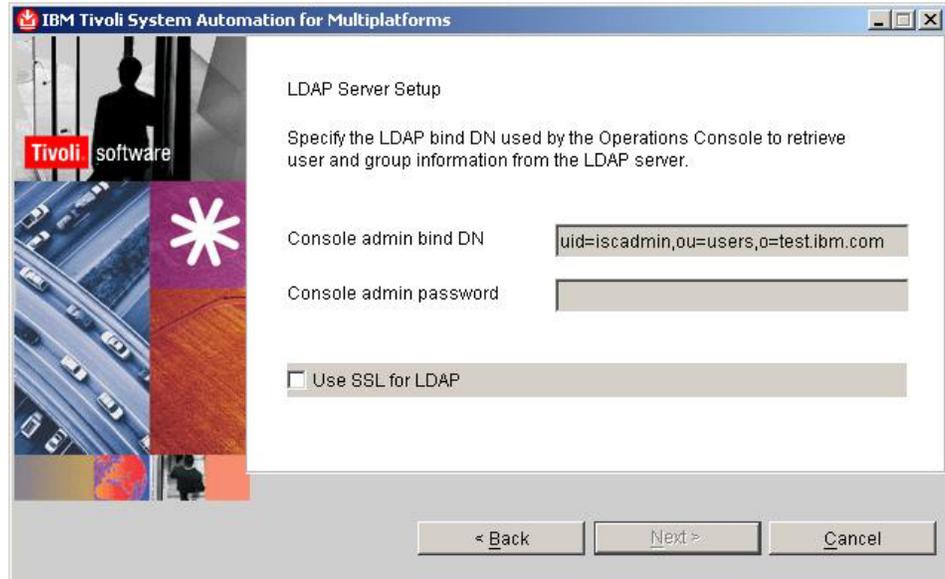
21. **This panel only appears if you are using LDAP as the user registry.** Specify the LDAP-specific data you have collected (refer to the table in "Installation parameters for LDAP" on page 74. In the table, the parameters are listed in the sequence in which you must specify them on this and the subsequent panels).

Note: On the panel, port number 389 is always displayed as default value for the LDAP server port, regardless of whether or not you are using SSL. If you will be using SSL, you must manually change the port number to the appropriate value.

Click **Next**.

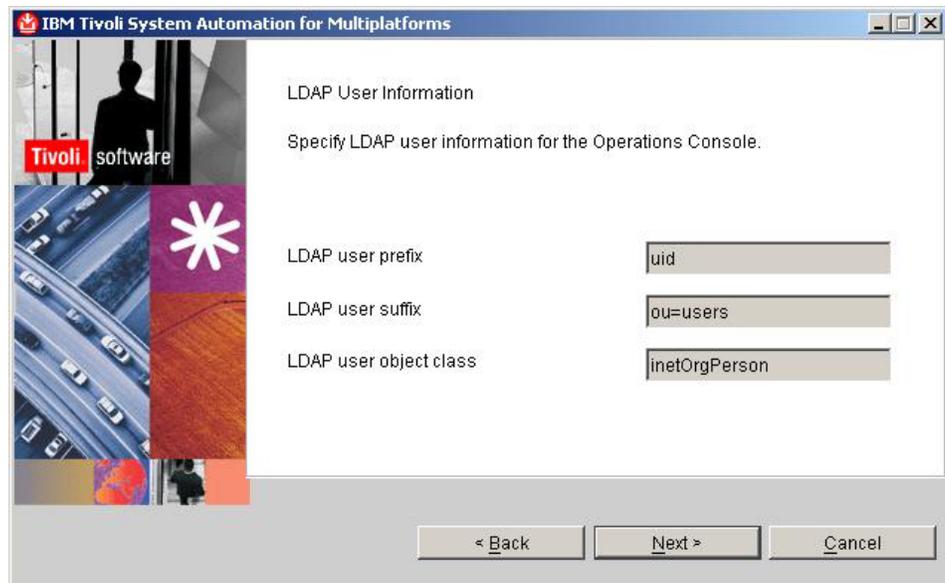


22. This panel only appears if you are using LDAP as the user registry. Specify the LDAP-specific data you have collected (refer to the table in “Installation parameters for LDAP” on page 74). Click **Next**.



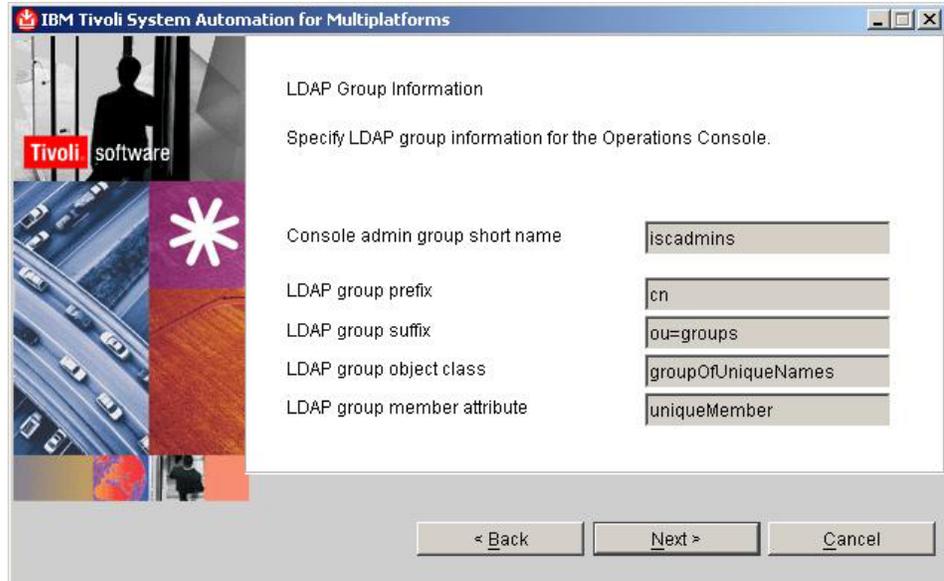
The screenshot shows a window titled "IBM Tivoli System Automation for Multiplatforms" with a sidebar on the left containing the Tivoli logo and various images. The main content area is titled "LDAP Server Setup" and contains the following text: "Specify the LDAP bind DN used by the Operations Console to retrieve user and group information from the LDAP server." Below this text are three input fields: "Console admin bind DN" with the value "uid=iscadmin,ou=users,o=test.ibm.com", "Console admin password" (empty), and a checkbox labeled "Use SSL for LDAP" which is unchecked. At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

23. This panel only appears if you are using LDAP as the user registry. Specify the LDAP-specific data you have collected (refer to the table in “Installation parameters for LDAP” on page 74). Click **Next**.



The screenshot shows a window titled "IBM Tivoli System Automation for Multiplatforms" with a sidebar on the left containing the Tivoli logo and various images. The main content area is titled "LDAP User Information" and contains the following text: "Specify LDAP user information for the Operations Console." Below this text are three input fields: "LDAP user prefix" with the value "uid", "LDAP user suffix" with the value "ou=users", and "LDAP user object class" with the value "inetOrgPerson". At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

24. This panel only appears if you are using LDAP as the user registry. Specify the LDAP-specific data you have collected (refer to the table in “Installation parameters for LDAP” on page 74). Click **Next**.



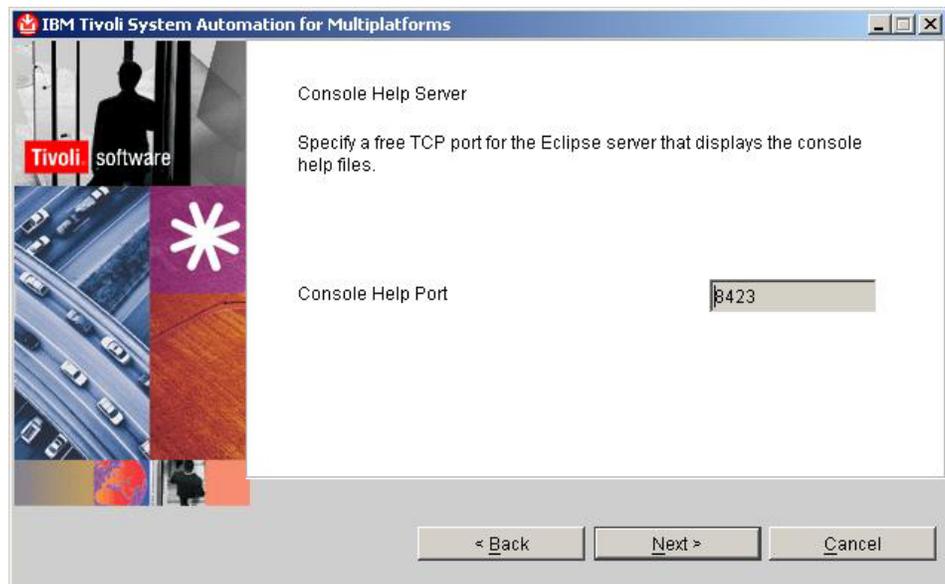
25. This panel only appears if you are using LDAP as the user registry. Specify the user ID and password of the Integrated Solutions Console administrator user. If you specify an operations console installation directory other than the default, note that the length of the installation path must be 32 characters or less (for further restrictions, refer to Table 14 on page 72). Click **Next**.



26. Specify the ports you want to use for the operations console or accept the default values (refer to “Installation parameters for the operations console” on page 76 for detailed information) and click **Next**.

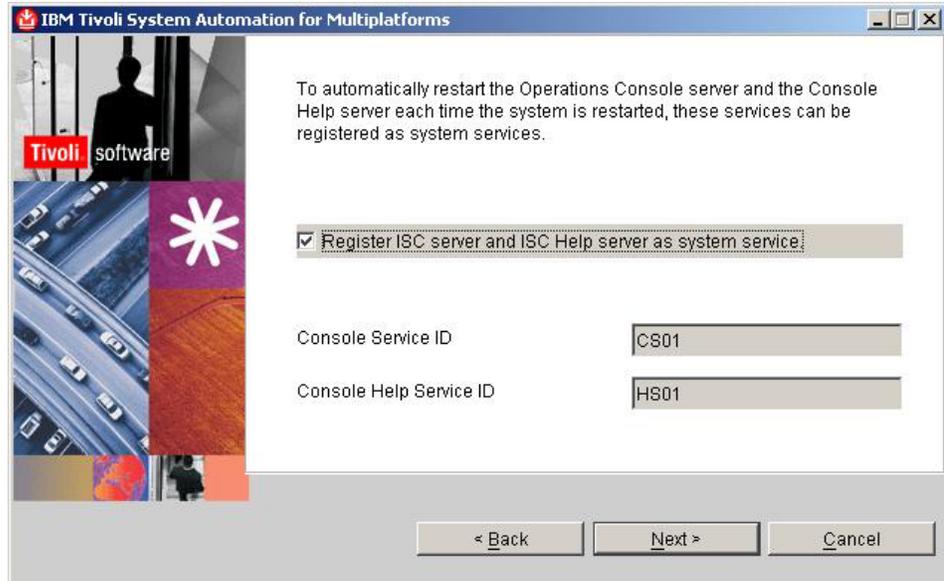


27. Specify the port number for the console help server or accept the default value and click **Next**.



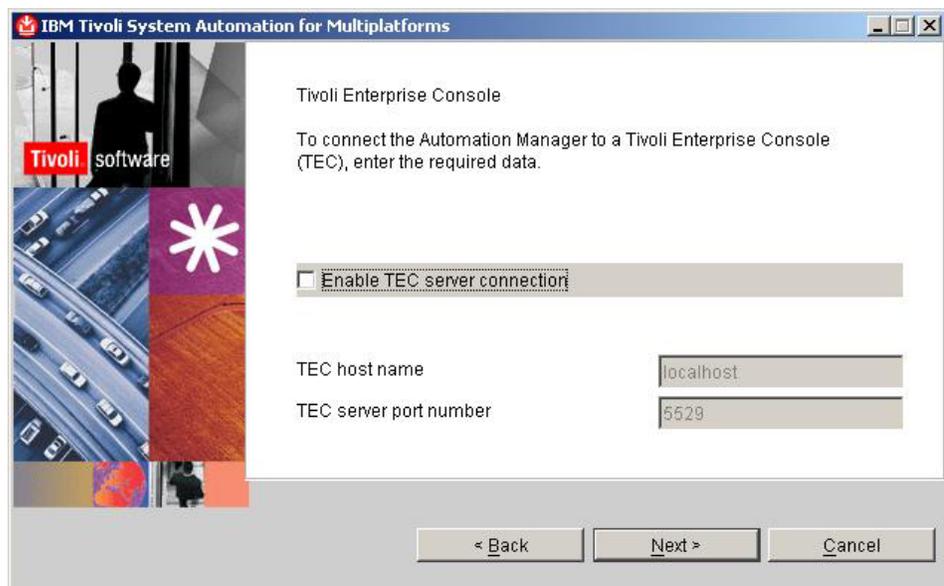
28. If you want to register the operations console server and the console help server as system services, specify service IDs or accept the default service IDs and click **Next**.

For more information, refer to “Installation parameters for the operations console” on page 76.



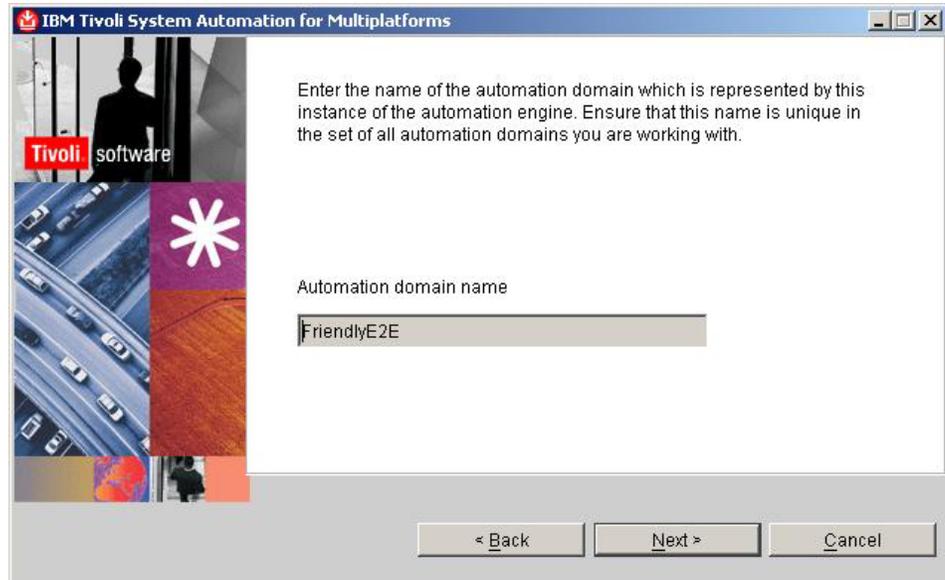
29. If you want to use Tivoli Enterprise Console to display end-to-end automation management events, select **Enable TEC server connection** and specify the host name and the port number of the console server. Click **Next**.

Note: You can also enable the connection using the WebSphere Application Server administrative console after the installation of the end-to-end automation management component is complete. This is described in Chapter 30, “Using Tivoli Enterprise Console with SA for Multiplatforms,” on page 199.

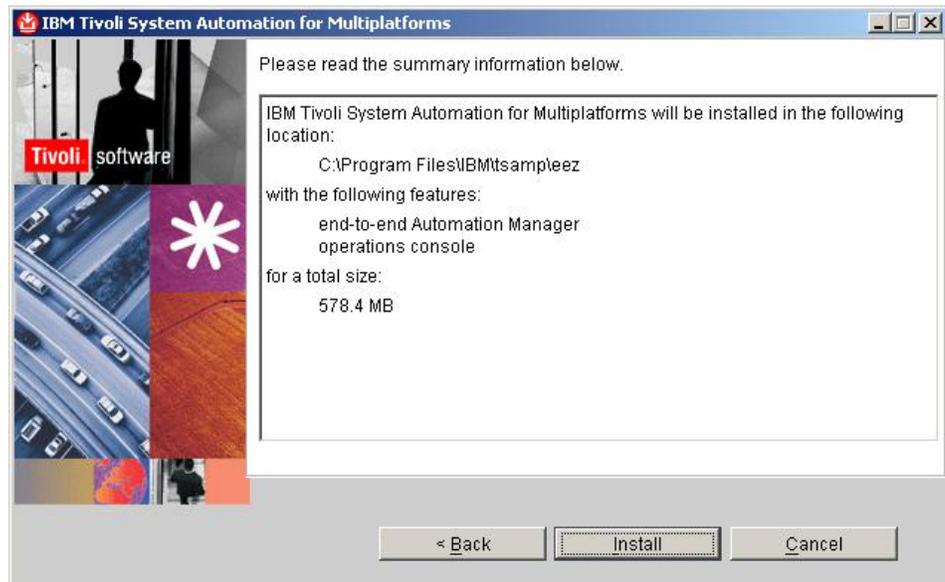


30. Specify the end-to-end automation domain name you want to use or accept the default name and click **Next**.

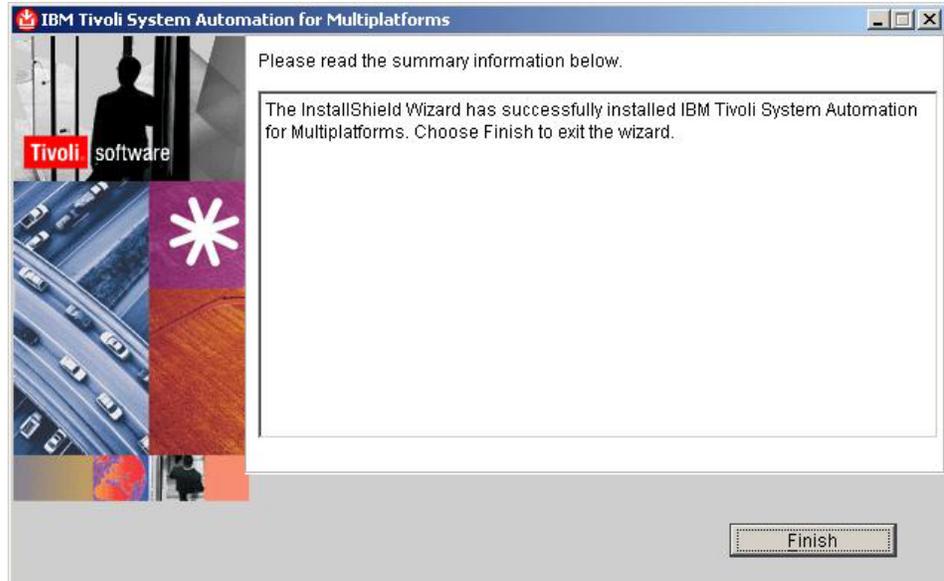
Note: Accept the default domain name (“FriendlyE2E”) if you want to use the sample end-to-end automation management environment to familiarize yourself with end-to-end automation management and the operations console. For more information, refer to Part 3, “First steps,” on page 125.



31. When you have specified all required information on the wizard panels, a summary panel appears. Click **Install**. The installation wizard begins installing the end-to-end automation management component. Note that the installation can take considerable time to complete.



32. When the end-to-end automation management component was installed successfully, a summary panel appears. Click **Finish** to close the installation wizard. For information on verifying the installation, refer to Chapter 11, “Verifying the installation,” on page 99.



Note: If problems were encountered during the installation, an error panel appears before the summary panel is displayed. In such a case, do this:

- On the error panel, click **Next**.
- On the summary panel that appears, click **Finish**.
- Use the log files that were created by the installation program to analyze and resolve the problems. For more information on the log files, refer to "Using the installation log files" on page 119.

Chapter 11. Verifying the installation

Automation manager

To verify that the automation manager was installed successfully, complete the tasks described in the following sections.

End-to-end automation database

Perform these steps to verify that the end-to-end automation database and the database tables were created successfully:

1. Ensure that DB2 is running.
2. Open the DB2 Control Center.
3. Navigate to Databases and expand the folder.
4. Expand EAUTODB.
5. Click Tables. The following database tables must be listed:
 - EEZAUTOMATIONACCESS
 - EEZAUTOMATIONRELATION
 - EEZDOMAINSUBSCRIPTION
 - EEZOPERATORDOMAINFILTER
 - EEZOPERATORDOMAINPREFERENCES
 - EEZOPERATORHIDDENDOMAIN
 - EEZRESOURCESUBSCRIPTION

Automation J2EE Framework

Perform these steps to verify that the automation J2EE framework was installed successfully:

1. In a Web browser window, specify the address `http://<your_host_name>:<your_was_admin_console_port>/admin` to display the WebSphere Application Server administrative console.
Typically, the default port number of the WebSphere Application Server administrative console is 9060.
2. On the login panel, enter the user ID and password of the Integrated Solutions Console administrator.
3. Navigate to Applications —> Enterprise Applications. The list of installed applications must contain the following entries:
 - EventServer
 - EventServerMdb
 - EEZEAR

Verifying that DB2 accepts WebSphere Application Server requests

Perform the following task to verify that DB2 accepts WebSphere Application Server requests:

1. In a Web browser window, specify the address `http://<your_host_name>:<your_was_admin_console_port>/admin` to display the WebSphere Application Server administrative console.

Typically, the default port number of the WebSphere Application Server administrative console is 9060.

2. On the login panel, enter the user ID and password of the Integrated Solutions Console administrator.
3. Navigate to **Resources** → **JDBC providers** → **DB2 Universal JDBC Driver (XA)** → **Data sources** → **EAUTODBDS**. Click **Test connection** to verify that DB2 accepts WebSphere Application Server requests. If the test is successful, the following message comes up:

Test connection for data source EAUTODBDS was successful.

If the test fails, check if the DB2 port number specified for **EAUTODB** is correct (for more information, refer to “WebSphere Application Server cannot connect to DB2” on page 305).

Automation engine

Perform these steps to verify that the automation engine was installed successfully:

1. Open a command prompt and change to the directory `<EEZ_INSTALL_ROOT>/bin`.
2. Issue the command `eezdmn -?`. When the installation of the automation engine was successful, the list of available command options is displayed.

Note: You can also use any of the other `eezdmn` command options to verify the installation of the automation engine. As long as you do not receive an exception, any message you receive verifies that the automation engine is installed correctly. For a complete list of the `eezdmn` command options, see Chapter 27, “Using the command-line interface of the automation engine,” on page 183.

Operations console

Perform the following steps to verify that the operations console feature was installed successfully:

1. In a Web browser window, specify the address `http://<your_host_name>:<your_isc_port>/ibm/console` to display the Login panel of Integrated Solutions Console. The default ISC port is 8421.

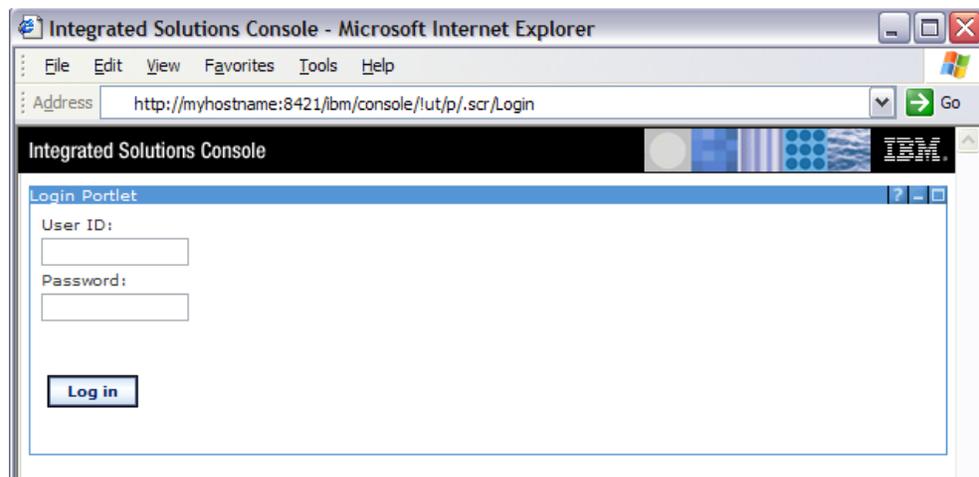


Figure 5. Log in panel of Integrated Solutions Console

2. Type the Integrated Solutions Console user ID and password that you specified during the installation and click **Log in**. The Welcome panel is displayed. On the Work Items page, you should see these entries:
 - Integrated Solutions Console
 - Tivoli System Automation for Multiplatforms

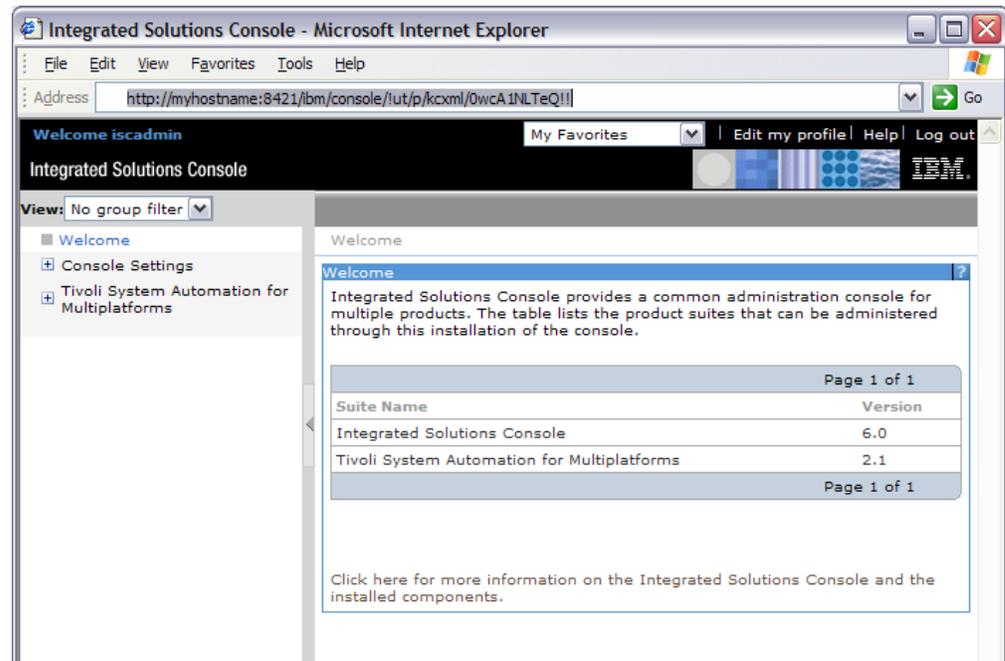


Figure 6. Welcome panel of Integrated Solutions Console

3. Expand the folder Tivoli System Automation for Multiplatforms.
4. Click **SA operations console**. When the panel for connecting to the operations console is displayed, the installation of the operations console was successful.

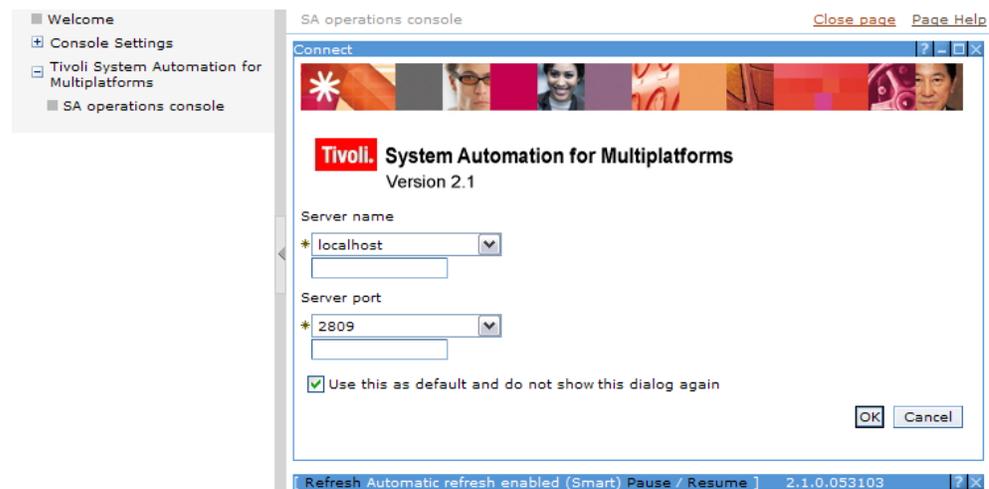


Figure 7. Panel for connecting to the operations console

5. To verify that communication can be established between the operations console and the automation J2EE framework, click **OK** on the Connect panel. The verification is successful when the operations console appears.

Chapter 12. Post-installations tasks

When you have verified the installation of the end-to-end management component, you need to perform a number of post-installation tasks:

- You must perform the post-installation tasks described in the following chapter:
 - Chapter 23, “Post-installation tasks for administrators,” on page 143
- You must enable the end-to-end automation manager to access the first-level automation domains.

To do this, you must specify the user credentials for the first-level domains on the User credentials page of the configuration dialog.

“Launching the configuration dialog” on page 193 describes how you launch the configuration dialog. For detailed information about the User credentials page, refer to the online help of the configuration dialog.

- Before users can work with the SA for Multiplatforms operations console, you must complete the following tasks regardless of whether the users will be using the operations console in end-to-end automation mode or first-level automation mode:
 - Create and authorize users. This is described in “Creating and authorizing users in Integrated Solutions Console” on page 153.
- To get end-to-end automation management operational, you must complete the following tasks:
 - Create and activate an automation policy. This is described in Chapter 25, “Creating and modifying policies,” on page 161 and “Activating a policy” on page 249.

Setting up SSL for the operations console

This is an optional post-installation task to allow for secure HTTP transport (HTTPS) between client browsers and the operations console.

To set up SSL for Integrated Solutions Console, perform the following procedure:

1. Configure the Web server to support HTTPS. If you are doing this in a production environment, you need to obtain a certificate from a certificate authority. For testing purposes, you can use IKEYMAN to generate a self-signed certificate. Refer to the WebSphere Application Server documentation for detailed instructions.

2. Add the virtual host defined in the Web server to the Web server virtual host alias list. Add a host alias for the SSL port that the Web server uses. To create the settings, perform the following steps:

- a. Use a text editor to open the following file:

```
<was_root>/profiles/default/config/cells/<cell>/virtualhosts.xml
```

where <was_root> is the WebSphere Application Server installation directory, and <cell> is the name of the WebSphere Application Server cell for the operations console installation.

- b. Add the following element before the ending element `</host:VirtualHost>` for the virtual host named `default_host`. Add the element to the list of aliases:

```
<aliases xmi:id="HostAlias_x" hostname="*" port="alias_port"/>
```

where `x` is the next number in the `HostAlias` sequence and `alias_port` is the value specified for the `HTTPS Port` parameter.

- c. Save the file.
-

3. Edit the `ConfigService.properties` file in `isc_runtime_root/PortalServer/shared/app/config/services`

where `isc_runtime_root` is the operations console installation directory.

Change the following parameters:

```
redirect.login.ssl = true
redirect.logout.ssl = true
host.port.https = alias_port
```

where `alias_port` is the port number used for the virtual host alias that you specified in step 2 on page 103.

4. Set the security constraints for the console URL. To do so, perform the following steps:

- a. Use an editor to open the file `web.xml`. It is located in the following directory:

```
<was_root>/profiles/default/config/cells/<cell>/
  applications/wps.ear/deployments/wps/wps.war/WEB-INF
```

where `<was_root>` is the WebSphere Application Server installation directory, and `<cell>` is the name of the WebSphere Application Server cell for the operations console installation.

- b. In the `web.xml` file, change the `<security-constraint>` element for the console URL to use HTTPS as shown in the following example:

```
<security-constraint id="SecurityConstraint_1">
  <web-resource-collection id="WebResourceCollection_1">
    <web-resource-name></web-resource-name>
    <url-pattern>/console/*</url-pattern>
    <http-method>DELETE</http-method>
    <http-method>GET</http-method>
    <http-method>POST</http-method>
    <http-method>PUT</http-method>
  </web-resource-collection>
  <auth-constraint id="AuthConstraint_1">
    <description></description>
    <role-name>All Role</role-name>*gt;
  </auth-constraint>
  <user-data-constraint id="UserDataConstraint_4">
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
    <!-- replace NONE with CONFIDENTIAL -->
  </user-data-constraint>
</security-constraint>
```

5. From a command prompt, change to the following directory:
-

```
<isc_runtime_root>/PortalServer/config
```

where <isc_runtime_root> is the operations console installation directory.

6. Edit the file `wconfig.properties` and set the `PortalAdminPwd` parameter to your console administrator user ID's password.
-

7. Run the following command from the command line:

```
WPSconfig.bat|sh init action-deploy-setupisc-ssl
```

8. From a command prompt, change to the following directory:

```
<isc_runtime_root>/PortalServer/bin
```

where <isc_runtime_root> is the operations console installation directory.

9. Run the following command from the command line:

Windows:

```
stopISC.bat ISC_Portal iscadmin iscpass
```

AIX, Linux:

```
./stopISC.sh ISC_Portal iscadmin iscpass
```

where `iscadmin` is the console administrator user ID and `iscpass` is the console administrator user ID's password.

10. Run the following command from the command line:

Windows:

```
startISC.bat ISC_Portal
```

AIX, Linux:

```
./startISC.sh ISC_Portal
```

11. Test your changes by launching the home page of the console in a Web browser. The login page that is displayed is not secure. However, when you click **Log in**, the credentials are encrypted and the session is directed to a secure connection.
-

Modifying the LTPA settings

After the installation of the end-to-end automation management component, you should check whether the LTPA settings are appropriate for your environment.

During installation, the following LTPA parameters are automatically set in WebSphere Application Server:

- LTPA Password is set to the password of the Integrated Solutions Console administrator user ID
- LTPA Timeout is set to 120 minutes

LTPA Timeout is a security-related timeout. Because this timeout is absolute, a user will be logged out and forced to log in to Integrated Solutions Console again when the LTPA timeout is reached even if the user is working with the operations console at the time.

To change the LTPA settings (for example, password and timeout) you use the WebSphere Application Server administrative console. On the administrative console, select **Security** —> **Global Security** —> **Authentication mechanism** —> **LTPA**.

Modifying the HTTP session timeout

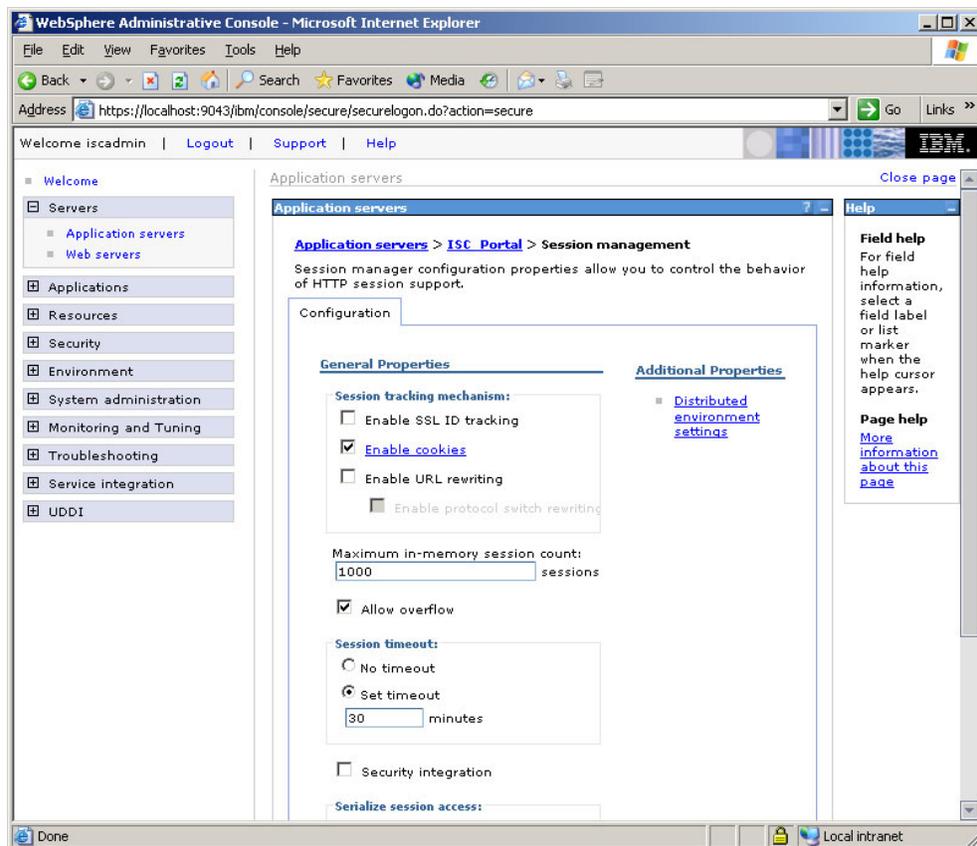
After the installation of the end-to-end automation management component, you should check whether the setting of the HTTP session timeout is appropriate for your environment.

The HTTP session time is an activity timeout. The value to which the HTTP session timeout is set defines after how many minutes of inactivity a user is logged out automatically.

During installation, the HTTP session timeout is set to 30 minutes.

To change the value, you use the WebSphere Application Server administrative console.

On the administrative console, select **Servers** —> **Application servers** —> **ISC_Portal** —> **Web container settings** —> **Session Management**:



Configuring how many users can connect to the automation manager using the operations console

During the installation of the end-to-end automation management component, a default value is set that defines how many users can simultaneously connect to the automation manager using the operations console. You can change the current setting by changing the **Maximum connections** value for the EEZTopicConnectionFactory from the WebSphere Application Server administrative console (**Resources** → **Resource adapters** → **SIB JMS Resource Adapter** → **J2C connection factories** → **EEZTopicConnectionFactory** → **Connection pool properties**).

If **Maximum connections** is set to 0, the number of concurrent connections that can be established is allowed to grow infinitely. If the specified number of maximum connections has been reached, the next connection attempt using the operation console will show the following error message:

```
EEZU0011E:  
Unable to set up the event path between the operations console  
and the management server:  
CWSIAD005E: The JCA runtime failed to allocate a connection.
```

Chapter 13. Uninstalling the end-to-end automation management component

This section describes how to uninstall the end-to-end automation management component. An uninstallation program is provided that removes the components that were installed by the installation wizard.

Launching the graphical uninstallation program on Windows

To launch the uninstallation program on Windows, you can either issue the command `<EEZ_INST_ROOT>/_uninst/uninstaller.exe` at a command prompt or perform the following steps:

1. Open the Control Panel (**Start** → **Settings** → **Control Panel**).
2. On the Control Panel, open **Add/Remove Programs**.
3. On the Add/Remove Programs panel, select **End-to-End Automation Management component of Tivoli System Automation for Multiplatforms** and click **Change/Remove**. This brings up the Welcome panel of the uninstallation program.

Launching the graphical uninstallation program on AIX and Linux

To launch the uninstallation program on AIX and Linux, enter the following command in a shell:

```
<EEZ_INST_ROOT>/_uninst/uninstaller.bin
```

This brings up the Welcome panel of the uninstallation program.

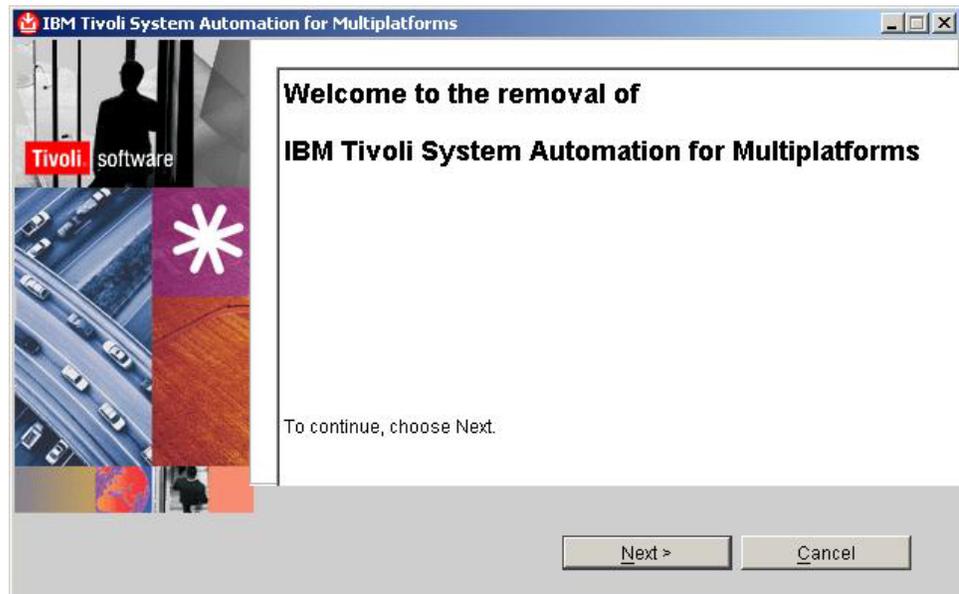
Using the uninstallation program

Before you begin:

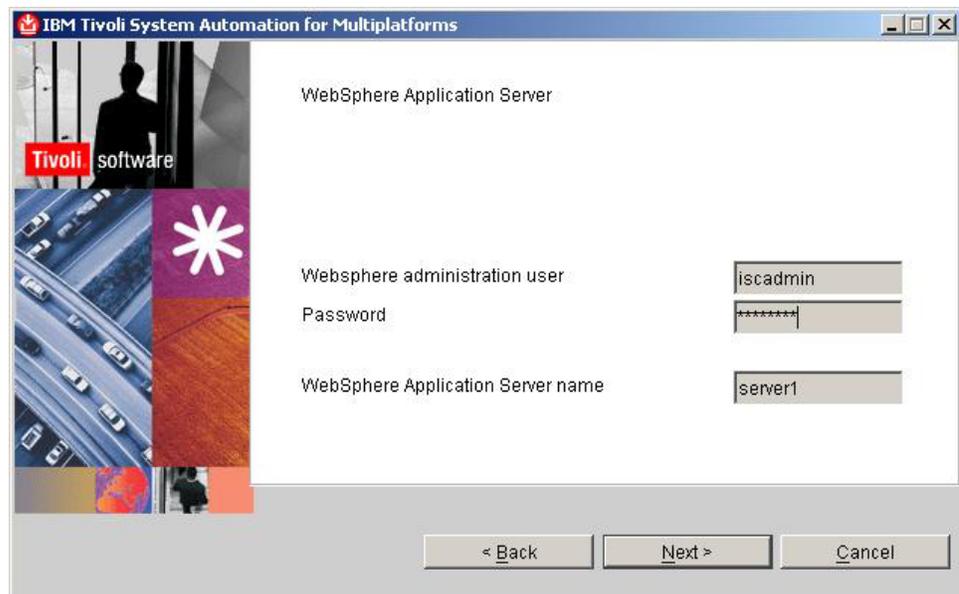
- Before starting the uninstallation of the end-to-end automation management component, make sure that the automation engine, the Integrated Solutions Console server, and the WebSphere Application Server "server1" are stopped. For information on how to stop the components, refer to the following sections:
 - To stop the automation engine, refer to Chapter 27, "Using the command-line interface of the automation engine," on page 183.
 - To stop the Integrated Solutions Console server, refer to "Starting and stopping the operations console" on page 190.
 - To stop WebSphere Application Server, refer to "Starting and stopping WebSphere Application Server" on page 189.
- During uninstallation, a number of panels may appear prompting you to confirm that specific files are to be deleted. Make sure that the files should be deleted before confirming the deletion.

Perform the following steps to uninstall the end-to-end automation management component:

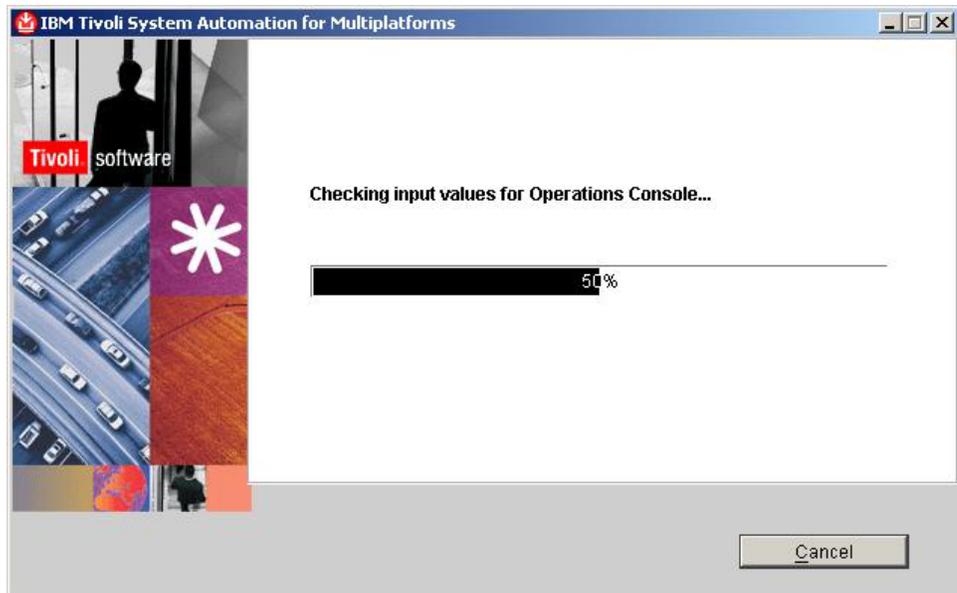
1. Launch the uninstallation program as described in the sections above.
2. On the Welcome panel of the uninstallation program, click **Next**.



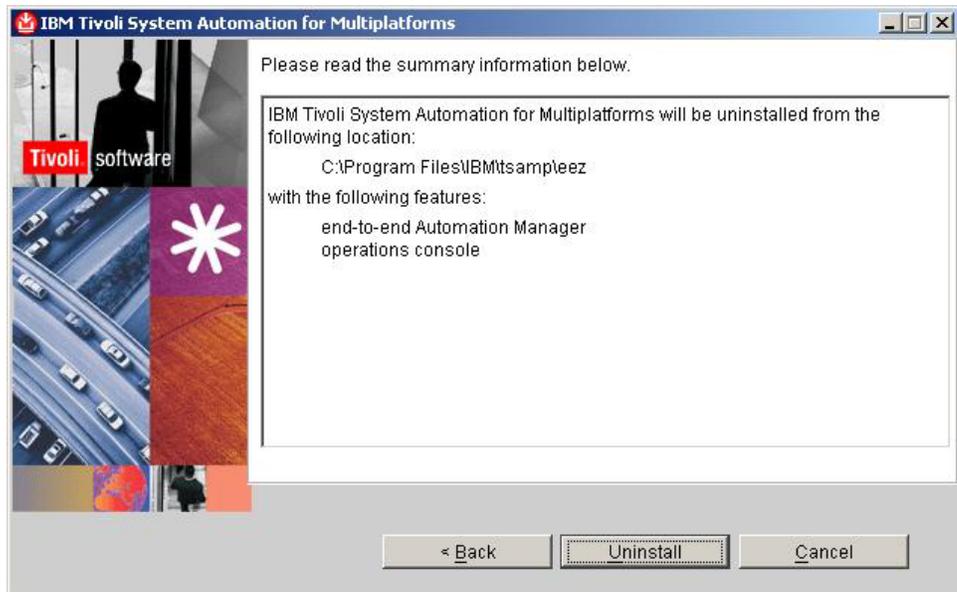
3. In the fields **WebSphere administration user** and **Password**, type the user ID and password of the Integrated Solutions Console administrator user. Click **Next**.



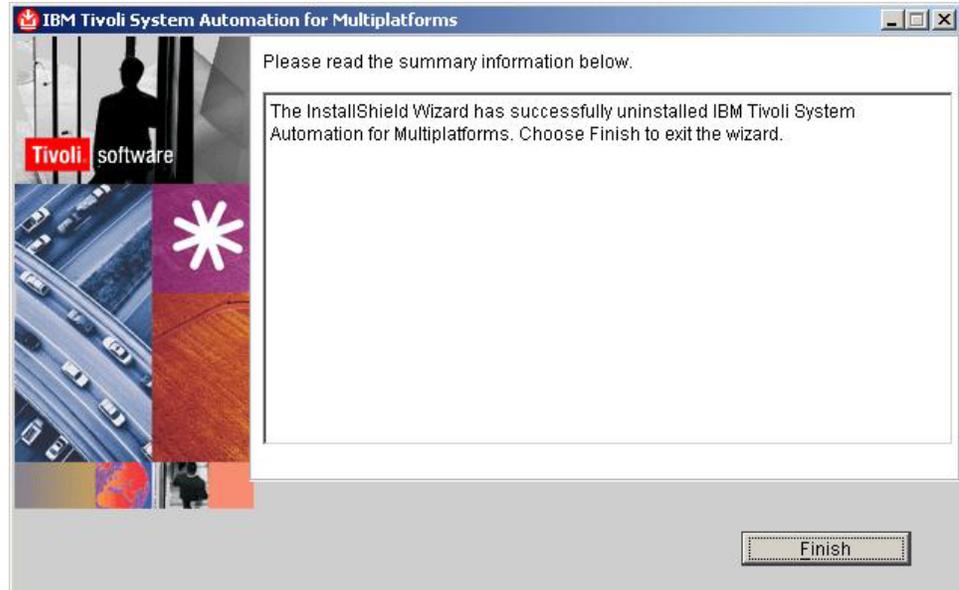
4. Some information panels are displayed while the uninstallation program checks your system for the information it needs for the uninstall. The following figure shows an example.



5. When all required information has been detected, a confirmation panel appears. To start the uninstallation, click **Uninstall**.
Note that the uninstallation can take considerable time to complete. Use the progress indicators that are displayed throughout the uninstallation to monitor the progress.



6. When the uninstallation is complete, a summary panel is displayed. On the panel, click **Finish**.



Note: If problems were encountered during the uninstallation, an error panel appears before the summary panel is displayed. In such a case, do this:

- On the error panel, click **Next**.
 - On the summary panel that appears, click **Finish**.
 - Use the log files that were created by the uninstallation program to analyze and resolve the problems. For more information on the log files, refer to "Using the installation log files" on page 119.
-

Chapter 14. Installing service

Installing service means applying corrective service fix packs to release 2.1.0 of IBM Tivoli System Automation for Multiplatforms or upgrading the software release level from release 2.1.0. In this documentation, the service fix packs that you use for updating the end-to-end automation management component are referred to as product fix packs.

Note: For some product fix packs, specific interim fixes or fix packs for WebSphere Application Server are required. In such a case, these fixes are available at the location from which you download the product fix pack. They must be installed **before** the product fix pack is installed. Detailed instructions for installing the fixes are provided in the release notes.

Do not install any WebSphere Application Server interim fixes or fix packs that are not mentioned in the release notes unless you are explicitly advised to do so by Tivoli System Automation support.

Product fix packs and interim fixes are delivered as:

- Self-extracting archives for Windows and AIX
- Archives in TAR-format for Linux

Where to obtain fix packs

Read the release notes to find out which fix packs are required for a release update. The release notes are available on the IBM Tivoli System Automation home page at:

www.ibm.com/software/tivoli/products/sys-auto-linux/

On the page, click **Technical Documentation** to display the list of available documentation.

The archives can be downloaded from the IBM Tivoli System Automation support site at:

www.ibm.com/software/sysmgmt/products/support/IBMTivoliSystemAutomationforLinux.html

Archive naming conventions

Naming convention for product fix pack archives:

2.1.0-TIV-SAE2E-<platform>-FP<fix_pack_number>.<archive_type>

Naming convention for WebSphere Application Server interim fix archives:

2.1.0-TIV-SAWAS-<platform>-FP<fix_pack_number>.<archive_type>

where

- <platform> represents the platform on which the end-to-end automation management component is installed
- <fix_pack_number> represents the fix pack number
- <archive_type> represents the platform-specific file extension of the archive

Example:

This is the tar archive that is used to install product fix pack 2 for IBM Tivoli System Automation for Multiplatforms 2.1 on PPC Linux platforms:

2.1.0-TIV-SAE2E-PPC-FP0002.tar

Naming conventions of the update installer location

The location at which you find the update wizard program for installing the product fix pack after unpacking an archive has the following syntax:

EEZ21<mf>E2E<platform>/<platform>/<update_wizard_file>

where

- <mf> represents modification level and fix level. For example, for fix pack 2110, the directory is named EEZ2110.
- <platform> represents the platform on which the end-to-end automation management component is installed
- <update_wizard_file> represents the update wizard program you use to install the product fix pack

Example:

This is where you find the update wizard after the Linux PPC archive for fix pack 2 for SA for Multiplatforms 2.1 is unpacked:

EEZ2110E2EPPC/ppc/update

Usage instructions for the platform-specific archives

These are the archives for applying service to the end-to-end automation management component.

Windows

Table 22. Windows platforms

| Archive name | Description |
|---|--|
| 2.1.0-TIV-SAE2E-WIN-FP<fix_pack_number>.exe | The archive is self-extracting. This is where you find the update installer program after unpacking the product fix pack archive: EEZ21<mf>E2EWindows/Windows/update.exe |
| 2.1.0-TIV-SAWAS-WIN-FP<fix_pack_number>.exe | The archive is self-extracting. For information about installing WebSphere Application Server interim fixes, refer to the release notes. |

AIX

Table 23. AIX platforms

| Archive name | Description |
|---|--|
| 2.1.0-TIV-SAE2E-AIX-FP<fix_pack_number>.bin | The archive is self-extracting. This is where you find the update installer program after unpacking the product fix pack archive: EEZ21<mf>E2EAIX/AIX/update |

Table 23. AIX platforms (continued)

| Archive name | Description |
|---|---|
| 2.1.0-TIV-SAWAS-AIX-FP<fix_pack_number>.bin | The archive is self-extracting. For information about installing WebSphere Application Server interim fixes, refer to the release notes. |

Linux on IBM x/Series

Table 24. Linux on IBM x/Series

| Archive name | Description |
|--|---|
| 2.1.0-TIV-SAE2E-I386-FP<fix_pack_number>.tar | For extracting the archive, GNU tar 1.13 or later is required. Use the tar -xf command to extract the files. This is where you find the update installer program after unpacking the product fix pack archive: EEZ21<mf>E2EI386/i386/update |
| 2.1.0-TIV-SAWAS-I386-FP<fix_pack_number>.tar | For extracting the archive, GNU tar 1.13 or later is required. Use the tar -xf command to extract the files. For information about installing WebSphere Application Server interim fixes, refer to the release notes. |

PPC Linux

Table 25. PPC Linux

| Archive name | Description |
|---|---|
| 2.1.0-TIV-SAE2E-PPC-FP<fix_pack_number>.tar | For extracting the archive, GNU tar 1.13 or later is required. Use the tar -xf command to extract the files. This is where you find the update installer program after unpacking the product fix pack archive: EEZ21<mf>E2EPPC/ppc/update |
| 2.1.0-TIV-SAWAS-PPC-FP<fix_pack_number>.tar | For extracting the archive, GNU tar 1.13 or later is required. Use the tar -xf command to extract the files. For information about installing WebSphere Application Server interim fixes, refer to the release notes. |

Linux on z/Series

Table 26. Linux on z/Series

| Archive name | Description |
|--|---|
| 2.1.0-TIV-SAE2E-S390-FP<fix_pack_number>.tar | For extracting the archive, GNU tar 1.13 or later is required. Use the tar -xf command to extract the files. This is where you find the update installer program after unpacking the product fix pack archive: EEZ21<mf>E2ES390/s390/update |
| 2.1.0-TIV-SAWAS-S390-FP<fix_pack_number>.tar | For information about installing WebSphere Application Server interim fixes, refer to the release notes. |

Steps for installing a product fix pack

Before you begin:

- Product fix packs are always cumulative.
- Release 2.1.0 must be installed before any product fix pack can be installed.
- To install a product fix pack, you must have root authority.

To install a product fix pack, perform the following steps:

1. Check the release notes to find out which archives are required.
2. Download the archives from the SA for Multiplatforms support site:
 - Archives for WebSphere Application Server fixes:
Follow the download instructions provided in the release notes.
 - Archives for product fix packs:
Typically, one archive is provided for each platform. Download the archive to a temporary directory.
3. If fixes for WebSphere Application Server must be installed, unpack and install the fixes as described in the release notes.
4. Unpack the product fix pack archive to a temporary directory. For information about how to unpack the archive for your platform, refer to “Usage instructions for the platform-specific archives” on page 114.
5. Before performing the subsequent steps, check the release notes for additional or deviating installation instructions.
6. Change to the directory in which the update wizard program is located. For information on where to find the update wizard program, refer to “Usage instructions for the platform-specific archives” on page 114.
7. Launch the update wizard.
When the wizard is launched successfully, the Welcome panel appears.
8. Follow the instructions on the wizard panels to install the product fix pack.

Chapter 15. Uninstalling service

Uninstalling service means that you have to uninstall the complete end-to-end automation management component as described in Chapter 13, “Uninstalling the end-to-end automation management component,” on page 109. After the uninstall procedure is complete, you need to reinstall the component and install the required service level (fix pack level).

Chapter 16. Troubleshooting installation problems

Installation wizard cannot find WebSphere Application Server on the system

The installation program for SA for Multiplatforms uses InstallShield to install code. InstallShield uses the `vpd.properties` file to track products that it installs.

WebSphere Application Server also uses InstallShield. WebSphere Application Server must be installed on the system before end-to-end automation management can be installed. If the installation wizard of Tivoli System Automation for Multiplatforms does not detect WebSphere Application Server on the system but you did install it, check whether the file `vpd.properties` exists and correct its contents.

This is where you find the `vpd.properties` file:

- **AIX:**
In the root directory or in the directory `/usr/lib/objrepos`
- **Linux:**
In the root directory
- **Windows**
In installation directory of the operating system, for example, `C:\WINNT` or `C:\windows`

If you cannot locate the file, your WebSphere Application Server installation may be damaged, in which case you need to re-install WebSphere Application Server. For more information about the `vpd.properties` file, refer to the WebSphere Application Server Help System.

The installation of the operations console fails

This may occur when an error is encountered during the creation of the operations console database. To check if this is the case, view the log file `ISCRuntimeInstall.log` and search for an error message like the one in the following example:

```
COM.ibm.db2.jdbc.DB2Exception: [IBM][CLI Driver] SQL30061N
The database alias or database name "OPCONDBD"
was not found at the remote node.  SQLSTATE=08004
```

The log file `ISCRuntimeInstall.log` is written to the system temporary directory (refer to Table 27 on page 120).

Using the installation log files

Installation log file directories

The following table lists the directories to which installation log files are written during the installation of the end-to-end automation management component.

Table 27. Location of the installation log files

| |
|--|
| <p>Location</p> <p>Installation directory of the end-to-end automation management component: <EEZ_INSTALL_ROOT></p> <p>Default directory:</p> <p>Windows: C:\Program Files\IBM\tsamp\eez</p> <p>AIX/Linux: /opt/IBM/tsamp/eez</p> |
| <p>Installation log directory of the end-to-end automation management component:</p> <p>Windows: <EEZ_INSTALL_ROOT>\install\logs</p> <p>AIX/Linux: <EEZ_INSTALL_ROOT>/install/logs</p> |
| <p>System temporary directory:</p> <p>Windows: %TEMP%</p> <p>Typically, this is: C:\Documents and Settings\<user_name>\Local Settings\temp</p> <p>where <user_name> is the user ID that is used for running the installer</p> <p>AIX/Linux: When the environment variable \$TEMP is set, the directory specified in the environment variable.</p> <p>Typically, this is: /tmp</p> <p>When the environment variable is not set, the logs are written to the directory /tmp.</p> |
| <p>WebSphere Application Server profile log directory:</p> <p>Windows: <was_root>\profiles\<profile>\logs</p> <p>AIX/Linux: <was_root>/profiles/<profile>/logs</p> <p>where <profile> is the profile you specified in the installation wizard during the installation of the end-to-end automation management component (refer to step 16 on page 89).</p> |

Installation log files

The tables in the following sections list the log files that are written during the installation of the end-to-end automation management component:

- Section “Pre-installation phase log files” on page 121 lists the log files that are written in the pre-installation phase, that is, before you click the **Install** button in the installation wizard.
- Section “Installation phase log files” on page 121 lists the log files that are written in the installation phase, that is, after you click the **Install** button in the installation wizard.

In the tables, the following acronyms appear:

- ISC** Integrated Solutions Console is the application server in which the operations console runs
- CEI** Common Event Infrastructure is part of the Tivoli Enterprise Console (TEC) server connection
- EIF** Event Instrumentation Framework is part of the Tivoli Enterprise Console (TEC) server connection

Pre-installation phase log files

During the pre-installation phase, that is, before you click the **Install** button in the installation wizard, various log files are created:

- **e2ewizard.log**

This is the first log file the installation wizard creates. Use it as the first source of information if you encounter problems during the pre-installation phase.

The file is located in the system temporary directory (refer to Table 27 on page 120).

At the beginning of the installation phase, after you click the Install button in the installation wizard, the contents of the file are copied to the installation log file `e2einstaller.log` (refer to “Installation phase log files”).

- **Task-specific pre-installation log files**

During the pre-installation phase, the installation wizard initiates numerous tasks. The messages that are generated while the tasks are performed are logged in task-specific log files.

To find out the names and locations of these log files, refer to the log file `e2ewizard.log`. It contains references that let you identify and locate these log files.

Installation phase log files

During the installation phase, that is, after you click the **Install** button, various log files are created:

- **e2einstaller.log**

This is the log file to which the installation wizard logs information during the installation phase. Use it as the first source of information if you encounter problems during the installation phase.

The file is located in the installation directory of the end-to-end automation management component (refer to Table 27 on page 120).

The file is created at the beginning of the installation phase. The contents of the pre-installation phase log file `e2ewizard.log` are copied to the file.

- **Task-specific installation-phase log files**

During the installation phase, the installation wizard initiates numerous tasks. The messages that are generated while the tasks are performed are logged in task-specific log files.

To find out the names and locations of these log files, refer to the log file `e2einstaller.log`. It contains references that let you identify and locate these log files.

- **Installation-phase log files in the WebSphere Application Server profile log directory or in the system temporary directory**

The log files listed in Table 28 on page 122 are written to either the WebSphere Application Server profile log directory or the system temporary directory or both (refer to Table 27 on page 120) during the installation phase:

Table 28. Installation logs in the WebSphere Application Server profile log directory or in the system temporary directory

| Log file name | Description |
|---------------------------|-------------------------------------|
| events_install_msg.log | CEI installation messages |
| events_install_trc.log | CEI installation traces |
| events_db_install_msg.log | CEI database configuration messages |
| events_db_install_trc.log | CEI database configuration traces |

- **Installation log files in the system temporary directory** (refer to Table 27 on page 120).

Table 29. Installation logs in the system temporary directory

| Log file name | Description |
|--------------------------------------|--|
| ISCAction.log | This log file is written while the ISC installer checks whether there is already an ISC installed in the WebSphere Application Server. |
| ISCRuntime.rsp | This response file contains all the information that was passed to the Integrated Solutions Console installer. |
| ISCRuntimeConfig.properties | This properties file contains all the information used by the Integrated Solutions Console installer. |
| ISCRuntimeInstallProgress.properties | This properties file contains a progress indicator for a running Integrated Solutions Console installation. For a completed installation, it contains the information whether the installation was successful. |
| ISCRuntimeInstall.log | The main log file for the Integrated Solutions Console installation. Calls the following sub-installation tasks: ISCPreInstallConfigTask ISCPostInstallConfigTask. |
| ISCPreInstallConfigTask.log | Installs WebSphere Portal. Called by ISCRuntimeInstall. |
| wpsinstalllog.txt | Log file for WebSphere Portal installation. |
| ISCPostInstallConfigTask.log | Prepares the installation of Integrated Solutions Console and performs the installation. Called by ISCRuntimeInstall. |
| PortalBasicConfig.log | Prepares WebSphere Portal for the installation of Integrated Solutions Console. Called by ISCPostInstallConfigTask. |
| ISCUdatePorts.log | Updates the ports of the Integrated Solutions Console Portal server. Called by ISCPostInstallConfigTask. |

Table 29. Installation logs in the system temporary directory (continued)

| Log file name | Description |
|------------------------------|--|
| ISCEnableSecurityDB.log | Enables WebSphere Application Server security. The log file is written when DB2 is used as the user registry. Called by ISCPPostInstallConfigTask. |
| ISCEnableSecurityLDAPNew.log | Enables WebSphere Application Server security. The log file is written when LDAP is used as the user registry. Called by ISCPPostInstallConfigTask. |
| ISCSSetupISCM1.log | Changes the WebSphere Portal configuration settings. Called by ISCPPostInstallConfigTask. |
| PortalDeployConfig.log | Changes the Integrated Solutions Console configuration settings. Called by ISCPPostInstallConfigTask. |
| WP_Fixes.log | Installs the WebSphere Application Server fixes needed by WebSphere Portal and Integrated Solutions Console. Called by ISCPPostInstallConfigTask. |
| PortalPostConfig.log | Changes the Integrated Solutions Console configuration settings. Called by ISCPPostInstallConfigTask. |
| ISCDeploy_AdminComponents | Contains return codes of deploy operations. |

Procedures for troubleshooting an installation

In general, the installation wizard displays an error message if the installation fails. The error message explains why the installation failed or points you to a log file. In the latter case, view the referenced log file and look for error messages or other indications of errors.

When the error occurred in the pre-installation phase, that is, before the **Install** button was clicked, use the log files described in “Pre-installation phase log files” on page 121 in the following way:

1. View the log file `e2ewizard.log`
2. Search for error indications
3. View all additional log files that are written in the pre-installation phase and search for error indications. The names and locations of the additional log files appear in the file `e2ewizard.log`.

When the error occurred in the installation phase, that is, after the **Install** button was clicked, use the log files described in “Installation phase log files” on page 121 in the following way:

1. View the file `e2einstaller.log`

2. Search for error indications, for example, unexpected return codes (RC) or exceptions. Search for the strings "RC" and "Exception". Locate the first problem that occurred.
3. Check for previous or subsequent messages that may explain why the problem occurred.
4. Check whether the failed step creates an additional log file. If this is the case, view the relevant log file. The names and locations of the additional log files appear in the file `e2einstaller.log`.
5. Repeat the procedure for every log file.

Gathering information for IBM Support

If you cannot resolve an installation problem and need to contact IBM Support, you can use the log file collector tool to collect the log files that were written during the installation and send the resulting archive to IBM Support.

Perform these steps to run the log file collector tool:

1. Change the directory to `<EEZ_INSTALL_ROOT>/install`.

2. Issue the following command:

Windows: **e2einstallerlogs**

AIX/Linux: **e2einstallerlogs.sh**

The installer log archive is created in the directory from which you invoke the collector tool. The name of the archive has the following format:

`e2einstallerlogs_<timestamp>.zip`

where `<timestamp>` is a concatenation of the current date and the time at which the collector tool was invoked.

Part 3. First steps

| | |
|---|-----|
| Chapter 17. Overview | 127 |
| Chapter 18. Starting the sample end-to-end automation domain | 129 |
| Chapter 19. Activating the sample end-to-end automation policy | 131 |
| Chapter 20. Creating and activating a new sample policy | 133 |
| Creating a new sample policy | 133 |
| Changing the domain name | 134 |
| Chapter 21. Displaying a first-level automation domain on the operations console | 137 |
| Checking the adapter configuration | 137 |
| Starting the adapter | 138 |
| Where to find the first-level automation domain on the operations console | 138 |
| Chapter 22. Creating a policy that references actual first-level resources | 139 |

Chapter 17. Overview

During the installation of the end-to-end automation management component, a sample end-to-end automation management environment is set up:

- The sample end-to-end automation domain “FriendlyE2E” is configured
- The sample policy file `sample.xml` is saved to the policy pool directory

The following chapters describe how you can use the sample end-to-end automation environment to learn more about the design of the operations console and the functionality it provides, and about the tasks you need to perform to create, change, and activate policies.

You can use the following chapters like a tutorial. When you follow the descriptions, you will use the sample end-to-end automation environment to obtain the following information:

- How to connect to an end-to-end automation domain (see Chapter 18, “Starting the sample end-to-end automation domain,” on page 129)
- How to activate a policy (see Chapter 19, “Activating the sample end-to-end automation policy,” on page 131)
- How to create and activate a new policy (see Chapter 20, “Creating and activating a new sample policy,” on page 133)
- How to display the first-level automation domains and the resources that are hosted by the domains on the operations console (see Chapter 21, “Displaying a first-level automation domain on the operations console,” on page 137)
- Which steps are required to adapt a policy and to activate the modified policy (see Chapter 22, “Creating a policy that references actual first-level resources,” on page 139)

Note: In the descriptions in the following chapters it is assumed that you accepted “FriendlyE2E” as name for the end-to-end automation domain when you installed the end-to-end automation management component. If you specified a different name for the end-to-end automation domain during installation or subsequently, you must first change the domain name you specified to “FriendlyE2E”. How you achieve this is described in “Changing the domain name” on page 134.

Chapter 18. Starting the sample end-to-end automation domain

Perform the following steps to launch the operations console of Tivoli System Automation for Multiplatforms and to display the sample end-to-end automation domain on the console:

1. Log in to the system on which the WebSphere Application Server instance is installed that hosts the automation J2EE framework.

2. Check that WebSphere Application Server, Integrated Solutions Console and the Eclipse Help System server are running.

3. Start the automation engine:
 - **Windows:**

On the task bar, click **start** —> **Run**, and click **Browse** to navigate to the start script of the automation engine (eezdmn.bat). Start the automation engine with the following command:

```
eezdmn.bat
```
 - **AIX and Linux:**

Change to the automation engine installation directory. Typically, this is /opt/ibm/tsamp/eez. The shell script for starting the engine is located in the bin directory, for example:

```
cd /opt/ibm/tsamp/eez/bin
```

Start the automation engine with the following command:

```
./eezdmn.sh
```

4. Open your Web browser and connect to Integrated Solutions Console. The address you enter has the following form:

```
http://<your_isc_server>:<isc_server_port>/ibm/console
```

If you accepted the default ports for Integrated Solutions Console during installation, the port number is 8421.

5. On the Welcome panel of Integrated Solutions Console, enter your user ID and password:
 - You can use the administrator user ID you created for Integrated Solutions Console during installation. If you accepted the default value, the user ID is `iscadmin`.
 - If you have already created and authorized end-to-end automation-specific user IDs, the user ID you use for logging on must belong to a group that allows you to activate a policy.

After entering your user ID and password, click **Log in**.

6. In the navigation tree on the left, expand Tivoli System Automation for Multiplatforms and click SA operations console to display the Connect panel.

7. On the Connect panel, accept the default port number and the default server name (localhost) that are displayed on the Connect panel by clicking **OK**.

Note: If you changed the port number when you installed the end-to-end automation management component, refer to “How to determine the server port number for connecting to the operations console” on page 296 for a description of how to find out the correct port number.

Results:

- The main panel of operations console is displayed:
 - The single icon that is displayed in the topology tree on the left of the main panel of the operations console represents the end-to-end automation domain “FriendlyE2E”. The domain was configured during the installation of the end-to-end automation management component.
 - The resource table is empty because no policy has been activated yet.

Next steps:

- To get an overview of the layout of the operations console, use one of the following approaches:
 - Read Chapter 32, “What you must know about the operations console,” on page 209 in this manual.
 - Read the description of the main panel and its components in the online help that is provided for the operations console. To display the online help, click ? on the main panel of the operations console.
- To be able to explore the operations console, you need to activate the sample policy. This is described in the following chapterChapter 19, “Activating the sample end-to-end automation policy,” on page 131.

Chapter 19. Activating the sample end-to-end automation policy

When you activate the sample policy, the resource table on the operations console will be populated with the dummy resources that have been specified in the policy.

During the installation of the end-to-end automation management component, the sample policy `sample.xml` was saved in the policy pool directory where all policy files must be available to be activated.

To activate the sample policy, perform the following steps:

1. In the topology tree, select the domain "FriendlyE2E".

2. In the information area which is displayed to the right of the navigation trees, click the **Policy** tab to open the Policy page.

3. On the Policy page, click **Activate new policy**. This brings up the Select an automation policy panel.

4. Select the policy "Sample E2E Policy" and click **Activate** to activate the policy.

Results:

- The automation manager activates the policy.
- The top-level resources defined in the policy are displayed in the resource table.

Next steps:

- Now you can explore the operations console. To learn about the layout of the console, to find out how to navigate it and what the displayed elements represent, refer to the descriptions in Chapter 32, "What you must know about the operations console," on page 209. The complete list of icons that appear on the console is available in Appendix F, "Operations console icons," on page 309.
- To understand how the resources that are displayed on the operations console map to the definitions in the XML policy file, you can look at the `sample.xml` policy file in "Sample automation policy" on page 282.
- To learn how to create and activate a new policy, perform the tasks described in Chapter 20, "Creating and activating a new sample policy," on page 133.

Chapter 20. Creating and activating a new sample policy

In this chapter you learn which tasks you need to perform to create and activate your own end-to-end automation policy for a new end-to-end automation domain.

The step-by-step descriptions provided in the sections of this chapter contain all the information you need to perform the tasks for a new sample policy. For detailed information about defining XML policies, refer to Chapter 25, "Creating and modifying policies," on page 161.

Creating a new sample policy

Perform the following steps to create a new sample policy:

1. Log in to the system where the end-to-end automation manager is installed.
2. Go to the policy pool directory and copy the file `sample.xml` to your working directory.
3. Open the copy of `sample.xml` in an XML editor.

Note: You can also use a text editor for creating and editing XML policy files. Whichever editor you choose, you must ensure that you can save the file in UTF-8 format. Policy files in any other format cannot be activated.

4. Change the `<PolicyInformation>` section in the file as shown in the following example (changes to the original `sample.xml` are marked in bold):

```
<PolicyInformation>
  <PolicyName> My sample policy </PolicyName>
  <AutomationDomainName> My Domain </AutomationDomainName>
  <PolicyToken>0.1</PolicyToken>
  <PolicyAuthor>Bob</PolicyAuthor>
  <PolicyDescription>My first policy</PolicyDescription>
</PolicyInformation>
```

5. Create a new dummy resource reference:

```
<ResourceReference name="My Reference">
  <DesiredState>Offline</DesiredState>
  <Description>My first resource reference</Description>
  <Owner>Bob</Owner>
  <InfoLink>http://www.example.com</InfoLink>
  <ReferencedResource>
    <AutomationDomain>MyFLADomain</AutomationDomain>
    <Name>MyResource</Name>
    <Class>ResourceGroup</Class>
  </ReferencedResource>
</ResourceReference>
```

6. Save the new policy as `MySamplePolicy.xml` and copy it to the policy pool directory.
-

Before you can activate the policy, you must change the domain name in the configuration dialog of the automation manager. This is described in the following section.

Changing the domain name

You can only activate an end-to-end automation policy if the domain name in the XML element `<AutomationDomainName>` in the XML policy file is identical to the name of the currently active end-to-end automation domain. The name of the currently active end-to-end automation domain is specified on the Domain page of the configuration dialog.

If you have edited the XML policy file according to the description in the previous section, you have changed the `<AutomationDomainName>` in the policy file to "My Domain". This is why you need to change the name of the end-to-end automation domain in the configuration dialog before you can activate the policy. This is described in the following procedure.

Perform the following steps:

1. Log in to the system on which the end-to-end automation manager is installed.
2. Stop the automation engine:
 - **Windows:**
On the task bar, click **start** —> **Run**, and click **Browse** to navigate to the stop script of the automation engine (`eezdmn.bat`). Stop the automation engine with the following command:

```
eezdmn.bat -shutdown
```
 - **AIX and Linux:**
Change to the end-to-end automation installation directory (`EEZ_INSTALL_ROOT`). Typically, this is `/opt/IBM/tsamp/eez`. The shell script to stop the engine is located in the `bin` directory, for example:

```
cd /opt/IBM/tsamp/eez/bin
```

Stop the automation engine with the following command:

```
./eezdmn.sh -shutdown
```
3. Start the configuration dialog and open the Domain page. For information how to start the configuration dialog, refer to "Launching the configuration dialog" on page 193.
4. On the Domain page, change the name in the field **Domain name** to "My domain".
5. Click **Save**.
6. Click **Cancel** to close the dialog.
7. Start the automation engine as described in Chapter 18, "Starting the sample end-to-end automation domain," on page 129.
Shortly after the automation engine has started, the new automation domain "My Domain" appears in the topology tree. The domain "FriendlyE2E" still

exists but is grayed out. The domain has left, as this state is described in the terminology of end-to-end automation management.

8. Activate the new policy by following the instructions provided in “Activating a policy” on page 249.
-
9. Select the domain “My Domain” in the topology tree to display the new resource “My Reference” in the resource table.
-

Chapter 21. Displaying a first-level automation domain on the operations console

This chapter describes the basic tasks you need to perform before you can access the resources that are hosted by a first-level automation domain from the operation console:

- “Checking the adapter configuration”
- “Starting the adapter” on page 138

Before you begin:

The end-to-end automation manager must be able to access the first-level automation domains. This is the case when the user credentials for the first-level automation domains have been specified on the User credentials page of the end-to-end automation manager configuration dialog.

To check that the credentials have been specified correctly, check the relevant settings on the User credentials page of the configuration dialog.

For information on how to launch the configuration dialog, refer to “Launching the configuration dialog” on page 193. For detailed information about the User credentials page, refer to the online help of the configuration dialog.

Checking the adapter configuration

The connection between the end-to-end automation manager and a first-level automation domain is established through the automation adapter. An automation adapter must be configured and started on each first-level automation domain.

The steps below give an overview of what you must do to check the adapter configuration. For detailed information about the automation adapters, refer to the automation adapter documentation for SA for Multiplatforms and SA z/OS.

Perform the following steps to verify that the configuration of an automation adapter is correct:

1. Log in to the system of the first-level automation domain where you want to start the adapter process.
2. Open the properties file of the automation adapter. Check whether the value of the property **eif-send-to-hostname** is correct. The value must be the IP address of the server where the end-to-end automation manager is running.
3. Verify the setting of the property **eif-send-to-port**. The port number must be identical to the port number that is specified on the Domain page of the end-to-end automation configuration dialog (field **Event port number**).

Results:

When you have verified that the properties are set correctly, the adapter process should be able to make initial contact with the end-to-end automation manager.

Starting the adapter

For detailed information about how you start the automation adapters on the first-level automation domains, refer to the automation adapter documentation for SA for Multiplatforms and SA z/OS.

Where to find the first-level automation domain on the operations console

Shortly after you have started the adapter, the first-level automation domain sends a so-called domain-join event to the end-to-end automation manager. This event contains all the data the automation manager needs to contact the first-level automation domain.

The new automation domain is displayed in the topology tree on the operations console:

- If no end-to-end automation policy is active or if the active end-to-end automation policy does not contain references to resources that are hosted by the first-level automation domain, the new first-level automation domain is displayed at the same tree level as the end-to-end automation domain.
- If an end-to-end automation policy is active and the policy contains references to resources that are hosted by the first-level automation domain, the domain is displayed as a child element of the end-to-end automation domain.

If a first-level automation domain of SA for Multiplatforms is not visible although it should appear in the topology tree, refer to “A System Automation for Multiplatforms domain is not displayed in the topology tree” on page 296 for information on how to resolve the problem.

Chapter 22. Creating a policy that references actual first-level resources

After an adapter on a first-level automation domain is configured, the resources that are hosted by this domain are available for being referenced in an end-to-end automation policy.

To create the resource references for the resources of the first-level automation domain, you can use the sample policy **My sample policy** that you created in section “Creating a new sample policy” on page 133, and modify it accordingly.

To gather the data about the first-level resources that you need for defining resource references, you can use the information provided for the resources of the first-level automation domains in the information area of the operations console.

Part 4. Administering the end-to-end automation management component

| | |
|--|-----|
| Chapter 23. Post-installation tasks for administrators | 143 |
| Access roles for end-to-end automation management | 143 |
| How users are given roles | 145 |
| Creating user groups in Integrated Solutions Console | 145 |
| Assigning access permissions to user groups in Integrated Solutions Console | 146 |
| Granting user groups access to the pages of Integrated Solutions Console | 146 |
| Granting user groups access to the operations console of Tivoli System Automation for Multiplatforms | 147 |
| Assigning the user ID of the automation engine to groups in Integrated Solutions Console | 148 |
| Assigning access roles to user groups in WebSphere Application Server. | 149 |
| Chapter 24. Managing users. | 153 |
| Managing users and user groups in Integrated Solutions Console | 153 |
| Creating and authorizing users in Integrated Solutions Console | 153 |
| Creating users in Integrated Solutions Console | 153 |
| Assigning users to groups in Integrated Solutions Console | 154 |
| Administering users and user groups in Integrated Solutions Console | 154 |
| Changing passwords for users in Integrated Solutions Console | 154 |
| Deleting user IDs in Integrated Solutions Console | 155 |
| Deleting groups in Integrated Solutions Console | 155 |
| Managing the user credentials of subcomponents of end-to-end automation management. | 156 |
| Modifying the default user ID of the automation engine. | 156 |
| Encrypting the passwords in the file sas.client.props | 156 |
| Managing the user ID used by the automation engine to access first-level automation domains | 157 |
| Modifying the default user ID used by the automation management server to access DB2 | 157 |
| Managing JMS authentication | 157 |
| Modifying the default JMS authentication entry for the automation engine | 158 |
| Modifying the default JMS authentication entry for the operations console | 158 |
| Modifying the default JMS authentication entry for the automation management server. | 159 |
| Chapter 25. Creating and modifying policies | 161 |
| What you must know before you define an end-to-end automation policy | 162 |
| The scope of end-to-end automation policies | 162 |
| Example 1 | 163 |
| Example 2 | 163 |
| Example 3 | 164 |
| Identifying cluster-spanning dependencies. | 164 |
| Grouping of resources | 164 |
| Relationships | 165 |
| Gathering the required data for defining a policy | 166 |
| Considerations for referencing SA for Multiplatforms resources | 167 |
| Restrictions for referencing SA z/OS resources | 167 |
| Defining an end-to-end automation policy. | 168 |
| Creating the XML policy file | 168 |
| Defining the resources of the end-to-end automation domain | 171 |
| Defining groups | 174 |
| Defining resource groups | 174 |
| Defining choice groups | 176 |
| Defining StartAfter, StopAfter, and ForcedDownBy relationships | 177 |
| Defining a StartAfter relationship. | 177 |
| Defining a StopAfter relationship. | 178 |
| Defining a ForcedDownBy relationship. | 178 |
| Saving the policy in the policy pool directory | 179 |
| Starting the policy checking tool from a command line | 179 |
| Chapter 26. Setting up information pages for operators | 181 |
| Chapter 27. Using the command-line interface of the automation engine. | 183 |
| eezdmn options quick reference | 184 |
| eezdmn options | 184 |
| -start | 184 |
| Return codes | 185 |
| -shutdown | 185 |
| Return codes | 185 |
| -monitor | 186 |
| -reconfig | 187 |
| Return codes | 187 |
| -co | 187 |
| Return codes | 187 |
| -xd | 188 |
| Return codes | 188 |
| -? | 188 |
| Chapter 28. Starting and stopping | 189 |
| Starting and stopping WebSphere Application Server | 189 |

| | |
|---|-----|
| Starting and stopping WebSphere Application Server on Windows | 189 |
| Starting and stopping WebSphere Application Server on AIX and Linux | 190 |
| Starting and stopping the automation J2EE framework | 190 |
| Starting and stopping the operations console | 190 |
| Starting and stopping the operations console on Windows | 191 |
| The servers are running as Windows services | 191 |
| The servers are not running as Windows services | 191 |
| Starting and stopping the operations console on AIX and Linux | 192 |
| Starting and stopping the automation engine | 192 |

| | |
|--|------------|
| Chapter 29. Configuring the end-to-end automation manager | 193 |
| Launching the configuration dialog | 193 |
| Configuration dialog pages | 193 |
| Domain page | 194 |
| User credentials page | 195 |
| Security page | 197 |
| Logger page | 198 |

| | |
|---|------------|
| Chapter 30. Using Tivoli Enterprise Console with SA for Multiplatforms | 199 |
| Configuring Tivoli Enterprise Console | 199 |
| Checking the Tivoli Event Integration Facility function | 200 |
| Enabling Tivoli Enterprise Console event filtering | 202 |
| Activating the default CEI filter | 202 |
| Customizing the default event filter | 204 |

Chapter 23. Post-installation tasks for administrators

After installing the end-to-end automation management component, you must perform the following tasks regardless of whether the operations console will be used in end-to-end automation mode or first-level automation mode:

1. Familiarize yourself with the end-to-end automation-specific access roles.

In end-to-end automation management, access control is based on WebSphere Application Server security implementation. The roles are described in “Access roles for end-to-end automation management.”

-
2. Create user groups in Integrated Solutions Console.

You need to create a user group for each of the end-to-end automation-specific roles. This is described in “Creating user groups in Integrated Solutions Console” on page 145.

-
3. Assign access permissions to the user groups in Integrated Solutions Console.

You need to grant the user groups access to the pages of Integrated Solutions Console and to the operations console of SA for Multiplatforms. The tasks you need to perform are described in “Assigning access permissions to user groups in Integrated Solutions Console” on page 146.

-
4. Add the user ID of the automation engine to the user groups EEZAdministratorGroup and EEZEndToEndAccessGroup.

To do this, follow the description in “Assigning the user ID of the automation engine to groups in Integrated Solutions Console” on page 148. If the user ID of the automation engine is not assigned to these groups, end-to-end automation management will be inoperable.

-
5. Assign the user groups to the access roles in WebSphere Application Server. This is described in “Assigning access roles to user groups in WebSphere Application Server” on page 149.
-

Access roles for end-to-end automation management

In end-to-end automation management, WebSphere Application Server access roles determine which actions an automation management user can perform on the operations console. Table 30 on page 144 describes the access roles that are available for end-to-end automation management.

For five of the six access roles you will create user groups in Integrated Solutions Console. In the administrative console of WebSphere Application Server, you assign access roles to the user groups.

When you create the groups in Integrated Solutions Console, you should specify group names that are similar to the names of the corresponding access roles. Group name recommendations are provided in the rightmost column of the following table.

Table 30. WebSphere Application Server access roles for end-to-end automation management

| Role | Permissions | Recommended group name |
|-------------------|--|------------------------|
| EEZMonitor | <p>Grants minimum access rights. Users who have this role can perform query-type operations.</p> <p>They cannot activate a policy or perform actions that modify the state of a resource, for example, submit start requests.</p> | EEZMonitorGroup |
| EEZOperator | <p>Extends the EEZMonitor role.</p> <p>Users who have this role can issue requests against first-level and end-to-end automation resources but cannot activate or deactivate policies.</p> <p>Operators managing both first-level and end-to-end automation resources must also have the role EEZEndToEndAccess.</p> | EEZOperatorGroup |
| EEZConfigurator | <p>Extends the EEZMonitor role.</p> <p>Users given this role cannot submit requests against resources.</p> <p>The role is required to be able to work with policies, for example, to activate policies. Currently, such actions can only be performed for end-to-end automation domains. This is why users who will perform these actions must also have the role EEZEndToEndAccess.</p> | EEZConfiguratorGroup |
| EEZAdministrator | <p>Extends the EEZOperator and EEZConfigurator roles.</p> <p>Users who have this role can perform all operations provided on the operations console for first-level and end-to-end automation domains.</p> <p>Administrators who manage both first-level and end-to-end automation domains must also have the role EEZEndToEndAccess.</p> <p>Note: The user ID of the automation engine must have the role EEZAdministrator. You achieve this by adding the user ID to the EEZAdministratorGroup.</p> | EEZAdministratorGroup |
| EEZEndToEndAccess | <p>User who do not have this role, can view and monitor the end-to-end automations domain and the resources hosted by the domain.</p> <p>This role is only required if a user needs to start or stop end-to-end automation resources or activate and deactivate policies.</p> <p>This means that this role determines which type of automation domain a user who has this role can or cannot access. It does not determine which operations can be performed by a user given this role.</p> <p>Note: The user ID of the automation engine must have the role EEZEndToEndAccess. You achieve this by adding the user ID to the EEZEndToEndAccessGroup.</p> | EEZEndToEndAccessGroup |

Table 30. WebSphere Application Server access roles for end-to-end automation management (continued)

| Role | Permissions | Recommended group name |
|----------|--|---|
| EEZAsync | RunAs role for all internal WebSphere methods of end-to-end automation management that are invoked by a timer or an event. | No group required. In WebSphere Application Server, the role will be mapped to All authenticated and to the WebSphere Application Server user ID. |

How users are given roles

Users may have to have more than one role to be able to perform the actions they are responsible for. For example, operators who need to be able to submit start and stop requests against end-to-end automation resources must have the roles EEZOperator and EEZEndToEndAccess:

- The EEZOperator role authorizes users to monitor resources, perform query-type operations, and submit requests from the operations console. User who only have this role can only submit requests against first-level automation resources.
- The EEZEndToEndAccess role authorizes them to also submit requests against end-to-end automation resources.

However, you do not assign access roles to users directly. Instead, you do the following:

1. For each access role in WebSphere Application Server, you create a user group in Integrated Solutions Console.

For example, for the access roles EEZOperator and EEZEndToEndAccess, you create the user groups EEZOperatorGroup and EEZEndToEndAccessGroup.

2. In WebSphere Application Server you assign an access role to the corresponding user group.

For example, you assign the access role EEZOperator to the user group EEZOperatorGroup, and the access role EEZEndToEndAccess to the user group EEZEndToEndAccessGroup.

3. In Integrated Solutions Console, you create users and assign the users to the user groups.

If a user must have more than one access role, you assign the user to the user groups that correspond to the access roles the user must have.

For example, when a user must have both the EEZOperator role and the EEZEndToEndAccess role, you assign the user to the groups EEZOperatorGroup and EEZEndToEndAccessGroup. This will give the user both of the required roles because you have mapped the user groups to the corresponding access roles in WebSphere Application Server.

Creating user groups in Integrated Solutions Console

In Integrated Solutions Console, you must create one user group for each of the end-to-end automation-specific WebSphere Application Server access roles.

You must create the following groups:

- EEZAdministratorGroup
- EEZConfiguratorGroup

- EEZOperatorGroup
- EEZMonitorGroup
- EEZEndToEndAccessGroup

Note: You can use different group names but the names should be similar to the names of the end-to-end automation-specific access roles in WebSphere Application Server.

To create the groups, perform the following steps:

1. Log in to Integrated Solutions Console as administrator (default: user ID iscadmin, group iscadmins)

2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

3. Select **User and Group Management**.

4. On the User and Group Management page, click **New Group**.

5. Type the name of the user group you want to create in the **ID** field and click **OK** to save the group and to return to the User and Group Management page.

Repeat steps 4 and 5 until you have created all the required groups.

Assigning access permissions to user groups in Integrated Solutions Console

After creating the user groups in Integrated Solutions Console, you must perform the following tasks:

- “Granting user groups access to the pages of Integrated Solutions Console”
- “Granting user groups access to the operations console of Tivoli System Automation for Multiplatforms” on page 147

Granting user groups access to the pages of Integrated Solutions Console

Perform the following steps:

1. Log in to Integrated Solutions Console as administrator (default: user ID iscadmin, group iscadmins)

2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

3. Click **Resource Permissions** to display the **Resource Types** list.

4. In the **Resource Types** list, click **Pages** to display the list of resources.

5. In the **Resources** list, click  for **Content Root**.

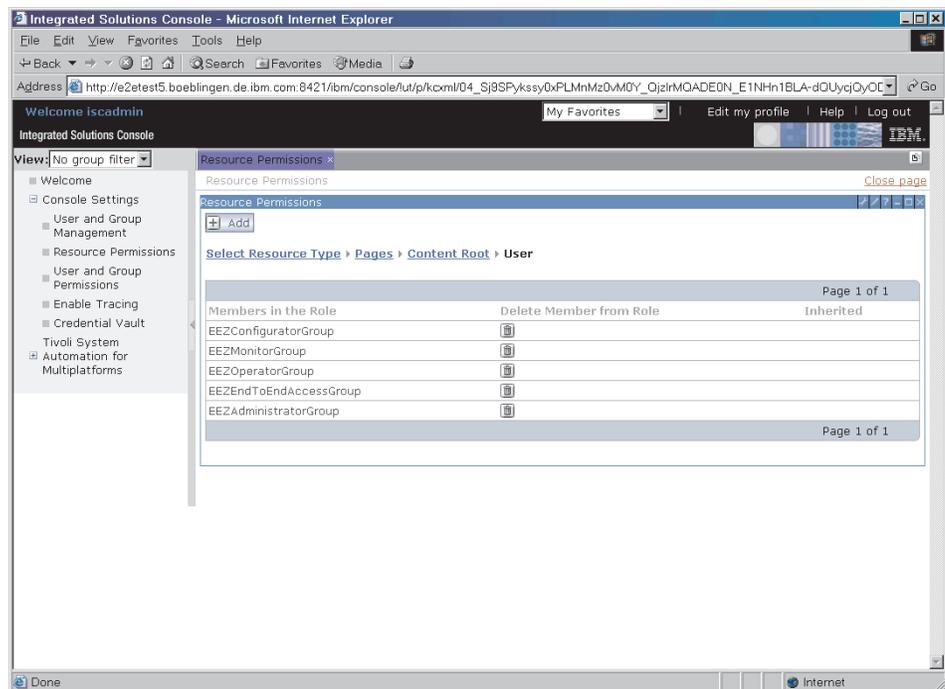
-
6. In the **Roles** list, click  for **User**.

 7. Click **Add**. On the page that appears, only the groups are listed for which access has already been granted. If you perform this task for the first time, the page may be empty.

 8. Click **Search**. (In the field **Search for Users or User Groups**, the entry **User Groups** must be selected.)

 9. In the **Users and User Groups** list, select the check boxes for the end-to-end automation-specific groups:
 - EEZAdministratorGroup
 - EEZConfiguratorGroup
 - EEZEndToEndAccessGroup
 - EEZMonitorGroup
 - EEZOperatorGroup

 10. Click **OK**. The following panel appears:



-
11. Restart Integrated Solutions Console.
-

Granting user groups access to the operations console of Tivoli System Automation for Multiplatforms

Perform the following steps:

1. Log in to Integrated Solutions Console as administrator (default: user ID iscadmin, group iscadmins)

-
2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

 3. Click **Resource Permissions** to display the **Resource Types** list.

 4. In the **Resource Types** list, click **Portlet Applications** to display the list of resources.

 5. In the **Resources** list, click  for **Tivoli System Automation for Multiplatforms Operations Console**.

 6. In the **Roles** list, click  for **User**.

 7. Click **Add**.

 8. Click **Search**. (In the field **Search for Users or User Groups**, the entry **User Groups** must be selected.)

 9. In the **Users and User Groups** list, select the check boxes for the end-to-end automation-specific groups:
 - EEZAdministratorGroup
 - EEZConfiguratorGroup
 - EEZEndToEndAccessGroup
 - EEZMonitorGroup
 - EEZOperatorGroup

 10. Click **OK**.

 11. Restart Integrated Solutions Console.
-

Assigning the user ID of the automation engine to groups in Integrated Solutions Console

The automation engine is only operable when the user ID of the automation engine is assigned to the user groups EEZAdministratorGroup and EEZEndToEndAccessGroup.

By default, the user ID of WebSphere Application Server is also used for the automation engine. The default user ID of WebSphere Application Server is iscadmin.

To assign user ID of the automation engine to the groups, perform the following steps:

1. Log in to Integrated Solutions Console as administrator (default: user ID iscadmin, group iscadmins)
-

2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

3. Select **User and Group Management**.

4. On the User and Group Management page, click **All Portal User Groups**. The list of user groups is displayed.

5. Select the group EEZAdministratorGroup.

6. Click **Add member**. The list of users who are not members of the group is displayed.

7. From the user list, select the user ID of the automation engine and click **OK**.

8. Select the group EEZEndToEndAccessGroup.

9. From the user list, select the user ID of the automation engine and click **OK**.

Assigning access roles to user groups in WebSphere Application Server

After you have created the end-to-end automation-specific user groups in Integrated Solutions Console, you must assign access roles to these groups in the administrative console of WebSphere Application Server. This will grant the members of a group all of the permissions the access role that is assigned to the group contains.

Note: You must make sure that the user ID of the automation engine has been added to the groups EEZAdministratorGroup and EEZEndToEndAccessGroup.

The following table shows how the access roles must be mapped to the user groups:

Table 31. Mapping of end-to-end automation management roles to groups in WebSphere

| Role | Map to |
|---------------------|---|
| EEZMonitor | EEZMonitorGroup |
| EEZOperator | EEZOperatorGroup |
| EEZConfigurator | EEZConfiguratorGroup |
| EEZAdministrator | EEZAdministratorGroup |
| EEZEndToEndAccess | EEZEndToEndAccessGroup |
| EEZAsync | map this to special All authenticated (leave the default as is) |
| RunAs role EEZAsync | map this to the WebSphere Application Server user ID |

The following description gives an overview of the steps you must perform to assign roles to groups. For detailed information, refer to the manual for WebSphere Application Server, Version 6, *Securing applications and their environment* (Chapter 10. Deploying secured applications → Assigning users and groups to roles).

Perform the following steps to assign access roles to groups:

1. Log in to the administrative console of WebSphere Application Server.

2. On the **Applications** menu, select **Enterprise Applications**.

3. Select **EEZEAR**.

4. Select **Map security roles to users/groups**.

5. Select one of the following roles: **EEZAdministrator**, **EEZConfigurator**, **EEZOperator**, **EEZMonitor**, **EEZEndToEndAccess**.

6. Click **Look up groups**.

7. Click **Search**.

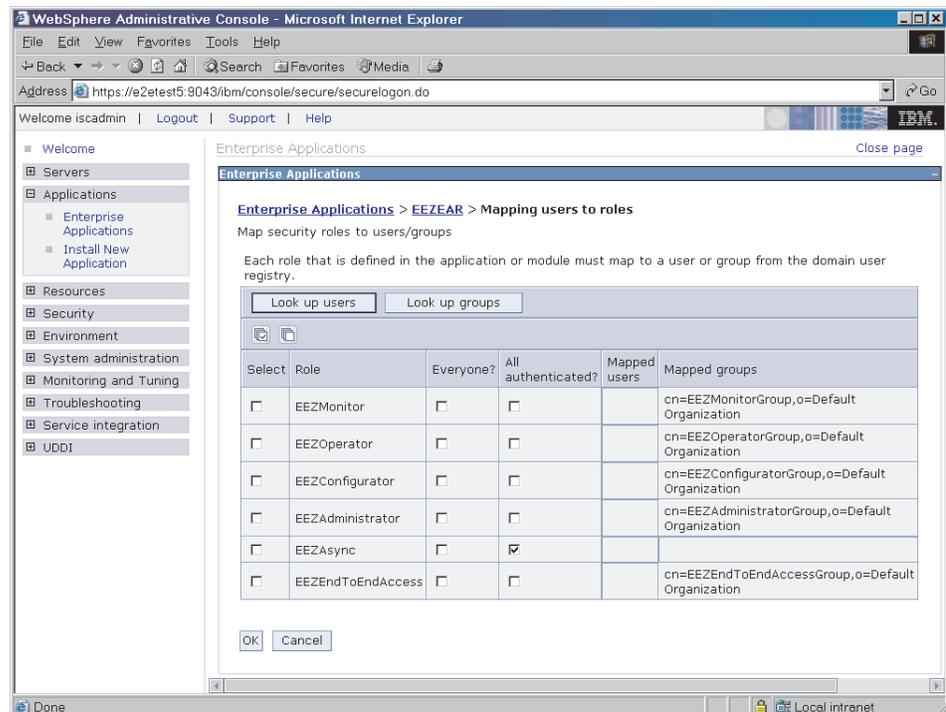
8. Select the group that matches the selected role.

9. Click **»**.

10. Click **OK**.

11. Repeat these steps starting from step 4 until you have assigned a role to each group.

12. Deselect the check box **Everyone** for all roles.
Deselect the check box **All Authenticated** for all roles but **EEZAsync**.



-
13. When you are done, click **OK**.

 14. Open the menu **Map RunAs roles to users** and enter the WebSphere Application Server user ID and password.

 15. Select the role EEZAsync and click **Apply**.

 16. Click **OK**.

 17. Open the **Save** menu and click **Save** to save the configuration.

 18. Restart WebSphere Application Server to activate the new configuration.

Chapter 24. Managing users

This chapter describes:

- How you create and administer users for end-to-end automation management in Integrated Solutions Console
- How you modify the user credentials for the subcomponents of end-to-end automation management

Managing users and user groups in Integrated Solutions Console

This chapter describes how you create, authorize, and administer users in Integrated Solutions Console and how you delete user groups. Additional information is available in Integrated Solutions Console help. To access the relevant help pages, open the console Help menu and navigate to **Console Basics** → **Console Settings** → **User and Group Management**.

Creating and authorizing users in Integrated Solutions Console

To authorize a user to work with end-to-end automation management, you create the user and assign the user to one or more user groups in Integrated Solutions Console. The relevant tasks you need to perform are described in this section.

Notes:

1. Before you create and authorize users, you must have completed the tasks described in Chapter 23, “Post-installation tasks for administrators,” on page 143.
2. Each user who manages first-level automation resources from the operations console must have a user ID on the first-level automation domain.

The operations console users can store and change their user credentials for the domains in the credential vault of Integrated Solutions Console (refer to “Managing your user credentials for first-level automation domains” on page 247 for more information).

Creating users in Integrated Solutions Console

To create a user in Integrated Solutions Console, perform the following steps:

1. Log in to Integrated Solutions Console as administrator (default: user ID `iscadmin`, group `iscadmins`)

2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

3. Select **User and Group Management**.

4. On the User and Group Management page, click **All Authenticated Portal Users**.

5. Click **New user**.

6. Enter the user ID and password, and the user's first name, last name, and e-mail address, and click **OK**.
-

Assigning users to groups in Integrated Solutions Console

Before you begin:

- You must assign each user you create in Integrated Solutions Console to at least one user group.
- The user inherits the access permissions that you have granted to the group in Integrated Solutions Console and those of the access role to which the group is mapped in WebSphere Application Server. To decide which access permissions a user needs, refer to "Access roles for end-to-end automation management" on page 143.
- Each user group must have at least one member.

To assign a user to a group, perform the following steps:

1. Log in to Integrated Solutions Console as administrator (default: user ID `iscadmin`, group `iscadmins`)

 2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

 3. Select **User and Group Management**.

 4. On the User and Group Management page, click **All Portal User Groups**. The list of user groups is displayed.

 5. Select the appropriate group.

 6. Click **Add member**. The list of users who are not members of the group is displayed.

 7. From the user list, select the user or users you want to add to the group and click **OK**.
-

If you want to add users to additional groups, repeat steps 5 through 7.

Administering users and user groups in Integrated Solutions Console

The following sections describe how to modify and delete users and user groups.

To create a new automation management user, perform the tasks described in sections "Creating and authorizing users in Integrated Solutions Console" on page 153 and "Assigning users to groups in Integrated Solutions Console."

Changing passwords for users in Integrated Solutions Console

Perform the following steps to change user passwords:

1. Log in to Integrated Solutions Console as administrator (default: user ID `iscadmin`, group `iscadmins`)
-

2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

3. Select **User and Group Management**.

4. On the User and Group Management page, select **All Authenticated Portal Users**.

5. Click the **Edit** button for the user ID you want to modify.

6. Type the new password in the entry field.

7. Click **OK**.

Deleting user IDs in Integrated Solutions Console

Perform the following steps to delete users:

1. Log in to Integrated Solutions Console as administrator (default: user ID iscadmin, group iscadmins)

2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

3. Select **User and Group Management**.

4. On the User and Group Management page, select **All Authenticated Portal Users**.

5. Click the **Delete** button for the user ID you want to delete.

6. Click **OK**.

Deleting groups in Integrated Solutions Console

Perform the following steps to delete a group:

1. Log in to Integrated Solutions Console as administrator (default: user ID iscadmin, group iscadmins)

2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

3. Select **User and Group Management**.

4. On the User and Group Management page, select **All Portal User Groups**.

5. Click the **Delete** button for the group you want to delete.

6. Click **OK**.

Managing the user credentials of subcomponents of end-to-end automation management

In end-to-end automation management, user authentication is required for accessing the automation management server, the first-level automation management servers, and DB2. Authentication is always performed using user ID-password pairs.

Modifying the default user ID of the automation engine

The file `<was_root>/properties/sas.client.props` contains authentication- and SSL-related information to allow external programs to access WebSphere Application Server.

The file also contains the authentication entry that allows the automation engine to access the automation management server. This authentication entry is configured automatically during the installation of the end-to-end automation management component. By default, the user ID of WebSphere Application Server is also used for the automation engine. The default user ID of WebSphere Application Server is `iscadmin`.

When you change the password of WebSphere Application Server or when you want to use a user ID for the automation engine other than the WebSphere Application Server user ID, you must modify the following properties:

```
com.ibm.CORBA.loginSource=properties
# RMI/IIOP user identity
com.ibm.CORBA.loginUserId=<was_userid>
com.ibm.CORBA.loginPassword=<password_for_was_userid>
```

where `<was_userid>` is the user ID of WebSphere Application Server and `<password_for_was_userid>` is the corresponding password.

To activate the changes, you must restart the automation engine.

Note: The file also contains SSL settings. It is not necessary to activate the SSL settings in the file because the automation engine runs on the same system as the automation management server. You should not activate the settings because performance suffers if SSL is active.

For detailed information about the file `sas.client.props`, refer to the manual for WebSphere Application Server, Version 6, *Securing applications and their environment* (Chapter 12. Administering security —> Configuring Common Secure Interoperability Version 2 and Security Authentication Service authentication protocols —> Common Secure Interoperability Version 2 and Security Authentication Service client configuration).

Encrypting the passwords in the file `sas.client.props`

You can use the WebSphere Application Server `PropFilePasswordEncoder` utility to encrypt the passwords in the file `sas.client.props`. For information on how to do this, refer to the manual for WebSphere Application Server 6, *Securing applications and their environment* (Chapter 6. Implementing security considerations at installation time —> Protecting plain text passwords).

Managing the user ID used by the automation engine to access first-level automation domains

The user IDs and the corresponding passwords the end-to-end automation engine needs to authenticate itself to first-level automation domains and to the WebSphere Application Server JMS Provider are stored in the domain identification file of the automation engine (`eez.automation.engine.dif.properties`).

The file must be protected by means of operating system mechanisms in such a way that the automation engine is still able to read the contents of the file.

You browse and edit the properties that are in the file on the User credentials page of the configuration dialog.

For information about the configuration dialog, refer to Chapter 29, “Configuring the end-to-end automation manager,” on page 193. Information about the properties that can be configured in the dialog is provided in the dialog help.

Modifying the default user ID used by the automation management server to access DB2

This authentication entry is required to allow the application EEZEAR to access the DB2 database.

Perform the following steps to modify the default authentication data the automation management server uses to access DB2:

1. Log in to the WebSphere Application Server administrative console.

2. Go to **Security** → **Global Security** → **JAAS Configuration** → **J2C Authentication Data**

3. In the table, select Alias `<hostname>/eAuto`

4. Change the password or both the user ID and the password and click **OK**.

5. From the menu, select **save**.

6. Click **save** to save and activate the new configuration. Do not restart WebSphere Application Server.

For more information, refer to the manual for WebSphere Application Server 6, *Securing applications and their environment* (Chapter 12. Administering security → Configuring application logins for Java Authentication and Authorization Service → Managing J2EE Connector Architecture authentication data entries).

Managing JMS authentication

The subcomponents of end-to-end automation management use a WebSphere Application Server JMS Provider to exchange asynchronous messages (events). In order to grant access to this messaging service, the same user ID and password for the following end-to-end automation components must be specified in different places:

- for the automation engine, the user ID and password must be specified on the User credentials page of the configuration dialog (fields **JMS User ID** and **JMS password**).
- for the operations console, the user ID and password must be specified in the credential vault of Integrated Solutions Console, entry `com.ibm.eez_AutomationManager`
- for the end-to-end automation management server, on the WebSphere Application Server administrative console - J2C Authentication Data Entries, entry `EEZAuth`

The user ID you specify must be the same in all of these entries and the user ID must be a valid WebSphere Application Server authenticated user. The default is the WebSphere Application Server server user ID `iscadmin`.

Modifying the default JMS authentication entry for the automation engine

Use the User credentials page of the configuration dialog of end-to-end automation management to browse and modify the user ID and password that the automation engine uses to authenticate itself to the JMS provider of WebSphere Application Server.

For information about the configuration dialog, refer to Chapter 29, “Configuring the end-to-end automation manager,” on page 193. Information about the properties that can be configured in the dialog is provided in the dialog help.

Modifying the default JMS authentication entry for the operations console

The entry in the credential vault of Integrated Solutions Console is required to authorize the operations console to access the JMS provider of WebSphere Application Server.

Perform the following steps to modify the entry in the credential vault:

1. Log in to Integrated Solutions Console as administrator (default user ID: `iscadmin`, group `iscadmins`).

2. In the navigation tree of Integrated Solutions Console, expand **Console Settings**.

3. Select **Credential vault**.

4. On the Credential Vault page, **Manage System Vault Slots**.

5. Select **Modify Shared Slot** for **Vault Slot Name** `com.ibm.eez_AutomationManager`.

6. Change the password or both the user ID and the password and click **OK**.

7. Click **Done** to activate the new settings.

8. Log out. Do not restart Integrated Solutions Console.

Modifying the default JMS authentication entry for the automation management server

The JMS authentication entry for the automation management server is required to allow the application EEZEAR to access the WebSphere Application Server JMS Provider. You modify the entry on the administrative console of WebSphere Application Server.

Perform the following steps:

1. Log in to the WebSphere Application Server administrative console.

2. Go to **Security** —> **Global Security** —> **JAAS Configuration** —> **J2C Authentication Data**.

3. In the table, select Alias <hostname>/EEZAuth.

4. Change the password or both the user ID and the password and click **OK**.

5. From the menu, select **save**.

6. Click **save** to save and activate the new configuration. Do not restart WebSphere Application Server.

Chapter 25. Creating and modifying policies

The policy is a core component of end-to-end automation management. The policy determines:

- which resources are managed by end-to-end automation management
- the behavior of the end-to-end automation manager

You specify the automation policy in an XML file. In the XML policy file, you make the following specifications:

- You define the resources that are to be managed by the end-to-end automation manager, namely, resource references, resource groups, and choice groups.
- You can define the default desired states, that is, the default automation goals that the end-to-end automation manager is to pursue.
- You define StartAfter, StopAfter, and ForcedDownBy relationships.

This chapter describes all the required steps for defining a policy. It is intended to serve as a roadmap that guides you through the process of policy definition. The following table lists the tasks that you need to perform in the recommended sequence and points you to the related description:

| Step | Task | Description | Associated topics and procedures |
|------|---|--|---|
| 1 | Identify candidate clusters and sysplexes, and the resources that are candidates for end-to-end automation management | Identify the first-level automation clusters and sysplexes that host resources that have relationships, and the relevant resources. You may want to complete this task in close cooperation with the persons responsible for the first-level automation domains. | "The scope of end-to-end automation policies" on page 162 |
| 2 | Identify relationships or group dependencies | Identify the relationships or group dependencies of the resources running on the sysplexes and clusters | "The scope of end-to-end automation policies" on page 162 |

| | | | |
|---|---|--|--|
| 3 | Gather information about the first-level automation resources | When you create the XML policy file in a later step, you will need resource-specific data, for example, the name of the resource, the name of the first-level automation domain it belongs to, its class, and the node on which it resides. | “The scope of end-to-end automation policies” and “Worksheet for gathering data for defining a policy” on page 265 |
| | | In addition, you should gather information about who can be contacted in case of problems, for example, the name and phone number of the person who is responsible for the resource. You should provide a short description of the resource, and, if at all possible, a URL where more information about the resource can be obtained. | |
| 4 | Define the automation policy in an XML file | Use a suitable XML editor or text editor to create the XML file and define the automation policy using the data you have collected in the previous steps. | “Defining an end-to-end automation policy” on page 168 |

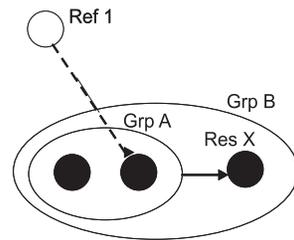
What you must know before you define an end-to-end automation policy

The scope of end-to-end automation policies

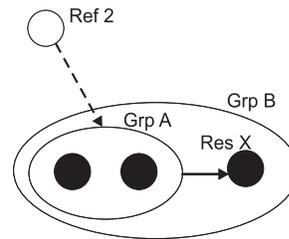
As described in Chapter 1, “What end-to-end automation management can do for you,” on page 3, end-to-end automation management is not intended to take over the role of first-level automation products. The main focus of first-level automation products is on ensuring the high availability of applications within a cluster of systems. This task must remain as close as possible to the resources for which high availability is to be ensured.

The scope of end-to-end automation policies starts where local first-level automation capabilities end - on the border of a first-level automation cluster. Consequently, end-to-end automation policies should only define cluster-spanning relationships and groups. The following examples provide some information on what you must consider when defining resource references for first-level automation resources.

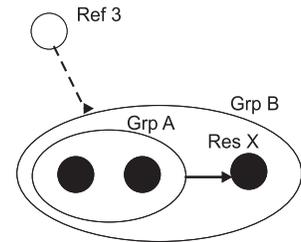
Example 1



Example 2



Example 3



The examples in the figure above show three resource references that were created for resources or resource groups that are hosted by a first-level-automation domain. These examples are described in the following sections.

Example 1

This example illustrates why it is not desirable to create resource references pointing to resources that are members of first-level automation groups if the integrity of first-level automation is to be ensured.

For this scenario, assume that:

- Resource reference "Ref 1" references an actual resource which is a member of the first-level automation domain group "Grp A".
- In the end-to-end automation policy, the desired state Online is defined for resource reference "Ref 1".
- In the first-level automation policy, the desired state Offline is defined for both "Grp A" and "Grp B".

When the end-to-end automation policy is activated, the end-to-end automation manager issues an Online request against the first-level automation resource that is referenced by "Ref 1". The first-level automation manager receives the request. If the referenced resource is offline, it will try to start the application.

If the referenced resource is started due to the request from the end-to-end automation manager, the observed state of "Grp A" changes accordingly. "Grp A" has been defined to be offline. This goal cannot be accomplished by the first-level automation manager because the request on the group member has a higher priority and will be fulfilled. As a result, the compound state of "Grp A" changes, indicating that a problem has occurred. The same is true for "Grp B".

An additional problem occurs because of the dependency between "Grp A" and the first-level automation resource "Res X". The administrator who created the first-level automation policy may have assumed that the relationship to "Res X" would always be evaluated before a member of "Grp A" is started. In such a scenario, however, this is not the case and the dependency will not be honored.

Example 2

In this example, resource reference "Ref 2" refers to "Grp A" which is hosted by the same first-level automation domain. This has the following two advantages over the constructs in Example 1:

1. All members of "Grp A" will be started or stopped in accordance with the desired group behavior. After the completion of the request from the end-to-end automation manager, "Grp A" changes to a normal end state and no problem will be indicated on the operations console.
2. The relationship to "Res X" will be evaluated when the request is sent to "Grp A". This ensures that all required actions will be performed by the first-level automation manager as defined by the administrator of the policy.

Only one problem remains: First-level automation cannot reach the desired state defined in the policy for "Grp B". However, in certain circumstances, referencing "Grp A" may reflect the desired behavior within in the scope of end-to-end automation. In such a case, the operator must understand that "Grp B" is in a problem state because end-to-end automation needed to start a member of this group in order to accomplish an end-to-end business goal.

Example 3

The two examples above show that creating an end-to-end automation policy which defines "Ref 3" will cause the least amount of undesired behavior. In this scenario, "Ref 3" references the outermost (or top-level) resource group defined in the first-level automation policy. No matter what desired state has been defined for "Ref 3", the first-level automation manager will act according to the request it receives from the end-to-end automation manager and all of the constructs defined in the first-level automation policy will remain in a satisfactory state.

Identifying cluster-spanning dependencies

This chapter is intended to give some advice on how to identify first-level automation resources that have cluster-spanning relationships. Such resources are candidates for being referenced in the end-to-end automation policy.

Two kinds of dependencies can be expressed in the constructs of an end-to-end automation policy:

1. Grouping concept: defines the general structure of resources and resource groups
2. Relationship concept: represents run-time dependencies between resources and resource groups

The following sections describe how you can find groups and relationships among automated resources that are hosted by different first-level automation domains.

Grouping of resources

Questions to ask:

- Which of the resources that are automated by different first-level automation domains need to be available at the same time?
- Which of the resources that are automated by different first-level automation domains can act as alternatives for other resources in case these fail?
- Which resources should be grouped together to ensure that their state can be easily monitored? For example, a group could comprise all resources that will be monitored by the same operator even if the resources are hosted by different first-level automation domains.

An enterprise application consists of multiple resources (for example, applications and IP addresses) that can belong to different business tiers and areas of responsibility.

In order to automate resources effectively, the resources need to be restructured from a technical and organizational point of view. This is why the grouping concept is introduced in end-to-end and first-level automation.

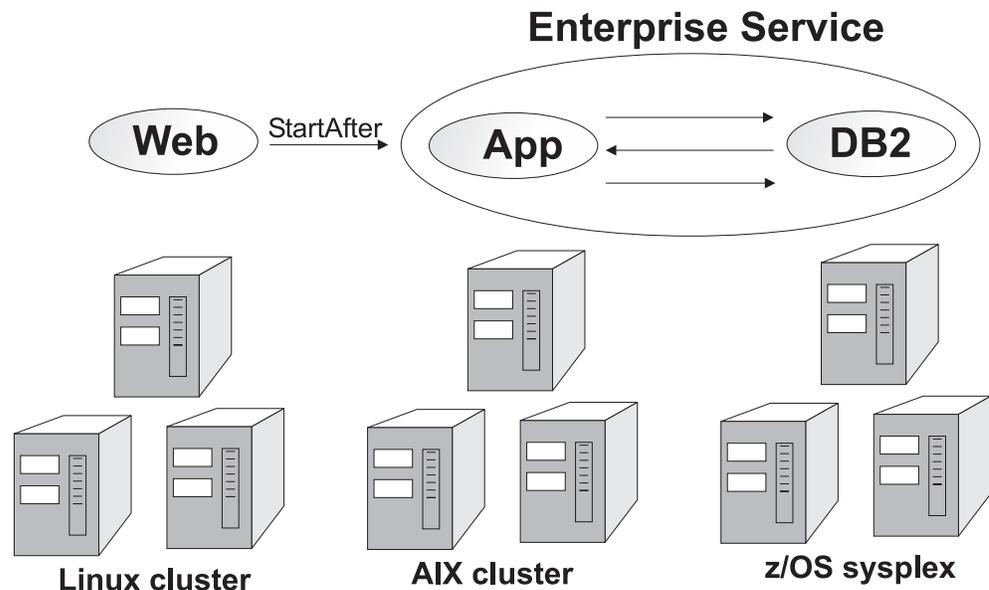
Organizing resources in groups has the following benefits:

- Groups are logical containers that can be controlled as one logical instance.
- Groups organize the automated resources in a hierarchical structure.
- A group can be composed of resource references and other end-to-end automation groups. The possibility of nesting groups allows you to structure complex environments into several layers.
- By encapsulating resources and nested groups within groups, you can organize your automated resources in a hierarchical structure that serves as the logical basis for an end-to-end automation policy.

Resources can be gathered in groups according to logical, technical, security, or responsibility criteria. For example:

- A resource group can be made up of resource references that reference all resources in an SAP environment
- A group can include all resources that have the same owner

End-to-end automation groups can be platform-spanning. This means that resource references for resources that are hosted by different first-level domains can be gathered in one group. As shown in the illustration below, the resource references that refer to a DB2 group on a first-level z/OS sysplex can be gathered in a group together with the application "App", which is physically hosted on an AIX cluster.



Relationships

Questions to ask:

- Which automated resource on a specific first-level automation domain needs which other resource on another automation domain in order to run?
- What are typical tasks for an operator to start or stop applications in order to start or stop some solution? Are workflow documents available which describe the sequence in which applications need to be started or stopped?

- How does an operator apply maintenance to specific applications? Are documents available that describe in which sequence an operator must shut down applications?
- In case of an unexpected failure of some critical applications on a first-level automation domain, do other applications on other automation domains need to be stopped as well?

Relationships represent dependencies between resources or groups. A relationship exists between a source and a target. Source and target can be either resource references or groups. For example, a relationship A StartAfter B ensures that resource A can only start when resource B is online.

Before you define a policy, you need to identify the relationships between the resources. When you identify the relationships that need to be defined in the policy, you should list the relationship information in the following sequence:

- source resource
- first-level automation domain name
- target resource
- first-level automation domain name
- relationship type

Example scenario: Stopping of a resource is triggered by the shutdown of another resource: The following example describes when a ForcedDownBy relationship between two resources is required.

In the description below, the following desired states are assumed for Resource A and Resource B:

- Resource A has the default desired state Online.
- Resource B has the default desired state Offline.

You need to define a ForcedDownBy relationship between source resource Resource A and target resource Resource B (Resource A ForcedDownBy Resource B) if you want to achieve the following behavior:

- Whenever Resource B is started, for example, due to an operator request, this should not have any effect on Resource A.
- Whenever Resource B was online and is stopping, for example, after it was started due to an operator's Online request and the request is canceled, or when Resource B fails while it is offline, Resource A must be bounced, that is, it has to be stopped and restarted again, for example, to allow Resource A to synchronize with Resource B.

Gathering the required data for defining a policy

This is the information you need for defining a policy:

- Resource identification data (for example, Name, Class, Location)
- Resource descriptions (Owner, InfoLink, short description)
- Information about cross-cluster relationships

Additionally, you should establish ownership for end-to-end automation resources and groups.

A worksheet for gathering the data you need for defining a policy is available in "Worksheet for gathering data for defining a policy" on page 265.

Considerations for referencing SA for Multiplatforms resources

When you create resource references for SA for Multiplatforms resources, the following considerations apply:

- Creating resources references for fixed resources that are constituents of a floating resource is not recommended because such resources cannot be controlled by end-to-end automation management and they can only be monitored but not managed from the operations console.
- You should avoid creating resource references for individual members of a SA for Multiplatforms group. For information about the effects that referencing such resources may have, refer to “The scope of end-to-end automation policies” on page 162.

Table 32. Recommendations for referencing SA for Multiplatforms resources in end-to-end automation policies

| RSCT classes IBM.* used in SA for Multiplatforms | Valid | Recommended |
|--|-------|-------------|
| IBM.NetworkInterface | X | |
| IBM.ResourceGroup | X | X |
| IBM.Equivalency | X | |
| IBM.Application | X | |
| IBM.ServiceIP | X | |
| IBM.Test | X | |

Restrictions for referencing SA z/OS resources

Resource references should not be created for the following SA z/OS resources:

- Resources that have external startup or shutdown set to ALWAYS should not be referenced.
The reason is that requests that are generated against such a resource reference always fail. As a result, the state of such a resource reference changes to Unrecoverable error as soon as the end-to-end automation manager generates the initial request after policy activation. For such resource references, the state cannot be resolved by using the Reset function.
- Passive application groups should not be referenced because operator requests against such resource references cannot be canceled from the operations console.
- Resources which have an agent or the manager automation flag set to NO should not be referenced because operator requests against such resource references cannot be canceled from the operations console in most cases.
- Resources for which the NOSTART option is specified during the agent start should not be referenced because the end-to-end automation manager will not honor the option.
This means that when the resource reference had the desired state Online, the referenced resource would be started after agent startup although the NOSTART option was specified.

Defining an end-to-end automation policy

When you have gathered the data for a new policy as described above, it is recommended that you complete the steps that are required for creating the policy in the following sequence:

Table 33. Steps for defining a new end-to-end automation policy

| Step | Task | This is where the task is described |
|------|--|--|
| 1 | Create the XML policy file | "Creating the XML policy file" |
| 2 | Define the resources of the end-to-end automation domain | "Defining the resources of the end-to-end automation domain" on page 171 |
| 3 | Define resource groups and choice groups | "Defining groups" on page 174 |
| 4 | Define StartAfter, StopAfter, and ForcedDownBy relationships | "Defining StartAfter, StopAfter, and ForcedDownBy relationships" on page 177 |

Notes:

1. To ensure that your XML policy file remains readable and maintainable, structure your file carefully by dividing it into sections. The following structure is recommended:
 - a. Resource references
 - b. Groups
 - c. RelationshipsYou can use comments in the policy file to separate the sections within the file.
2. An example of a complete XML policy file is provided in "Sample automation policy" on page 282.
3. Do not edit an XML policy file in the policy pool directory. Always use a copy of the XML file, edit it in a working directory, and update the PolicyToken before you save the file to the policy pool directory.
4. The following chapters assume that you have a good basic knowledge of XML.

Creating the XML policy file

This section describes the basic elements an XML policy file contains. Some of these elements are required and the policy cannot be activated if they are omitted. Some of the optional elements should not be omitted because they can be used to provide important meta-information about the policy (for example, the name of the owner of the policy and the date when the policy was last changed).

When you create an XML file with just the elements described in this section, you have a template you can use to create XML policy files. However, it is recommended that you use the official XML policy file template that you find in the following directory:

```
EEZ_INST_ROOT/policyPool/template.xml
```

To use the template, copy the file to your working directory and rename it according to your file naming conventions.

To create the XML policy file, you can use any commercial, shareware, or free-ware XML or ASCII editor as long as the editor allows you to save the file in UTF-8 format. XML files in any other format will be rejected by the policy checking tool.

If you use an XML editor to create the XML policy file, the editor will create the basic XML policy template for you. Additionally, most XML editors have a validation function that ensures that your XML code conforms to the relevant schema. When you want to use these functions, you must ensure that the XML editor knows where to find the relevant schema. This is where the schema for the end-to-end automation policy files is located:

EEZ_INST_ROOT/policyPool/EEZPolicy.xsd

Here is an example of the basic elements that all policy documents should contain (the required elements are marked in **bold**):

```
<?xml version="1.0" encoding="UTF-8"?>
<AutomationPolicy version="1.0"
  xmlns="http://www.ibm.com/TSA/Policy.xsd"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.ibm.com/TSA/Policy.xsd EEZPolicy.xsd ">
  <PolicyInformation>
    <PolicyName>Sample E2E Policy</PolicyName>
    <AutomationDomainName>FriendlyE2E</AutomationDomainName>
    <PolicyToken>1.0.1</PolicyToken>
    <PolicyAuthor>Michael Atkins</PolicyAuthor>
    <PolicyDescription>
      Policy for the end-to-end automation domain FriendlyE2E.
      Last Update: 09/16/05
      Last Editor: Michael Atkins
      Change History:
      -----
      Date      Name      Description
      -----
      09/16/05  Michal Atkins  Initial Policy
      -----
    </PolicyDescription>
  </PolicyInformation>
  ...
</AutomationPolicy>
```

The elements have the following meaning:

XML declaration

The XML file must begin with the following XML declaration and the encoding statement:

```
<?xml version="1.0" encoding="UTF-8"?>
```

Element AutomationPolicy

The complete XML policy must be enclosed in an `<AutomationPolicy>` element. The closing tag `</AutomationPolicy>` must be the last element and the last line in the XML policy file.

Use the following declarations in your policy file:

```
<AutomationPolicy version="1.0"
  xmlns="http://www.ibm.com/TSA/Policy.xsd"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.ibm.com/TSA/Policy.xsd EEZPolicy.xsd">
  ...
</AutomationPolicy>
```

The four attributes of the `AutomationPolicy` element and their values must be specified exactly as shown in the example above. Here is an explanation of what the attributes specify:

version

This is the minimum version of the end-to-end automation management component required for this policy.

xmlns This is the name space declaration.

xmlns:xsi

This is the XML schema format used for this XML policy.

xsi:schemaLocation

This is the XML schema that defines the XML syntax to which this policy XML file must conform and against which the policy checking tool checks the validity of the XML file before you can activate the policy.

Element PolicyInformation and its subelements

You use the element PolicyInformation and its children to provide important information about the policy.

The element is required and must occur only once in an XML policy file.

PolicyInformation has three required subelements that uniquely identify the policy (namely, PolicyName, AutomationDomainName, and PolicyToken).

Two additional subelements (namely, PolicyAuthor and PolicyDescription) are optional but declaring them and maintaining them carefully throughout a policy's life-cycle simplifies the maintenance and administration of XML policy files.

Here is an example of a PolicyInformation definition (the required elements are marked in **bold**):

```
<PolicyInformation>
  <PolicyName>Sample E2E Policy</PolicyName>
  <AutomationDomainName>FriendlyE2E</AutomationDomainName>
  <PolicyToken>1.0.1</PolicyToken>
  <PolicyAuthor>Michael Atkins</PolicyAuthor>
  <PolicyDescription>
    Policy for the end-to-end automation domain FriendlyE2E.
    Last Update: 09/16/05
    Last Editor: Michael Atkins
    Change History:
    -----
    Date      Name      Description
    -----
    09/16/05  Michal Atkins  Initial Policy
    -----
  </PolicyDescription>
</PolicyInformation>
```

These are the required subelements of PolicyInformation:

PolicyName

Assign a meaningful name to your policy. When you have more than one policy in your policy pool directory, especially if you change policies frequently, a meaningful PolicyName makes it easy to determine the policy's purpose and usage.

The PolicyName can have up to 64 characters.

AutomationDomainName

This is the name of the end-to-end automation domain for which the policy will be used. The automation domain name is specified in the end-to-end automation configuration dialog (page Domain,

field **Domain name**). Only if the domain name in the policy file matches the domain name on the configuration dialog page will the policy be accepted for activation. The AutomationDomainName can have up to 64 characters.

PolicyToken

Careful versioning of policy files is important to be able to keep track of your changes. You use the PolicyToken element to identify the version in the XML policy file. The format is optional. The policy checking tool will only verify that the PolicyToken element is available in the XML policy file. The content will not be checked. The PolicyToken can have up to 64 characters.

The PolicyInformation element has these optional subelements:

PolicyAuthor

Use this element to identify the author of the policy. A maximum of 64 characters is supported.

PolicyDescription

This element may contain free text, for example, comments, or a table for the change history as in the example above. A maximum of 1024 characters is supported.

When you have created an XML file with the elements described above, you should give the file a meaningful name and save it to your working directory before you start defining the resources of the end-to-end automation domain in the file.

Defining the resources of the end-to-end automation domain

You define the resources of the end-to-end automation domain by declaring a ResourceReference element for each first-level automation resource that you want to include in end-to-end automation management.

This is an example of a complete resource reference definition (the elements marked in **bold** are required for resource references pointing to actual resources that are managed by SA for Multiplatforms or SA z/OS):

```
<ResourceReference name="Enterprise DB2">
  <DesiredState>Offline</DesiredState>
  <Description>Database Enterprise DB2 on FEPLEX2</Description>
  <Owner>Bob Owens
    phone: 555-3677
    e-mail: b.owens@example.com
  </Owner>
  <InfoLink>http://www.example.com/help/DB2</InfoLink>
  <ReferencedResource>
    <AutomationDomain>FEPLEX2</AutomationDomain>
    <Name>DB2</Name>
    <Class>ResourceGroup</Class>
    <Node>node1</Node>
  </ReferencedResource>
</ResourceReference>
```

To create a resource reference, you need the following information about the first-level automation resource it points to (the so-called referenced resource):

- The name of the first-level automation domain that hosts the resource.
- The name by which the resource is known in the first-level automation domain.

- The Class element is optional. In some cases, however, the class to which the resource belongs must be specified.
- The Node element is optional. Only specify the Node element when creating a resource reference for a fixed resource. Do not specify the node for any other type of first-level automation resource.

Here is a description of the element ResourceReference and its subelements:

ResourceReference

This is the element that will be used to create the end-to-end automation resource that will be managed by the end-to-end automation manager and that can be monitored and managed by the end-to-end automation operator from the operations console.

The name you define for the resource in its **name** attribute must be unique within the policy, the same name cannot be used for another ResourceReference, ResourceGroup, or ChoiceGroup in the policy.

As operators can set name filters to see only selected resources in the resource table of the operations console, your naming conventions for resource references should support filtering by name, for example, by using common prefixes.

The name can have a maximum of 64 characters. Do not use more than one blank to separate strings within the name. Duplicate blanks will be ignored.

Description

Use this element to enter a description of the resource.

The description will appear on the operations console when an operator selects the resource in the resource table. The element is optional. The free text you type can have up to 1024 characters.

Owner

Use this element to enter the name of the owner of the resource and to provide information on how the owner can be contacted.

The information will appear on the operations console when an operator selects the resource in the resource table. The element is optional. The owner information you provide can have up to 1024 characters.

InfoLink

Use this optional element to specify a URL that points to additional information about the resource, for example, to an HTML page. The link will be available on the operations console when an operator selects the resource in the resource table.

The URL can have up to 1024 characters.

DesiredState

You can use this element to define the default desired state for the resource reference. Valid states are *Online* and *Offline*.

The element DesiredState is optional. The default value is Online. For information on how the default desired state of a resource is calculated when it is a member of a reference group or choice group, refer to Chapter 5, “Automation concepts,” on page 27.

ReferencedResource

ReferencedResource is a container element. You use its subelements to specify which first-level automation domain resource or resource group is to be included in end-to-end automation management.

The element `ReferencedResource` consists of the subelements `AutomationDomain`, `Name`, `Class`, and `Node`.

Here is an example of a resource reference for a resource that is managed by SA for Multiplatforms (required elements are marked in **bold**):

```
<ReferencedResource>
  <AutomationDomain>FEClusterSAP</AutomationDomain>
  <Name>SAP AppServer</Name>
  <Class>IBM.Application</Class>
  <Node>node1.ibm.com</Node>
</ReferencedResource>
```

Here is an example of a resource reference for a resource that is managed by SA z/OS (required elements are marked in **bold**):

```
<ResourceReference name="NFS Server">
  <DesiredState>Offline</DesiredState>
  <Description> Resource reference NFS Server </Description>
  <Owner>Bob Owens</Owner>
  <InfoLink>file://X:/help/NFS.pdf</InfoLink>
  <ReferencedResource>
    <AutomationDomain>FEPLEX1</AutomationDomain>
    <Name>NFS Server</Name>
    <Class>ResourceGroup</Class>
    <Node>node3</Node>
  </ReferencedResource>
</ResourceReference>
```

The subelements of `ReferencedResource` have the following meaning:

AutomationDomain

Use this element to specify the name of the first-level automation domain that hosts the referenced resource.

The domain name can have up to 64 characters.

The element is required.

Name This is the name by which the referenced resource is known in its first-level automation domain. The name can have up to 64 characters.

The element is required.

Class This is the resource class of the referenced resource in the first-level automation domain. The name of the resource class can have up to 64 characters. The element is optional, but must be defined for resources that are automated by SA for Multiplatforms or SA z/OS.

Node This is the name of the host (SA for Multiplatforms) or the name of the system (SA z/OS) in the first-level automation domain on which the referenced resource is located.

Restrictions:

- Maximum number of characters supported: 256
- Host name or system name must be specified in first-level automation domain syntax.
- The `Node` element is optional. Only specify the `Node` element when you create a resource reference for a fixed resource. Do not specify the node for any other type of first-level automation resource.

Note that creating resource references for fixed resources is **not** recommended.

Defining groups

You can define two different types of groups:

Resource groups

You use a resource group to gather resources in one group that share these characteristics:

- They are functionally related (for instance, they are components of a distributed business application).
- They have the same desired state (either *Online* or *Offline*) and should be managed and monitored as one unit.
- Typically, the members of a resource group are hosted by different first-level automation domains.

For information on how you define a resource group in an XML policy file, see “Defining resource groups.”

Choice groups

Choice groups make it easy to manage alternatives of redundant applications or application groups. For example, operators can switch from the production setup to the test setup of an application or application group without having to know how the applications are started or stopped.

Choice groups ensure that only one member of the group (the preferred member) is online at any given time. When an operator switches to an alternative, end-to-end automation management ensures that the old preferred member is brought into an offline state and is stopped before the new preferred member is started.

For information on how you define a choice group in an XML policy file, see “Defining choice groups” on page 176.

Defining resource groups

This is an example of a resource group definition in an XML policy file (the required elements are marked in **bold**):

```
<ResourceGroup name="Friendly Computer Shop" >
  <DesiredState>Online</DesiredState>
  <Description>Resource group Friendly Computer Shop</Description>
  <Owner>Jerry Owens</Owner>
  <InfoLink>http://www.example.com/help/policy/compshop.html</InfoLink>
  <Members>
    <ResourceGroup name="mySAP Solutions"/>
    <ResourceReference name="WebSphere AE"/>
  </Members>
</ResourceGroup>
```

The elements have the following meaning:

ResourceGroup

This is the element that will be used to create an end-to-end automation resource group.

Members of a resource group can be other resource groups or resource references.

The name you define for the resource group in its **name** attribute must be unique within the policy, the same name cannot be used for any other ResourceGroup, ChoiceGroup, or ResourceReference in the policy.

As operators can set name filters to see only selected resources in the resource table on the operations console, your naming conventions for resource groups should support filtering by name.

The name can have a maximum of 64 characters. Do not use more than one blank to separate strings within the name. Duplicate blanks will be ignored.

Note:

- Resource groups can be nested, but one resource group cannot be a member of more than one resource group.
- Making a choice group a member of a resource group is not recommended. If you do, a warning will be issued during policy activation.

The ResourceGroup element has the following subelements:

DesiredState

You can use this element to define the default desired state for the resource group. Valid states are *Online* and *Offline*.

The element is optional. You only need to define the desired state if the resource group is to be kept offline. When you do not define the desired state here, the default value (Online) is used.

Description

Use this optional element to provide a description of the resource group. The description will appear on the operations console when an operator selects the resource group. The description can have up to 1024 characters.

Owner

Use this optional element to enter the name of the owner of the resource group and to provide information on how the owner can be contacted. The information will appear on the operations console when an operator selects the resource group. The owner information you provide can have up to 1024 characters.

InfoLink

Use this optional element to specify a URL that points to additional information about the resource, for example, to an HTML page. The link will be available on the operations console when an operator selects the group. The link can have up to 1024 characters.

Members

You use this container element to define which of the resource references or resource groups that you have defined in the policy make up the resource group. To define the members, you must use the element definition for the resource reference or resource group that is to become a member of the group.

```
<Members>  
  <ResourceGroup name="mySAP Solutions"/>  
  <ResourceReference name="WebSphere AE"/>  
</Members>
```

Note: A resource reference that is a member of a resource group cannot be a member of a choice group.

Defining choice groups

This is an example of a choice group definition in an XML policy file (the required elements are marked in bold):

```
<ChoiceGroup name="HTTP Server">
  <DesiredState> Offline </DesiredState>
  <Description>Choice group for choosing one HTTP Server</Description>
  <Owner>Jenny Parker</Owner>
  <InfoLink>http://www.example.com/choice</InfoLink>
  <Members>
    <ResourceReference name="HTTP Server Prim" preferred="true"/>
    <ResourceReference name="HTTP Server Backup"/>
  </Members>
</ChoiceGroup>
```

The elements have the following meaning:

ChoiceGroup

This is the element that will be used to create a choice group.

Resource groups and resource references can be members of a choice group.

The name you define for the choice group in its **name** attribute must be unique within the policy, the same name cannot be used for another ChoiceGroup, ResourceReference, or ResourceGroup in the policy.

As operators can set name filters to see only selected resources in the resource table of the operations console, your naming conventions for choice groups should support filtering by name.

The name can have a maximum of 64 characters. Do not use more than one blank to separate strings within the name. Duplicate blanks will be ignored.

Notes:

1. Making a choice group a member of a resource group is not recommended. If you do, a warning will be issued during policy activation.
2. Making a choice group a member of another choice group is not recommended. If you do, a warning will be issued during policy activation.

The ChoiceGroup element has the following sub-elements:

DesiredState

The DesiredState is the automation goal that the automation manager will try to achieve. Valid states are *Online* and *Offline*.

For choice groups that are to be kept online, the element is optional, because *Online* is the default that will be used when you do not declare the desired state in the XML file.

When the desired state is Online, the automation manager will try to keep the so-called preferred member of the group online and will try to keep the other member or members offline.

When the desired state is Offline, you must declare the DesiredState element. Then the automation manager will try to keep all members of the group offline.

Description

Use this optional element to enter a description of the choice group. The description will appear on the operations console when an operator selects

the choice group but will also facilitate the maintenance of the policy document itself. The free text you type can have up to 1024 characters.

Owner

Use this optional element to enter the name of the owner of the choice group or of the resources that make up the choice group and to provide information on how the owner can be contacted. The information will appear on the operations console when an operator selects the choice group. The owner information you provide can have up to 1024 characters.

InfoLink

Use this optional element to specify a URL that points to additional information about the choice group, for example, to an HTML page. The link will be available on the operations console when an operator selects the choice group. The link can have up to 1024 characters.

Members

You use this container element to define which of the resource references or resource groups that you have defined in the policy make up the choice group.

To define the members, you must use the element definition for the resource reference or resource group that is to become a member of the group. Additionally, one of the members in the list of group members must have the attribute *preferred="true"*. This is the member that will be kept online by the automation manager if the desired state of the choice group is *Online*. For all other members, the attribute can be omitted, because the default is *false*.

```
<Members>
  <ResourceReference name="HTTP Server Prim" preferred="true"/>
  <ResourceReference name="HTTP Server Backup"/>
</Members>
```

Note: A resource reference that is a member of a choice group cannot be a member of a resource group.

Defining StartAfter, StopAfter, and ForcedDownBy relationships

Defining a StartAfter relationship

This is an example where IMS Connect is started first when a start request is submitted against Banking Application:

```
<Relationship>
  <Source>
    <ResourceReference name="Banking Application"/>
  </Source>
  <Type>StartAfter</Type>
  <Target>
    <ResourceReference name="IMS Connect"/>
  </Target>
</Relationship>
```

The elements have the following meaning:

Source

This is container element that contains the resource reference or end-to-end automation group that can only be started if the resource or group that is specified in the Target element is online.

To define the source resource, use the ResourceReference, ResourceGroup or ChoiceGroup definition.

Type Type must be set to StartAfter.

Target This is container element that contains the resource reference or end-to-end automation group that will be automatically started first if an operator submits a start request against the resource or group that is specified in the Target element and the target resource is not online.

To define the target resource, use the ResourceReference, ResourceGroup or ChoiceGroup definition.

Defining a StopAfter relationship

This is an example where Banking Application is stopped first when a stop request is submitted against IMS Connect:

```
<Relationship>
  <Source>
    <ResourceReference name="IMS Connect"/>
  </Source>
  <Type>StopAfter</Type>
  <Target>
    <ResourceReference name="Banking Application"/>
  </Target>
</Relationship>
```

The elements have the following meaning:

Source

This is container element that contains the resource reference or end-to-end automation group that can only be stopped if the resource or group that is specified in the Target element is offline.

To define the source resource, use the ResourceReference, ResourceGroup or ChoiceGroup definition.

Type Type must be set to StopAfter.

Target This is container element that contains the resource reference or end-to-end automation group that will be automatically stopped first if an operator submits a stop request against the resource or group that is specified in the Target element and the target resource is not offline.

To define the target resource, use the ResourceReference, ResourceGroup or ChoiceGroup definition.

Defining a ForcedDownBy relationship

When two resources have a ForcedDownBy relationship, one of the resources is forced down by the automation manager if the other resource goes offline unexpectedly or is forced down itself.

This is an example where Banking Application is brought offline when IMS Connect goes offline unexpectedly:

```
<Relationship>
  <Source>
    <ResourceReference name="Banking Application"/>
  </Source>
  <Type>ForcedDownBy</Type>
  <Target>
    <ResourceReference name="IMS Connect"/>
  </Target>
</Relationship>
```

The elements have the following meaning:

Source

This is the container element that defines which resource reference or group will be forced offline if the target resource:

- goes offline unexpectedly after having been online, or
- fails, regardless of its former state

To define the source resource, use the ResourceReference, ResourceGroup or ChoiceGroup definition.

Type Type must be set to ForcedDownBy.

Target If the the resource reference or group contained in this container element goes offline unexpectedly or is forced down, this will trigger the force down of the source resource

To define the target resource, use the ResourceReference, ResourceGroup or ChoiceGroup definition.

Saving the policy in the policy pool directory

XML policy files must be saved to the policy pool directory. To find out where the policy pool directory is located, launch the configuration dialog, open the Domain page and click **Advanced**. The default is <EEZ_INSTALL_ROOT>/policyPool.

For the files in the policy pool directory, the following recommendations apply:

- Make backup copies of all XML policy files. The XML file in the policy pool directory and its backup copy must be identical.
- Do not modify an XML policy file in the policy pool directory, especially not the one in which the currently active policy is defined. If the automation engine needs to be restarted, it will reload the same policy file from the policy pool directory. If the policy file has been modified, problems may occur, especially, if the changes are incorrect or not valid.
- When you update an XML policy file, use a copy of the file to make the changes and update the PolicyToken tag in the policy file before you save it to the policy pool directory.

When you have saved the XML policy to the policy pool directory, you use the operations console to activate the policy. This is described in “Activating a policy” on page 249. When you try to activate a policy, the validity of the policy is checked automatically.

Alternatively, you can start the policy checking tool from a command line. This is described in the following section.

Starting the policy checking tool from a command line

Perform the following steps:

1. Open a command window and change the directory to EEZ_INSTALL_ROOT/bin.

-
2. Issue the following command to start the tool:

On Windows: eezpolicychecker.bat <policy_file_name>

On AIX and Linux: ./eezpolicychecker.sh <policy_file_name>

If the policy file you want to check is not in the policy pool directory, you must enter the fully qualified file name.



Chapter 26. Setting up information pages for operators

In the information area of the operations console, you can make an info link available for each resource and group. The operator can follow the link to display information pages that provide additional information about the automated application. For resources of the end-to-end automation domain, you define the URL of the link in the InfoLink element of the XML policy.

If you have not yet set up such information pages, here are some suggestions for what they could include:

- A description of the managed application
- Procedures for analyzing and fixing problems (for example, where the logs are located, what to look for in the logs, where to find check scripts)
- Information about the primary and secondary contacts for the application
- Information about service periods and service level agreements

Chapter 27. Using the command-line interface of the automation engine

You use the script files `eezdmn.bat` (on Windows systems) and `eezdmn.sh` (for AIX and Linux systems) for the following purposes:

- starting the automation engine

Note: The way in which you start the automation engine determines in which mode the operations console runs:

- To run the operations console in end-to-end automation mode, the automation engine must be started with the command `eezdmn` or `eezdmn -start`.
 - To run the operations console in first-level automation mode, the automation engine must be started with the command `eezdmn -co`.
- stopping the automation engine
 - monitoring its current state
 - refreshing its configuration at runtime

To perform these tasks, do the following:

1. Log in to the system on which the automation manager is installed.
2. Change the directory to `EEZ_INSTALL_ROOT/bin`
3. Enter the command for the function you want to use. The command has the following syntax:
 - **On Windows:**
`eezdmn <option>`
For example:
`eezdmn -shutdown`
 - **On AIX and Linux:**
`./eezdmn.sh <option>`
For example:
`./eezdmn.sh -shutdown`

Table 34 on page 184 provides an overview of the available options. A detailed description is provided in the following sections of this chapter.

Note: When the automation engine is running on a Windows server, the automation engine is stopped when you log off from Windows, switch to a different user ID, or set the system to **Stand by** or **Hibernate**. To ensure that end-to-end automation is active continuously, do not use any of these functions. To prevent unauthorized access, only lock your computer.

eezdmn options quick reference

The following table presents an overview of the options that are available for the command.

Table 34. Command line options for the automation engine

| Option | Short form | Description |
|-----------|------------|---|
| -start | | Starts the automation engine. This is the default that is used when no option is specified. |
| -shutdown | -shutd | Stops the automation engine. |
| -monitor | -m | Retrieves the current state of the automation engine. |
| -reconfig | -r | Refreshes the credentials the automation manager uses to contact referenced resources that are hosted by first-level automation domains. You must always invoke the command with this option when you have modified configuration properties in the configuration dialog. |
| -co | | Starts the automation engine in conversion-only mode. In this mode, only the EIF-to-JMS conversion functionality is activated, the process will not act as automation engine. End-to-end automation management will not be performed. You must invoke the command with this option if you want to run the operations console in first-level automation mode. |
| -xd | | Dumps internal information into a specified file. This debug option generates detailed information that IBM support can use for debugging the automation states of resources. |
| -? | | Displays the version identifier of the automation engine and a help text that lists the command options. |

eezdmn options

This section provides a detailed description of the options you can use with the **eezdmn** (Windows) or **eezdmn.sh** (AIX and Linux) command.

-start

The option **-start** is the default value that is used when you enter the command **eezdmn** without specifying an option. The command starts the automation engine. During startup, the automation engine reads-in and processes the configuration parameters you specified on the Domain and User credentials pages of the configuration dialog (see Chapter 29, “Configuring the end-to-end automation manager,” on page 193).

When the automation engine has started successfully, the end-to-end automation domain is displayed on the operations console. The domain has the name that is defined in the **Domain name** field on the Domain page of the configuration dialog.

When you start the automation engine for the first time after you installed the end-to-end automation management component, you must subsequently activate an end-to-end automation policy.

If a policy for the domain had previously been active, the last active policy will be reactivated automatically if it is found in the policy pool directory.

Note: After you start the automation engine, you will always receive the message that the automation engine is in IDLE state and that no policy is activated even if the last active policy is available in the policy pool directory. This is because it takes time to load the last active policy.

Return codes

The following table lists the return codes that are returned by the command **eezdmn -start**.

| Code | Meaning |
|------|---|
| 0 | The automation engine was started successfully or was already running. |
| 2 | Error: No valid license key was found on the system. The automation engine could not be started. |
| 8 | Error: Incorrect attributes were specified. The automation engine could not be started. |
| 9 | Error: The automation engine could not be started. Check the automation engine log file for details. |
| 10 | Severe error: Required components are missing or corrupted. The automation engine could not be started. |

-shutdown

Use the option **-shutdown** to stop the automation engine in a controlled way. When the automation engine is stopped, end-to-end automation for the resources that are defined in the end-to-end automation policy will stop as well.

If you stop an automation engine that was started in conversion-only mode, the contact to the first-level automation domains will be lost. In this case, events are no longer received and the state information that is displayed for the resources on the operations console will be outdated shortly after the engine has stopped.

Return codes

The following table lists the return codes that are returned by the command **eezdmn -shutdown**.

| Code | Meaning |
|------|--|
| 0 | The automation engine was stopped successfully. |
| 1 | The automation engine had already been stopped. |
| 8 | Error: Incorrect attributes were specified. The automation engine could not be stopped. |
| 9 | Error: The automation engine could not be stopped. Check the automation engine log file for details. |

-monitor

Use the option -monitor to retrieve information about the current state of the automation engine. When you issue the command, the following message is displayed:

State of the EEZ automation engine is: <state-related information>

where <state-related information> stands for one of the states described in the following table.

Table 35. Messages and return codes returned by the automation engine

| Code | State-related information in the message | Description |
|------|---|--|
| 1 | RUNNING – Policy is activated | This is the normal state after a policy has been activated and end-to-end automation is running. |
| 2 | STARTING – Automation engine is not ready yet | The automation engine is being started. It cannot be contacted as a domain yet. |
| 3 | STOPPING – Automation engine does not accept requests anymore | The automation engine is being stopped. |
| 4 | IDLE – No policy is activated | <p>The automation engine is running. Before end-to-end automation can start, a policy must be activated.</p> <p>After you start the automation engine, you will always receive the message that the automation engine is in IDLE state and that no policy is activated. This is because it takes time to load the last active policy. As soon as the policy is loaded, the state of the automation engine will change.</p> |
| 5 | Process is only converting EIF messages | This informational message appears when the automation engine was started in conversion-only mode (with the command line option -co). It indicates that the automation engine is running but end-to-end automation is not being performed. |
| 6 | NOT AVAILABLE – Automation engine probably not started | No contact to the automation engine can be established. It is assumed that it has not been started yet. |
| 7 | No state-related information is displayed. | Unknown |
| 8 | No state-related information is displayed. | Incorrect attributes were specified. The command could not be processed. |
| 9 | PROBLEM – See message log for details | Problems have been detected. Check the message log file for information on the problems that have occurred. If you cannot resolve the problems, contact IBM support. |

Table 35. Messages and return codes returned by the automation engine (continued)

| Code | State-related information in the message | Description |
|------|--|--|
| 10 | SEVERE – See message log for details | Severe problems have been detected. Check the message log file for information about the problems that have occurred. If you cannot resolve the problems, contact IBM support. |

-reconfig

Use the option `-reconfig` to activate new configuration settings. You must invoke the command with this option in the following cases:

- After modifying configuration properties in the configuration dialog (for more information, refer to Chapter 29, “Configuring the end-to-end automation manager,” on page 193).
- When a security exception was reported while the automation manager tried to access a first-level automation domain, and the problem has been resolved.

Return codes

The following table lists the return codes that are returned by the command `eezdmn -reconfig`.

| Code | Meaning |
|------|---|
| 0 | The automation engine was reconfigured successfully. |
| 8 | Error: Incorrect attributes were specified. The reconfiguration could not be performed. |
| 9 | Error: The automation engine could not be contacted, it may not be running. The automation engine must be up and running in order to be reconfigured. |

-co

Use this option to start the automation engine in conversion-only mode. This is required when you want to use the operations console in first-level automation mode, because in this case, the EIF-to-JMS functionality of the automation engine is required but end-to-end automation management must not be performed. For more information about using the operation console for first-level automation management only, refer to “The operations console is used in first-level automation mode” on page 24.

Return codes

The following table lists the return codes that are returned by the command `eezdmn -co`.

| Code | Meaning |
|------|--|
| 0 | The automation engine was started successfully or was already running. |
| 2 | Error: No valid license key was found on the system. The automation engine could not be started. |
| 8 | Error: Incorrect attributes were specified. The automation engine could not be started. |
| 9 | Error: The automation engine could not be started. Check the automation engine log file for details. |

| Code | Meaning |
|------|---|
| 10 | Severe error: Required components are missing or corrupted. The automation engine could not be started. |

-xd

Use this command option only when IBM requests debugging information for one or more resources that are hosted by the end-to-end automation manager. The command will dump the debugging information into a file.

When you enter the command, you must provide additional parameters. This is the complete syntax of the command:

```
eezdmn -xd ("*"|"<resource_name>[,<resource_name>"])"<name_of_dump_file>
```

The parameters have the following meaning:

- * Specify this parameter when you want to dump information about all resources of the end-to-end automation domain into the file <name_of_dump_file>. Depending on the number of resources defined in the active policy, the resulting dump file can be large.

<resource_name>

To only write information about specific resources to the file <name_of_dump_file>, list the names of all relevant resources, separated by commas, and enclose the list in quotation marks. This is an example of the syntax of such a command:

```
eezdmn -xd ("Resource_A,Resource_B")dump1.txt
```

Return codes

The following table lists the return codes that are returned by the command **eezdmn -xd**.

| Code | Meaning |
|------|--|
| 0 | The operation completed successfully. |
| 8 | Error: Incorrect attributes were specified. The operation could not be performed. |
| 9 | Error: The automation engine could not be contacted. Check the automation engine log file for details. |

-?

Use this option to display the following help text:

```
IBM Tivoli System Automation end-to-end automation engine
Version: 2.1.0.051501, NO_APARUsage:
```

```
eezdmn [option]
```

```
-START           Starts the automation engine
-SHUTDOWN      -SHUTD  Stops the automation engine
-MONITOR       -M      Displays the current state
-RECONFIG      -R      Re-configures the automation engine
-CO            Starts only the EIF2JMS conversion thread
-XD ("*" | "<RES_NAME>[,<RES_NAME>"])" <DUMPFIL>
                Dumps (all | specific) resources to a file
```

When no option is specified, start is used

Chapter 28. Starting and stopping

This chapter describes how to start and stop the subcomponents of the end-to-end automation management component and the applications needed for operating it:

- WebSphere Application Server for end-to-end automation management
Section “Starting and stopping WebSphere Application Server” describes how to start and stop the server.
- Automation J2EE framework
The automation J2EE framework is started or stopped automatically when WebSphere Application Server is started or stopped. Section “Starting and stopping the automation J2EE framework” on page 190 describes how to start and stop the framework manually from the WebSphere Application Server administrative console.
- Operations console
To be able to use the operations console and to display the online help for the console, both the Integrated Solutions Console server and the Eclipse Help System server must be started. This is described in section “Starting and stopping the operations console” on page 190.
- Configuration dialog of end-to-end automation management
For information about starting the configuration dialog, refer to Chapter 29, “Configuring the end-to-end automation manager,” on page 193.
- Automation adapters
An automation adapter must be started on each first-level automation domain hosting resources that are referenced in the end-to-end automation policy. For information on starting and stopping the automation adapters, refer to the adapter documentation for the first-level automation product.

Note: For starting and stopping the automation engine, you use the **eezdmn** command. For information on how to use the command, refer to Chapter 27, “Using the command-line interface of the automation engine,” on page 183.

Starting and stopping WebSphere Application Server

The WebSphere Application Server instance for end-to-end automation management is started in the same way as any other WebSphere Application Server instance. The following sections describe how you use the scripts to start or stop WebSphere Application Server.

Starting and stopping WebSphere Application Server on Windows

When you are running WebSphere Application Server on a Windows system, you usually start and stop WebSphere Application Server by clicking the relevant icons on your desktop. If the icons are not available, you can start and stop the server from the Windows Start menu:

start —> **Programs** —> **IBM WebSphere** —> **Application Server v6** —> **Profiles**
—> <profile> —> **Start the server**

Alternatively, you can use the start and stop scripts that are available in the directory `<was_root>\bin`:

- To start WebSphere Application Server, open a command prompt and issue the following command:

```
<was_root>\bin\startServer <ServerName>
```

For example:

```
C:\Program Files\IBM\WebSphere\AppServer\bin\startServer server1
```

- To stop WebSphere Application Server, open a command prompt and issue the following command:

```
<was_root>\bin\stopServer <ServerName> -user <user ID> -password <password>
```

Starting and stopping WebSphere Application Server on AIX and Linux

To start WebSphere Application Server on AIX and Linux systems, issue this command from a command line:

```
<was_root>/bin/startServer.sh <ServerName>
```

To stop WebSphere Application Server on AIX and Linux systems, issue this command from a command line:

```
<was_root>/bin/stopServer.sh <ServerName> -user <user ID> -password <password>
```

Starting and stopping the automation J2EE framework

The automation J2EE framework is started and stopped automatically when WebSphere Application Server is started or stopped.

Alternatively, you can start and stop the automation J2EE framework from the administrative console of WebSphere Application Server as you would any other application that is running in a WebSphere Application Server environment. The name of the automation J2EE framework on the console is EEZEAR.

Starting and stopping the operations console

To be able to use the operations console and to display the online help for the console, both the Integrated Solutions Console server and the Eclipse Help System server must be started. The following sections describe how you start and stop the servers.

Note: The Integrated Solutions Console server and the Eclipse Help System server must be started as described in the following sections regardless of the mode in which the operations console is to run.

You determine in which mode the operations console runs when you start the automation engine:

- To run the operations console in end-to-end automation mode, the automation engine must be started with the command `eezdmn` or `eezdmn -start`.
- To run the operations console in first-level automation mode, the automation engine must be started with the command `eezdmn -co`.

For more information, refer to Chapter 27, “Using the command-line interface of the automation engine,” on page 183.

Starting and stopping the operations console on Windows

How you start and stop the Integrated Solutions Console server and the Eclipse Help System Server on Windows depends on whether or not you are running the servers as Windows services.

The servers are running as Windows services

If you are running the servers as Windows services, use one of the following approaches:

- You can start and stop the servers from the Windows **Services** panel. These are the relevant entries in the services list:
 - CS01 (ID of the Integrated Solutions Console server)
 - HS01 (ID of the Eclipse Help System server)
- If you want to start or stop the servers from a command prompt when you are running the servers as Windows services, you must start and stop the servers separately. To ensure that the status of the servers is reflected in Windows Services, use the commands described below to start and stop the servers.

Note: Do not use the scripts StartEclipse.bat and StopEclipse.bat to start or stop the Eclipse Help System server, because then the status of the servers will not be reflected in Windows Services.

Starting the servers:

- To start the Integrated Solutions Console server, use this command:

```
<was_root>\bin\startserver ISC_Portal
```

For example:

```
C:\Program Files\IBM\WebSphere\AppServer\bin\startserver ISC_Portal
```

- To start the Eclipse Help System server, use this command:

```
<isc_runtime_root>\PortalServer\ISCEclipse\EclipseServiceStart.bat
```

For example:

```
C:\Program Files\IBM\ISC\PortalServer\bin\EclipseServiceStart.bat
```

Stopping the servers:

- To stop the Integrated Solutions Console server, use this command:

```
<was_root>\bin\stopserver ISC_Portal -user <user_ID> -password <password>
```

where <user_ID> and <password> are the user credentials of the Integrated Solutions Console administrator.

- To stop the Eclipse Help System server, use this command:

```
<isc_runtime_root>\PortalServer\ISCEclipse\EclipseServiceStop.bat
```

For example:

```
C:\Program Files\IBM\ISC\PortalServer\bin\EclipseServiceStop.bat
```

The servers are not running as Windows services

If you are *not* running the servers as Windows services, use the following commands to start or stop both the Integrated Solutions Console server and the Eclipse Help System server.

To start the servers, use this command:

```
<isc_runtime_root>\PortalServer\bin\startISC.bat ISC_Portal
```

For example:

```
C:\Program Files\IBM\ISC\PortalServer\bin\startISC.bat ISC_Portal
```

To stop the servers, use this command:

```
<isc_runtime_root>\PortalServer\bin\stopISC.bat ISC_Portal <user_ID> <password>
```

Starting and stopping the operations console on AIX and Linux

To start the Integrated Solutions Console server and the Eclipse Help System server, use this command:

```
<isc_runtime_root>/PortalServer/bin/startISC.sh ISC_Portal
```

For example:

```
/opt/IBM/ISC/PortalServer/bin/startISC.sh ISC_Portal
```

To stop the Integrated Solutions Console server and the Eclipse Help System server, use this command:

```
<isc_runtime_root>/PortalServer/bin/stopISC.sh ISC_Portal <user_ID> <password>
```

For example:

```
/opt/IBM/ISC/PortalServer/bin/stopISC.sh ISC_Portal iscadmin pw4iscadmin
```

Starting and stopping the automation engine

The **eezdmn** command and the command options you use for starting and stopping the automation engine are described in Chapter 27, “Using the command-line interface of the automation engine,” on page 183.

Chapter 29. Configuring the end-to-end automation manager

The basic configuration of the end-to-end automation manager is performed during the installation of the end-to-end automation management component. The configuration properties are in properties files. To browse or change the properties, you use the configuration dialog. You should not edit the properties files.

This chapter describes how you launch the dialog and provides an overview of the pages of the dialog. Detailed information about the configuration dialog and the configuration properties is available in the online help of the dialog.

Launching the configuration dialog

Before you begin:

The user ID you use to launch the dialog must meet the following requirements:

- It must be in same group as the user ID you used for installing the end-to-end automation management component. The group permissions for `cfgeezdmn.sh` must be set to EXECUTE.
- The user ID must have read-write permissions on the following directory:
 - **Windows:** `<EEZ_INSTALL_ROOT>/cfg`
 - **AIX and Linux:** `/etc/<EEZ_INSTALL_ROOT>/cfg`

Perform the following step to launch the configuration dialog:

1. Log in to the system where end-to-end automation management is installed.

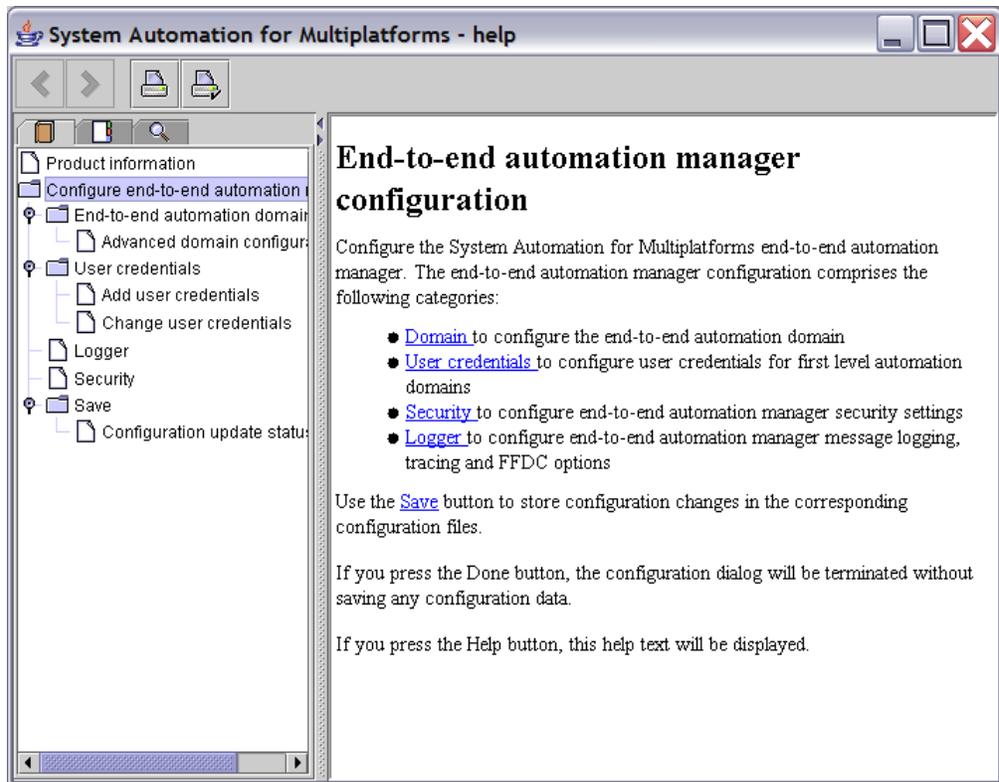
2. Change the directory to `<EEZ_INSTALL_ROOT>/bin`.

3. To launch the dialog, enter the following command:
 - **Windows:** `cfgeezdmn.bat`
 - **AIX and Linux:** `cfgeezdmn.sh`The configuration dialog is displayed.

After you have changed any of the configuration properties, you must activate the new settings by invoking the command `eezdmn` with the option `-reconf`. For more information about the command `eezdmn` and its options, refer to Chapter 27, "Using the command-line interface of the automation engine," on page 183.

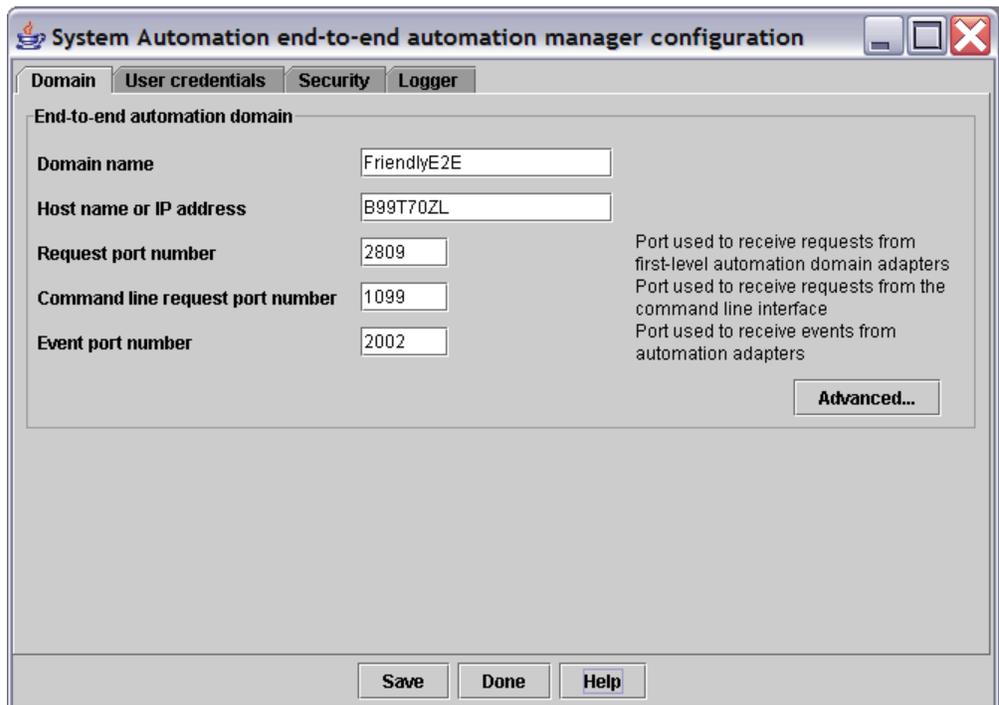
Configuration dialog pages

The following sections provide an overview of the pages that are available in the configuration dialog. Detailed information about the configuration dialog and the properties you can configure is available in the online help of the dialog. Click the **Help** button on a page of the dialog to launch the online help:



Domain page

Use the page to browse or change the configuration properties of the end-to-end automation domain.



Fields on the Domain page:

Domain name

The name of the end-to-end automation domain. The name specified here must be identical with the name specified in each XML policy file for the domain in the element <AutomationDomainName>.

Host name or IP address

Name or IP address of the system that host the end-to-end automation manager.

Request port number

The port on which the automation engine receives all requests from the automation manager.

Command line request port number

The port on which the automation engine receives command line interface requests.

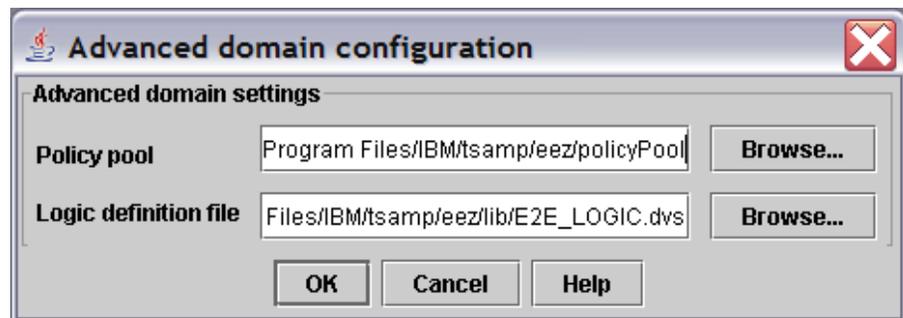
Event port number

The port on which the EIF message converter listens for events from the first-level automation domains. The port number specified here must match the port number specified in the properties files of the automation adapters on the first-level automation domains (property eif-send-to-port).

For SA for Multiplatforms, this is the Event port specified in the adapter configuration dialog.

Buttons on the Domain page:**Advanced**

Click **Advanced** to bring up the Advanced domain settings panel:



Use the Advanced panel to specify the paths to the policy pool directory and to the logic definition file.

When you click **Browse**, you can select the paths in the selection panel that appears.

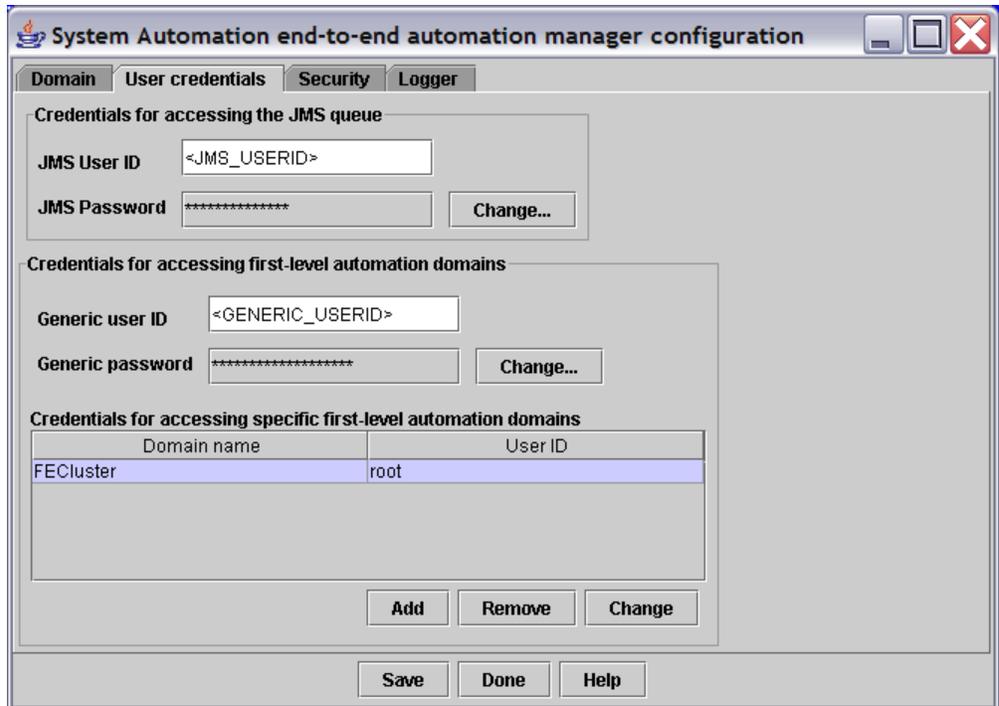
Save Click **Save** to save your changes.

Done Click **Done** to close the dialog.

Help Click **Help** to open the dialog help.

User credentials page

Use the page to browse or change the user credentials of the end-to-end automation manager. The automation manager uses these credentials to authenticate itself.



Fields on the User credentials page:

JMS User ID

The user ID that is used to access the JMS queue of the end-to-end automation manager.

JMS Password

The password for the JMS user ID.

Generic user ID

The user ID the automation manager uses to authenticate itself to a first-level automation domain when no credentials are specified for the domain in **Credentials for accessing specific first-level automation domains**.

Generic password

The password for the generic user ID.

Credentials for accessing specific first-level automation domains

User IDs and passwords for specific domains. Use the **Add**, **Remove** and **Change** buttons to create or modify the credentials of a first-level automation domain.

Buttons on the User credentials page:

Change

Click **Change** to open a dialog on which you can change the password.

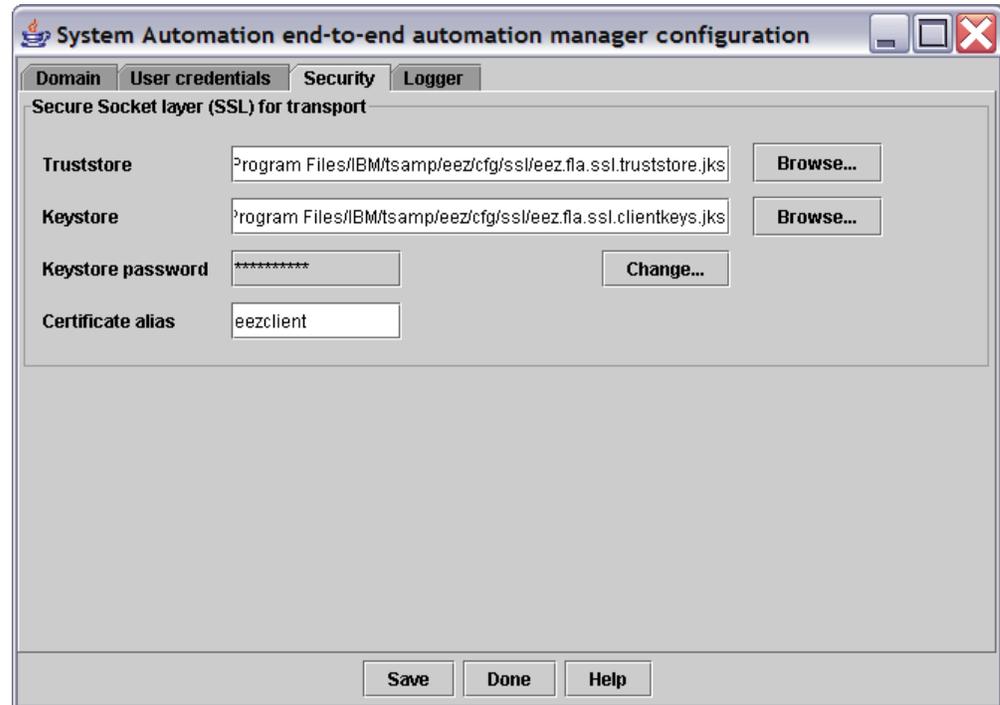
Save Click **Save** to save your changes.

Done Click **Done** to close the dialog.

Help Click **Help** to open the dialog help.

Security page

Use the page to browse and change the configuration properties for the SSL connection to the first-level automation domains.



Fields on the Security page:

Truststore

The fully qualified file name of the truststore file that is used for SSL.

Keystore

The fully qualified file name of the keystore file that is used for SSL.

Keystore password

The password of the keystore file.

Certificate alias

The alias name of the certificate to be used by the server.

Buttons on the Security page:

Browse

Click **Browse** to display a panel where you can select the file.

Change

Click **Change** to display a panel where you can change the keystore password.

Save Click **Save** to save your changes.

Done Click **Done** to close the dialog.

Help Click **Help** to open the dialog help.

Logger page

System Automation end-to-end automation manager configuration

Domain User credentials Security **Logger**

Maximum log/trace file size

Message logging level

Error

Warning

Information

Trace logging level

Off

Minimum

Medium

Maximum

Settings for first failure data capture (FFDC)

Recording level

Off

Minimum

Medium

Maximum

Disk space

Maximum disk space

Space exceeded policy

Ignore

Auto-delete

Suspend

Message IDs

Filter mode

Passthru Block

Message ID list

EEZD*E

Save Done Help

Use the Logger page to specify which information is written to the log and trace files of the automation engine. For detailed information on the properties you can specify on the page, see the dialog help.

Chapter 30. Using Tivoli Enterprise Console with SA for Multiplatforms

Configuring Tivoli Enterprise Console

If you have not activated or configured the Tivoli Enterprise Console (TEC) function during the installation of SA for Multiplatforms, you can do so by performing the following steps:

1. Activate the Common Event Infrastructure (CEI) service when server1 is started by selecting the check box **Enable service at server startup** on the administrative console of WebSphere Application Server:



Figure 8. Common Event Infrastructure Service panel

Save Master configuration and restart the WebSphere Application Server server1.

Note: Alternatively, you can start both CEI services manually:

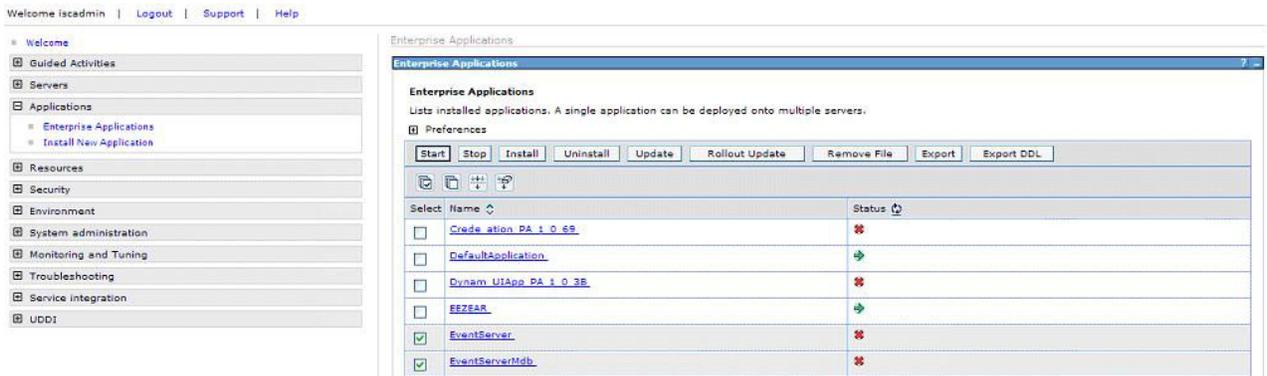


Figure 9. Enterprise Applications panel

Start the applications EventServer and EventServerMdb.

2. Navigate to the TEC settings:

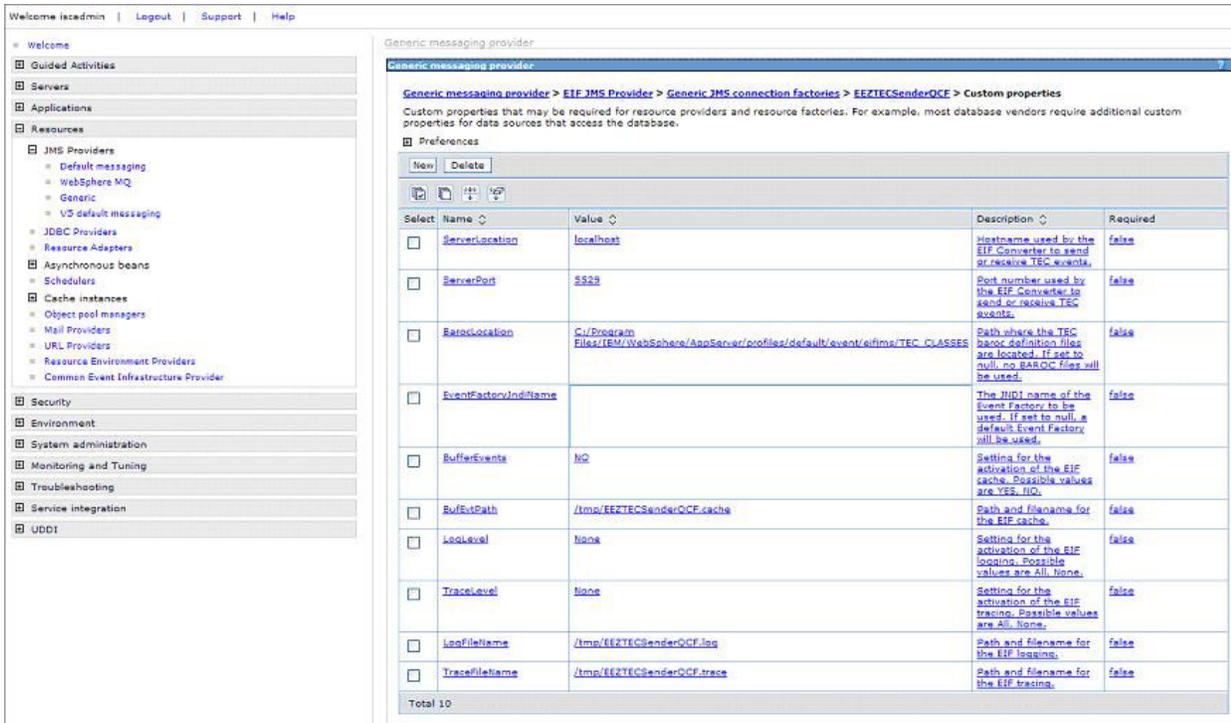


Figure 10. Custom properties panel

On the panel, do this:

- Set ServerLocation to your TEC server name.
- Set ServerPort to the following values:
 - TEC server runs on Windows: 5529
 - TEC server runs on AIX or Linux: 0
- Install the file SystemAutomation.baroc on your TEC server.

Checking the Tivoli Event Integration Facility function

This section describes how you verify that the Tivoli Event Integration Facility (EIF) is installed and configured correctly by sending an event to the event server. If the event appears on the Tivoli Enterprise Console, the configuration is correct.

Prerequisites:

- WebSphere Application Server is running
- The Tivoli Enterprise Console server is running
- Common Event Infrastructure (CEI) is installed
- EIF is installed
- CEI and EIF are configured:
 - In the WebSphere administrative console, navigate to **Resources > JMS Providers > Generic > EIF JMS Provider > JMS connection factories** and do this:
 - Verify that EEZTECSenderQCF exists.
 - Select EEZTECSenderQCF and navigate to **Custom Properties**. Ensure that the value for the ServerLocation property contains the host name or

address of the TEC server. In addition, ensure that the value for the ServerPort property contains the number of the port on which the TEC server is listening.

- Check that the SystemAutomation.baroc file is located in the following directory:

- **Windows:** <EEZ_CONF_ROOT>

For example:

C:\Program Files\IBM\tsamp\eez\cfg

- **AIX and Linux:** <EEZ_CONF_ROOT>

For example:

/etc/opt/IBM/tsamp/eez/cfg

- Ensure that the SystemAutomation.baroc file is known to Tivoli Enterprise Console

For information on how to import, compile, load, and activate the BAROC file on the Tivoli Enterprise Console server, refer to the manual *IBM Tivoli Enterprise Console Rule Developer's Guide Version 3.9, SC32-1234* (Chapter 1, Rule development fundamentals - Rules - Rule bases - Rule base manipulation procedures using the rule builder).

- A test event must be available. To create a test event, perform the following steps:

1. Create the file eif_test_event.xml and copy the following test event into the file:

```
<?xml version="1.0" encoding="UTF-8"?>
<CommonBaseEvents xmlns="http://www.ibm.com/AC/commonbaseevent1_0_1"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.ibm.com/AC/commonbaseevent1_0_1
  commonbaseevent1_0_1.xsd">
  <CommonBaseEvent creationTime="2005-02-18T15:02:26.066Z"
    extensionName="SystemAutomation_Base"
    msg="EIFConverter testevent" severity="30"
    version="1.0.1">
    <extendedDataElements name="sa_domain_name" type="string">
      <values>EIFConverter Testdomainname</values>
    </extendedDataElements>
    <sourceComponentId application="N/A"
      component="IBM Tivoli System Automation for Multiplatform"
      componentIdType="N/A" location="N/A" locationType="Hostname"
      subComponent="" componentType="N/A" />
    <situation categoryName="AvailableSituation">
      <situationType xsi:type="AvailableSituation"
        reasoningScope="EXTERNAL"
        operationDisposition="N/A"
        processingDisposition="N/A"
        availabilityDisposition="N/A" />
    </situation>
  </CommonBaseEvent>
</CommonBaseEvents>
```

2. Save the file eif_test_event.xml to the following directory:

<was_root>/profiles/<profilename>/event/samples

To send the event to CEI, perform the following steps:

1. Open a command line.
2. Change the directory to <was_root>/bin.
3. Issue the following command:

```
wsadmin -f ../profiles/<profilename>/event/bin/emitevent.jacl
  -xml ../profiles/<profilename>/event/samples/eif_test_event.xml
```

A message appears, for example:

```

D:\prog\AppServer\event\bin>wsadmin
-f emitevent.jacl -xml ../samples/eif_test_event.xml
WASX7209I: Connected to process "server1"
on node BKDFN1CL using SOAP connector;
The type of process is: UnManagedProcess
Successfully submitted event(s) with global instance id(s):
CEE94F9747BC143EA171A356D084B411D9

```

To check whether the event was received by Tivoli Enterprise Console, perform the following steps:

1. Open a command line
2. Issue the following command:

```
wtdump r1
```

Sample output:

```

1~1969~1~1109064656(Sep 30 10:30:56 2005)
### EVENT ###
SystemAutomation_Base;hostname=N/A;severity=30;
date='Sep 28 15:02:26 2005';
sa_domain_name='EIFConverter Testdomainname';
source='IBM Tivoli System Automation for Multiplatform';
msg='EIFConverter testevent';END
### END EVENT ###
PROCESSED

```

Enabling Tivoli Enterprise Console event filtering

When you use the event console of the Tivoli Enterprise Console (TEC) product to display events, all end-to-end automation events are sent to the event console by default. To limit the scope of events that are forwarded to the event console, you can use the default Common Event Infrastructure (CEI) event filter that is provided for end-to-end automation management in order to achieve the following goals:

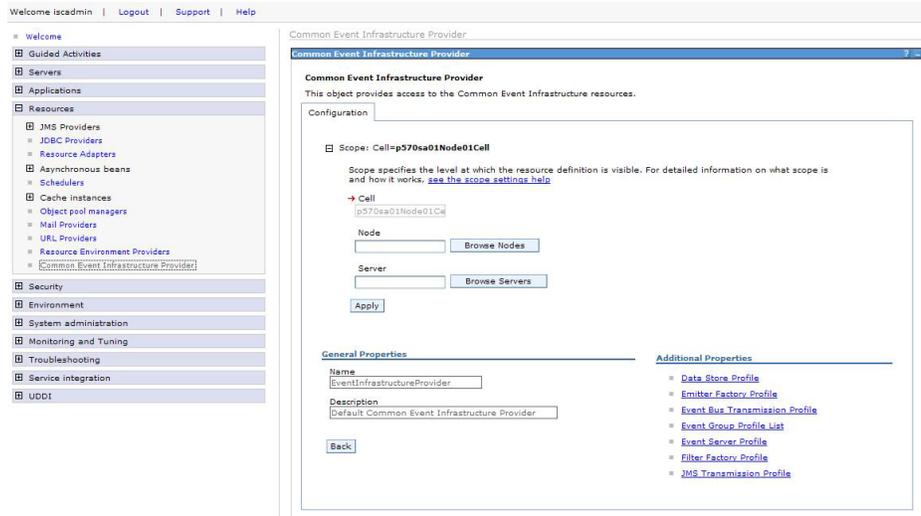
- All domain events and all operator request events are sent to the event console.
- Only resource events with severity level Critical are sent to the event console.

Resource events with severity level Warning or lower are dropped and not displayed on the event console.

The following sections describe how you activate and customize the default filter.

Activating the default CEI filter

Note: To perform the steps described below, you must select a cell scope in the administrative console panel shown in the screen capture below (**Resources** → **Common Event Infrastructure Provider**). To select a cell scope, empty the **Node** and **Server** fields.



To activate the default event filter, perform the following steps:

1. Open the WebSphere Application Server administrative console and navigate to **Resources** → **Common Event Infrastructure Provider** → **Filter Factory Profile**

2. Click **New** to create a new filter factory profile.

3. Enter the following values in the fields on the page:

Name Type EEZDefaultEventFilter in the field.

JNDI name

Type the following string in the field:
com/ibm/eez/aab/tec/EEZDefaultEventFilter

Description

Type the following description in the field:
EEZ Default Event Filter

Filter Configuration String

To specify that all domain events and operator request events are forwarded to the event console but resource events are to be forwarded only if they have the severity level Critical, type the following configuration string in the field:

```
CommonBaseEvent[(@severity > 30 and extendedDataElements
    [ @name = "sa_event_category"
      and @values = "ResourceEvent" ])
or extendedDataElements
    [ @name = "sa_event_category"
      and @values = "DomainEvent" ]]
```

The string specifies an XPath event selector that describes the events you want to use for filtering events. Events matching this event selector are sent to the event server; events that do not match will be discarded.

4. Click **OK** and save the configuration.

5. Navigate to the **Common Event Infrastructure Provider** —> **Emitter Factory Profile** page.
6. Go to the default emitter factory profile.
7. In the **Filter Factory JNDI Name** field, specify the JNDI name of the new filter factory profile you created:
com/ibm/eez/aab/tec/EEZDefaultEventFilter
8. Select the **Filtering enabled** check box.
9. Click **OK** and save the configuration.
10. Restart WebSphere Application Server to activate the filter.

Customizing the default event filter

You customize the default event filter by modifying the XPath event selector in the field **Filter Configuration String** on the Filter Factory Profile page (see previous section).

When modifying the XPath operators, remember the following rules:

- When used to compare XML dateTime values, the comparison operators perform logical comparisons that recognize time zone differences.
- Logical operators and function names must be specified using all lowercase letters (for example, *and* rather than *AND*).
- Operators must be separated with white space from the surrounding attribute names and values (`@severity > 30` rather than `@severity>30`).
- Parentheses can be used to change operator precedence.

The following examples are valid XPath event selectors.

Table 36. Valid XPath event selectors

| XPath event selector | Description |
|---|---|
| CommonBaseEvent[@extensionName = 'ApplicationStarted'] | All events with the extensionName attribute ApplicationStarted. |
| CommonBaseEvent[sourceComponentId/@location = "server1"] | All events containing a sourceComponentId element with the location attribute server1 |
| CommonBaseEvent[@severity] | All events with a severity attribute, regardless of its value. |
| CommonBaseEvent[@creationTime < '2003-12-10T12:00:00-05:00' and @severity > 30] | All events created before noon EST on 10 December 2003 and with severity greater than 30 (warning): |
| CommonBaseEvent[contains(@msg, 'disk full')] | All events with the phrase disk full occurring within the msg attribute. |
| CommonBaseEvent[(@severity = 30 or @severity = 50) and @priority = 100] | All events with a severity attribute equal to 30 or 50, and a priority attribute equal to 100. |

Part 5. Monitoring and managing automated resources

| | |
|---|-----|
| Chapter 31. Overview | 207 |
| Chapter 32. What you must know about the operations console | 209 |
| Configuring your Web browser | 209 |
| Logging on | 209 |
| Steps for accessing the operations console | 210 |
| Understanding the layout of the operations console | 212 |
| What you must know about the topology tree | 214 |
| Navigating the topology tree | 215 |
| Selecting an element in the topology tree | 216 |
| Limiting the scope of the topology tree | 216 |
| What is displayed in the topology column | 216 |
| What you can see in the Status column | 217 |
| What you can see in the Located here column | 217 |
| What you must know about the resources section | 217 |
| Section header | 218 |
| View and Search | 218 |
| Resource table views | 218 |
| Group hierarchy view | 219 |
| Search results view | 221 |
| What you must know about the information area | 222 |
| What you must know about the Smart refresh bar | 223 |
| What you must know about the main menu | 223 |
| Using links to quickly jump to a specific element | 224 |
| Chapter 33. Monitoring resources | 225 |
| State information provided on the operations console | 225 |
| Compound state and operational state | 225 |
| Compound state values | 225 |
| Compound state icons | 226 |
| State information provided for domains | 226 |
| Operational state descriptions provided on the General page | 227 |
| Domain state | 229 |
| Communication state | 230 |
| State information provided for nodes | 231 |
| State information provided for resources | 231 |
| Operational state descriptions provided on the General page | 232 |
| Observed state | 234 |
| Desired state | 235 |
| Monitoring tasks | 236 |
| Locating a resource | 236 |
| Switching between resource references and referenced resources | 236 |
| Identifying which first-level automation resource is referenced by a resource reference | 236 |
| Identifying the resource reference that references a first-level automation resource | 237 |
| Finding out to which groups a resource belongs | 237 |
| Displaying relationships | 237 |
| Viewing log files | 237 |
| Displaying operator instructions using the info link | 238 |
| Displaying owner contact information | 238 |
| Limiting the scope of the resource table | 238 |
| Displaying only resources that are in an error or warning state | 238 |
| Searching for resources | 239 |
| Defining a name filter | 239 |
| Applying an existing name filter | 241 |
| Administering name filters | 241 |
| Displaying only resources against which operator requests were submitted | 243 |
| Hiding domains | 243 |
| Using non-top-level resources as domain health indicators | 244 |
| Refreshing the operations console | 245 |
| Switching to a different end-to-end automation manager | 246 |
| Steps for connecting to a different end-to-end automation manager from the operations console | 246 |
| Managing your user credentials for first-level automation domains | 247 |
| Chapter 34. Managing resources | 249 |
| Working with policies | 249 |
| Activating a policy | 249 |
| Steps for checking the validity of a policy from the operations console | 249 |
| Steps for activating a policy | 250 |
| Deactivating a policy | 250 |
| Modifying a policy | 251 |
| Working with requests | 251 |
| Submitting start requests | 251 |
| Submitting stop requests | 252 |
| Displaying information about an operator request | 252 |
| Displaying request lists | 253 |
| Steps for viewing a request list and request details | 253 |
| Canceling requests | 253 |
| Steps for canceling requests | 254 |
| Resetting a resource from an unrecoverable error | 254 |
| Steps for resetting a resource | 254 |
| Including a node in automation and excluding a node from automation | 255 |
| Steps for excluding a node from automation | 255 |
| Steps for including a node in automation | 255 |
| Working with choice groups | 256 |
| Steps for starting the preferred member of a choice group | 256 |
| Steps for starting a different member of a choice group | 257 |

Chapter 31. Overview

This part of the guide is intended for operators. It describes the operations console of the end-to-end management component and how it can be used for monitoring and managing resources.

This is the information that is provided in the individual chapters of this guide:

- Chapter 32, “What you must know about the operations console,” on page 209 gives you an overview of the information provided on the operations console and the actions you can perform.
- Chapter 33, “Monitoring resources,” on page 225 describes how you can use the operations console to monitor resources and to analyze and resolve the problems that may occur.
- Chapter 34, “Managing resources,” on page 249 describes how you start and stop resources from the operations console, include nodes in automation and exclude them from automation, explains the procedures for working with choice groups, and describes how you have to proceed when a resource that has encountered an unrecoverable error should be included in automation again.

Note: Most of the information presented in this part of the guide applies to all operations console modes.

Resource references, referenced resources, and choice groups are resources that are managed by the end-to-end automation manager. References to these types of resources are applicable only when you are running the operations console in end-to-end automation mode.

For more information about the different console modes, refer to Chapter 3, “Operations console modes,” on page 15.

Chapter 32. What you must know about the operations console

This chapter gives you an overview of the information provided on the operations console and the actions you can perform.

All operators are recommended to read through this section in order to understand the basic concepts of using the operations console.

Configuring your Web browser

To be able to display the operations console in your Web browser, the following settings are required:

- JavaScript must be enabled in all Web browsers.
- For Microsoft Internet Explorer, the following settings are required:
 - Set the security level to medium.
Do not set the security level to high. If high security is required, ensure that the entry **ActiveX controls and plugins - Initialize and Script ActiveX controls not marked as safe** on the Security settings page is set to **Enable**. Otherwise, the information displayed on the operations console is not updated automatically.
 - Set **Scripting - Active Scripting** to **Enable** on the Security settings page. Otherwise, navigating the operations console is not possible.

Logging on

The operations console is a browser-based graphical user interface that runs in Integrated Solutions Console and is displayed in a Web browser window.

To access the operations console, you have to perform the following steps:

1. Open Integrated Solutions Console in a Web browser window.
2. Log in to Integrated Solutions Console using your user ID and password.
3. Connect to the operations console.

The following section describes how to log in to Integrated Solutions Console and connect to the operations console.

Note: It is recommended that you do not use multiple browser windows on the same client system simultaneously to connect to the same Integrated Solutions Console, because browser types other than Microsoft Internet Explorer will share a single HTTP session between multiple browser instances if these instances are running on the same system and connect to the same Integrated Solutions Console.

Working with multiple browser instances using the same HTTP session will cause unexpected results. The same situation occurs if you open multiple Microsoft Internet Explorer browser windows using **File** → **New Window** (or Ctrl+N) from an existing Integrated Solutions Console session, because in this case the new browser window and the one from which it was opened will also share the same session.

Steps for accessing the operations console

To access the operations console, perform the following steps:

1. Open a Web browser window and type the address of Integrated Solutions Console in the **Address** field.

The entry must have the following form:

```
http://<hostname>:<port>/ibm/console
```

where <hostname> is the name of the host on which Integrated Solutions Console is running and <port> is the specific port number of Integrated Solutions Console. The default port is 8421.

The log in panel of Integrated Solutions Console is displayed in the browser window:

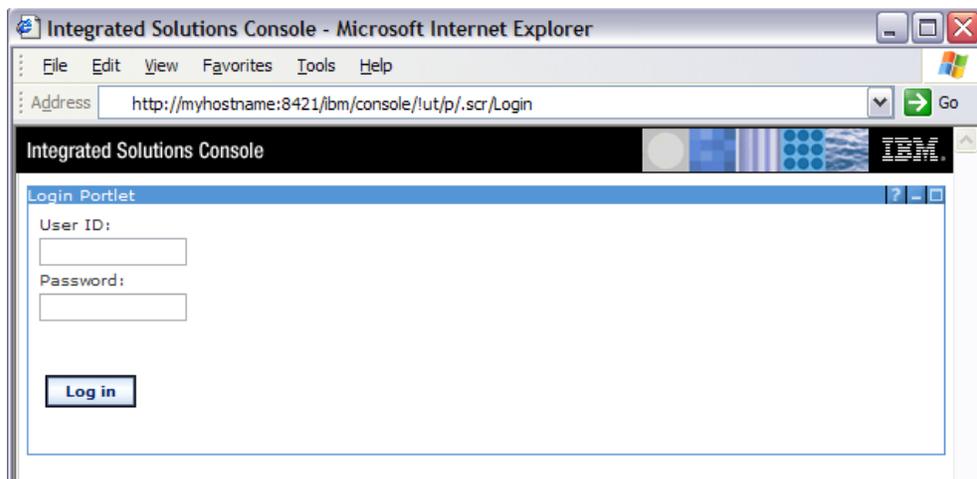


Figure 11. Log in panel of Integrated Solutions Console

2. On the log in panel, specify your user ID and password and click **Log in**.
The Welcome page of Integrated Solutions Console comes up:

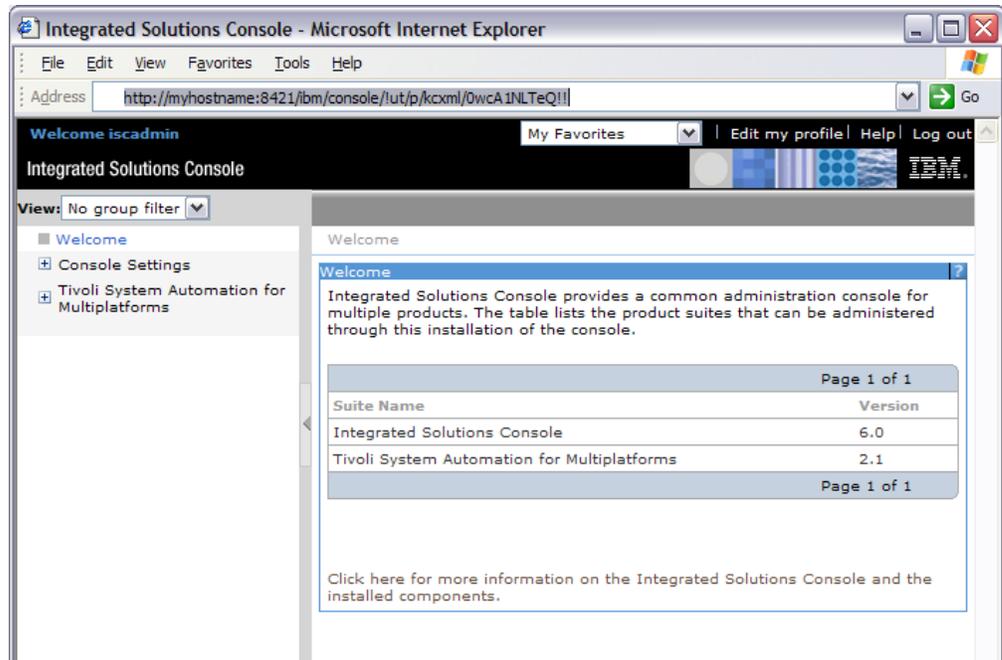


Figure 12. Welcome panel of Integrated Solutions Console

- In the navigation tree on the left, expand the folder **Tivoli System Automation for Multiplatforms**:

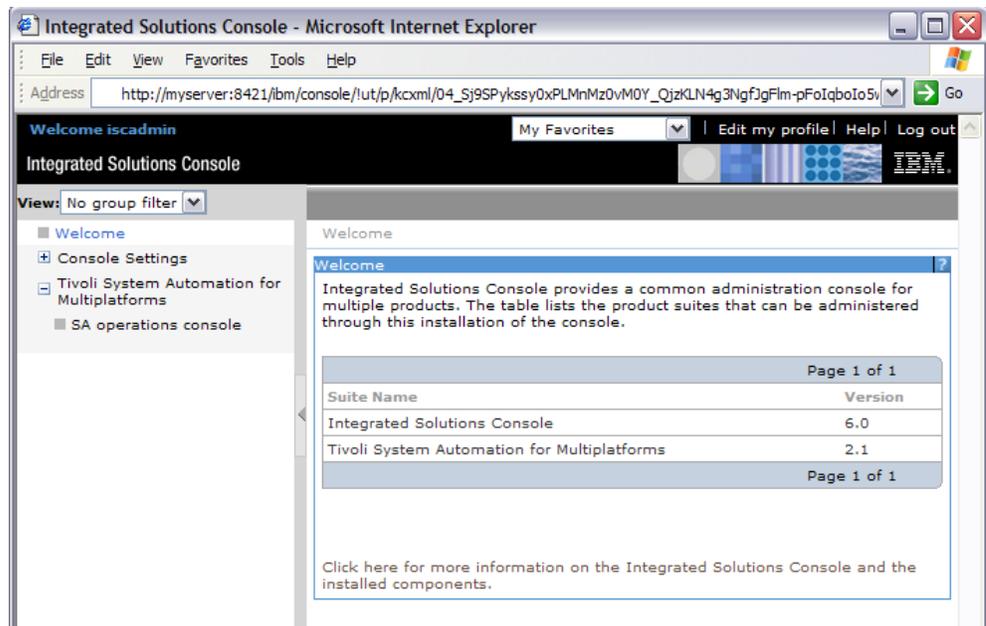


Figure 13. Operations console entry in the navigation tree

- Click **SA operations console**.
The Connect panel is displayed:

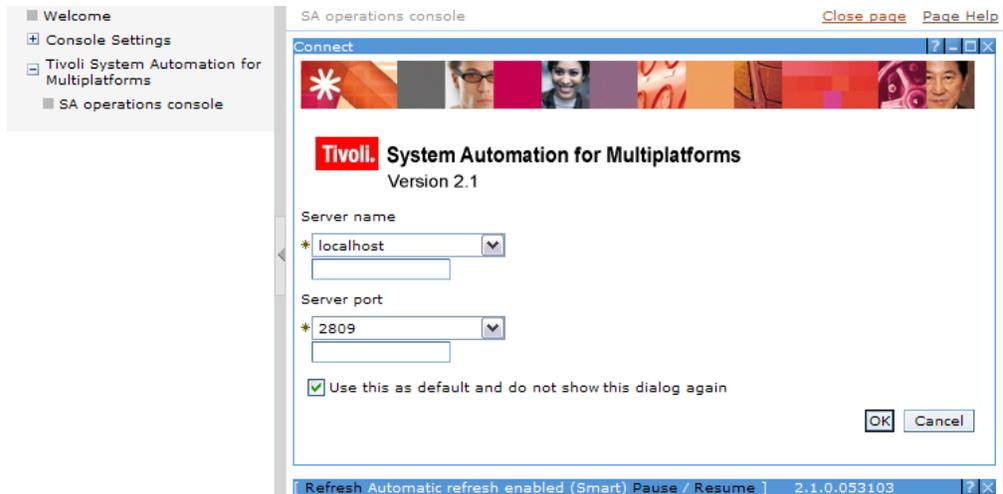


Figure 14. Connect panel

5. On the Connect panel, do the following:
 - a. In the field **Server name**, leave the default value (localhost) as is.
 - b. In the field **Server port**, the default value is displayed. Accept the default or select a port number from the drop-down list.
 If the appropriate port is not listed, select **Use entry from below** from the **Server port** drop-down list and type the port number of the server in the entry field below the **Server port** field.
 For a description of how to find the correct port number, see “How to determine the server port number for connecting to the operations console” on page 296.
 - c. If you want to set your entries as defaults, select the check box **Use this as default and do not show this dialog again**. In this case, this panel will not appear any more after you have logged in to Integrated Solutions Console.
 - d. Click **OK**. The main panel of the operations console is displayed.

Understanding the layout of the operations console

The main panel of the operations console of SA for Multiplatforms is divided into several areas:

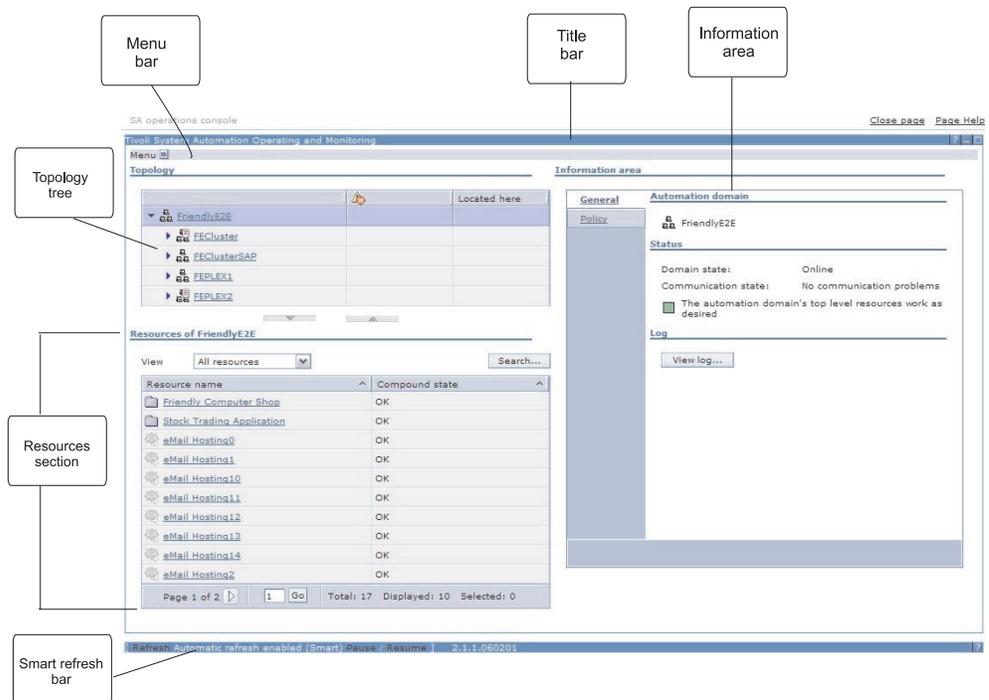


Figure 15. Main panel of the operations console

Title bar

Use the controls on the title bar to display the online help page for the panel you are displaying and for minimizing and maximizing the panel.

Menu bar

Use the entries in the Menu which is available on the menu bar to update the information displayed in the topology tree and the resource table, to change your user preferences, and to display information about the version of the operations console you are using. For more information, refer to “What you must know about the main menu” on page 223.

Information bar

The information bar is not shown in the figure above. It is displayed below the menu bar when you have performed an action on an element in the operations console. It displays a message confirming that the request or command has been submitted for processing. The message on the information bar only confirms the initial action, it is not updated while the command or request is being processed. The results of the system actions that are performed due to the request or command are reflected on the operations console itself. There you can see, for example, that the status of a resource has changed.

The confirmation message is replaced with a new message whenever you perform an action against an element in the operations console. Clicking **Clear** on the information bar hides the information bar from view. It reappears with a new confirmation message when you perform an action on an element.

Topology tree

The topology tree shows the automation domains and the nodes that belong to the domains. The topology tree displays state-related information, allows you to select and work with domains and nodes, and

is used to control what is displayed in the resource table. For more information, refer to “What you must know about the topology tree.”

Resources section

Use the areas of the resources section to work with resources:

View and Search

The View and Search functions allow you to limit the scope of the resource table. **View** allows you to display only resources that are in an error or warning state. **Search** allows you to display only resources whose name contains a filter phrase you specify, and lets you limit the scope of the resource table to resources against which operator requests have been submitted.

Resource table

Displays a list of resources and their states. You use it to select and work with resources. For more information about the resources section, refer to “What you must know about the resources section” on page 217.

The resource table has two views:

Search results view

When you use **Search** to see only a specific set of resources in the resource table, the search results are displayed in the search results view. For more information, refer to “Search results view” on page 221.

Group hierarchy view

The group hierarchy view is displayed when you are not displaying the results of a search. For more information, refer to “Group hierarchy view” on page 219.

Information area

Use the pages in the information area to obtain information about the element you have selected in the topology tree or resource table, and to perform actions against the element. For more information, refer to “What you must know about the information area” on page 222.

Smart refresh bar

On the Smart refresh bar, you can invoke an immediate smart refresh of the operations console, suspend and reactivate the smart refresh function, and you can see whether smart refresh is enabled or suspended. For more information, refer to “What you must know about the Smart refresh bar” on page 223.

What you must know about the topology tree

The following figure shows the topology tree and the resources section.

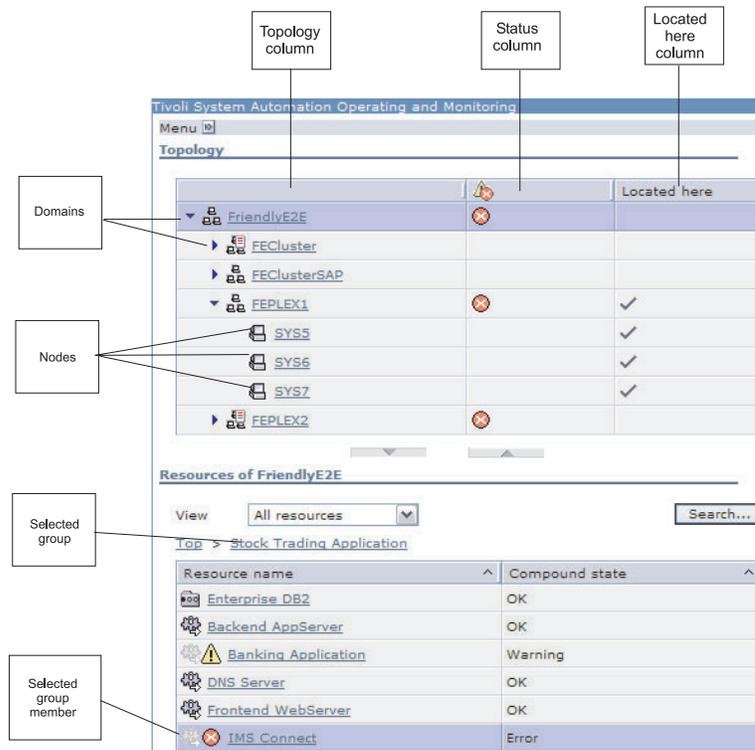


Figure 16. Topology tree and resources section

The topology tree is divided into three columns (see Figure 16):

- The **Topology** column shows the automation domains  and the nodes  that belong to a domain in a hierarchical view (see “What is displayed in the topology column” on page 216).
- The **Status** column shows the health status of the domain (see “What you can see in the Status column” on page 217).
- The **Located here** column is used to identify by which domain a resource is hosted and on which node or nodes it is located (see “What you can see in the Located here column” on page 217).

Figure 16 shows the following scenario:

- In the topology tree, the end-to-end automation domain (“FriendlyE2E”) is selected. The icon in the Status column indicates that at least one resource that is hosted by “FriendlyE2E” is in an error state.
- The resource table, in the resources section, shows the resources of the resource group “Stock Trading Application”.
- In the resource table, the resource reference “IMS Connect” is selected. The check marks in the **Located here** column of the topology tree indicate that the resources that are referenced by the resource reference “IMS Connect” are hosted by the first-level automation domain “FEPLEX1” and show on which nodes they are located.

Navigating the topology tree

You click the twistie in front of a domain icon to expand or collapse the nodes belonging to the domain.

Selecting an element in the topology tree

To select an element in the topology tree, you click the name of the element.

When you select a domain or node, you influence what is displayed in the resource table and in the information area:

- The resource table shows the top-level resources of the domain or node.
- The pages in the information area show information about the element that is selected in the topology tree. Depending on which type of element you have selected, buttons are enabled on the pages that let you perform actions against the element.

Limiting the scope of the topology tree

By default, all automation domains are displayed in the topology tree. When you are not interested in seeing all automation domains or if you are not authorized to access particular domains, you can hide domains from view (for more information, refer to “Hiding domains” on page 243).

What is displayed in the topology column

In the topology column you see the automation domains and the nodes that are managed by each first-level automation domain. When an end-to-end automation policy is active, the first-level automation domains whose resources are referenced in the policy appear below the end-to-end automation domain icon.

The following icons are used to identify the elements in the topology tree:

Table 37. Icons used for the elements of the topology tree

| Icon | Description |
|---|--|
|  | An automation domain. When the domain is not online or its state is unknown, the icon is grayed-out. |
|  | A node that belongs to a first-level automation domain. When a node is not online, the icon is grayed-out. |

The icons change their appearance if something happens that you need to be informed of. The following table provides some examples. For a complete overview of all icons that can be displayed on the operations console, see Appendix F, “Operations console icons,” on page 309.

Table 38. Some flavors of topology tree icons

| Icon | Description |
|---|---|
|  | At least one event was lost. Events inform you of a change to a resource, for example, a change in state of a first-level automation resource. This icon indicates that such an event could not be received, for example, because the network was down when the event was sent. This means that the information displayed on the operations console may not be correct for all resources. To resolve the problem, perform a Refresh all (Menu → Refresh all) to update the information on the operations console. |
|  | The first-level automation domain is online and commands and queries can be issued against the domain but no resource events are received. |

Table 38. Some flavors of topology tree icons (continued)

| Icon | Description |
|---|--|
|  | The first-level automation domain is online and resource events are still received from the domain but commands and queries cannot be issued against the domain. |
|  | There are new severe errors in the log file of the domain. |

What you can see in the Status column

The Status column is used to inform you of the health status of a domain. When the domain is healthy, the column is empty.

By default, a domain is considered healthy if none of the top-level resources that are hosted by the domain has encountered a problem that may require your attention. However, you can also define that a different set of resources is to be used to indicate whether a domain is healthy or not (refer to “Using non-top-level resources as domain health indicators” on page 244).

If a resource that is used as domain health indicator has encountered a problem, one of the following icons appears in the Status column:

Table 39. Icons in the Status column of the topology tree

| Icon | The icon indicates ... |
|---|--|
|  | A warning has been issued. The problem may still be resolved automatically, but the element should be monitored carefully. |
|  | The red error icon indicates that an error has occurred. To resolve the error, operator intervention is required. |
|  | The black error icon indicates that an unrecoverable error has occurred. To resolve the problem, urgent operator intervention is required. |

As the topology tree informs you of problems in a domain or on a node, you can use it as an entry point for monitoring resources.

What you can see in the Located here column

You use the **Located here** column to find out which domain hosts a resource or the members of a group and on which nodes the resources are located.

To determine the location of a resource, select the resource in the resource table. When you have made your selection, check marks in the **Located here** column indicate the hosting domain. Additionally, if you have expanded the domain, in which case the node hierarchy is displayed, check marks identify the node or nodes on which the resource is located (see Figure 16 on page 215).

What you must know about the resources section

The following figure shows the layout of the resources section.

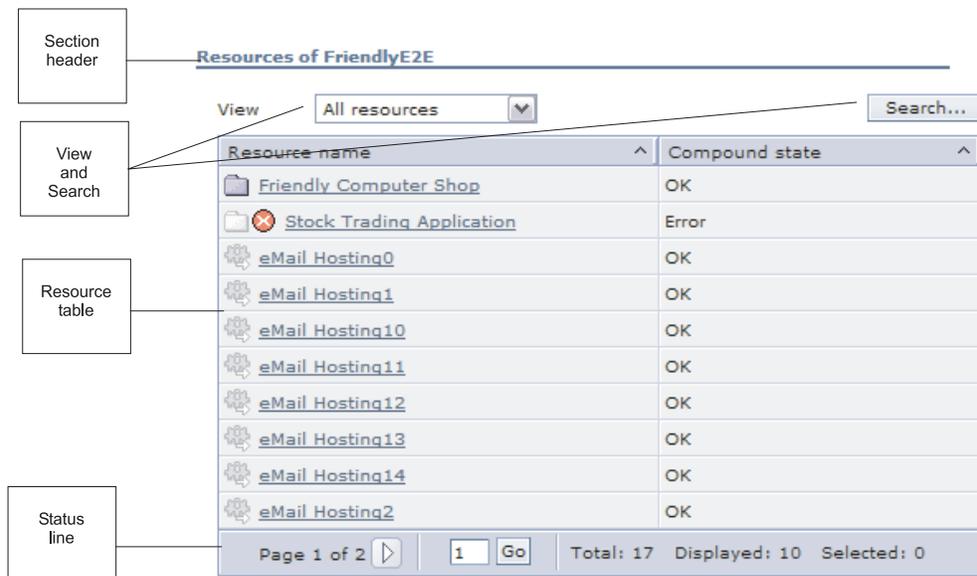


Figure 17. Layout of the resources section

The resources section has the following areas:

Section header

The section header displays the name of the domain or node that is currently selected in the topology tree.

View and Search

The View and Search functions allow you to limit the scope of the resource table: The View and Search functions allow you to limit the scope of the resource table:

View Select the **Errors and warnings** item from the View drop-down list to display only resources that are in an error or warning state. The view is always applied to the list of resources which is currently displayed in the resource table.

Search

Provides the following options:

- You can search for resources by name to display only resources whose name contains a filter phrase you specify.
- You can search for resources against which operator requests were submitted.
- You can combine both search methods.

Resource table views

The resource table has two views, which are described in the sections below. In both views, you can perform the following basic actions:

Select a resource

To select a resource, you click its name in the Resource name column.

Control the sort order of the resource table

You can sort the resource table on any column by clicking the sort arrow in the column header.

A solid sort arrow in a column header indicates that the table is currently sorted on the column. The direction in which the solid sort arrow is pointing indicates the current sort order (ascending or descending). By clicking on the solid sort arrow, you can toggle between ascending and descending sort order.

When you position the cursor over a sort arrow, a hover help text appears showing the current sort status of the column and the sort order that will result when the sort arrow is clicked.

Page through the resource table

The resource table may extend over multiple pages. To page through the table or to go to a specific resource you use the controls that are available in the status line below the table.

Group hierarchy view

The group hierarchy view is displayed when you are not displaying the results of a search. In the following figure, the top-level resources of the automation domain "Friendly E2E", which is selected in the topology tree, are displayed in the resource table.

Resources of FriendlyE2E

View

| Resource name ^ | Compound state ^ |
|---|------------------|
| Friendly Computer Shop | OK |
| Stock Trading Application | Error |
| eMail Hosting0 | OK |
| eMail Hosting1 | OK |
| eMail Hosting10 | OK |
| eMail Hosting11 | OK |
| eMail Hosting12 | OK |
| eMail Hosting13 | OK |
| eMail Hosting14 | OK |
| eMail Hosting2 | OK |

Page 1 of 2 Total: 17 Displayed: 10 Selected: 0

When you select a group in the resource table, the members of the group are displayed in the resource table. In the area above the table, a bread crumb trail appears. On the trail, the name of the group whose members are listed in the resource table is highlighted, indicating that the group is selected.



The bread crumb trail is useful for navigation and orientation:

- When you drill down into the group hierarchy, an entry is added to the trail for each group you select.
- The last entry on the trail identifies the group whose members are currently displayed in the resource table. When the group name is highlighted, the group is selected and the group details are displayed in the information area.
- When you click **Top** on the bread crumb trail, the top-level resources of the automation domain or node that is selected in the topology tree are again displayed in the resource table and the bread crumb trail disappears.
- When the bread crumb trail starts to get deeper than three levels, an ellipsis symbol (...) replaces all but the last two entries on the trail.

The ellipsis symbol cannot be clicked. To navigate upward through the group hierarchy, click an available group name on the trail until the group you want to view appears again, and select the group name on the trail to display the group members in the resource table.

Resource name column: The Resource name column lists the resources of the selected element, which is either an automation domain, a node, or a group.

- To sort the resources alphabetically by name, click the sort arrow in the column header.
- The resource icon to the left of the resource name indicates both the resource type and its online status: when the resource is online, its icon is active, when the resource is offline, the icon is grayed out.
- When a resource is in a warning or error state, the resource icon is highlighted with a warning or error icon.
- An operator icon  indicates that an operator request was submitted against the resource. The color of the operator icon changes while the request is being processed, yellow indicates that the request has been submitted, green indicates that the request was completed successfully.

The following table lists the resource icons that appear in the resource column.

| Icon | Description |
|---|---|
|  | A resource that is hosted by a first-level automation domain for which no resource reference is specified in the end-to-end automation policy |
|  | An end-to-end automation resource reference that references a first-level automation resource |
|  | A first-level automation resource that is referenced by a resource reference |

| Icon | Description |
|---|--|
|  | A resource group |
|  | A choice group or a first-level automation domain move group |

The following table lists the warning and error icons that appear in resource name column when a resource is in an error or warning state.

| Icon | Description |
|---|--|
|  | The yellow warning icon indicates that the resource is in warning state. |
|  | The red error icon indicates that the resource is in an error state. |
|  | The black error icon indicates that the resource has encountered an unrecoverable error. |

For the complete list of resource icons that appear in the operations console, refer to “Resource icons” on page 310.

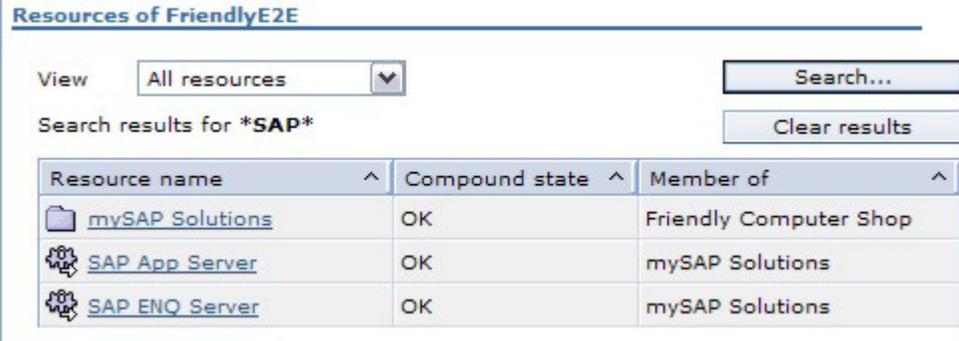
Compound state column: The column shows the compound state of the resource. By sorting on this column, you can group the resources by state.

The compound state can have one of the following values:

| State | Description |
|---------|--|
| OK | The resource is working as desired. |
| Warning | The resource is in warning state. |
| Error | The resource is in an error state. |
| Fatal | The resource has encountered an unrecoverable error. |

Search results view

When you use **Search** to see only a specific set of resources in the resource table, the search results are displayed in the search results view. In the area above the resource table, the search criteria that were used for the search are displayed. In this view, the resource table has the following layout:



| Resource name | Compound state | Member of |
|---|----------------|------------------------|
|  mySAP Solutions | OK | Friendly Computer Shop |
|  SAP App Server | OK | mySAP Solutions |
|  SAP ENQ Server | OK | mySAP Solutions |

To limit the scope of resources that are currently displayed in the resource table to those that are in an error or warning state, you can additionally apply the **Errors and warnings** view that is provided in the **View** field.

Resource table columns: In the search results view, the resource table has three columns:

Resource name column

In the column, the resources that match the search criteria are listed.

- To sort the resources alphabetically by name, click the sort arrow in the column header.
- If a resource is in a warning or error state, the resource icon is highlighted with a warning or error icon.
- If an operator request was submitted against the resource, an operator icon is displayed.
- Clicking a resource selects the resource and its details are displayed in the information area.

Note: When you select a group in the search results view, the group details will be displayed in the information area, but the resource table will not switch to the group hierarchy view to display the group members.

To display the group's members in the group hierarchy view, you must select the group and click **Clear results** (see “Clearing the search results”).

Compound state column

The column shows the compound state of the resource. By sorting on this column, you can group the resources by state.

Member of column

If a resource is a member of a group, the name of the group is displayed in this column. When you sort the resource table on this column, the resources that are members of the same group are listed next to each other.

Clearing the search results: When you click **Clear results**, the resource table switches back to the group hierarchy view. Which resources are then displayed in the group hierarchy view, depends on your selection in the search results view:

- No resource was selected: the top-level resources of the automation domain or node that is selected in the topology tree are displayed.
- A resource group was selected: The group members are displayed. On the bread crumb trail, the name of the group is highlighted, the group details are displayed in the information area.
- A resource that is a member of a group was selected: The group members are displayed, the group name is displayed on the bread crumb trail but is not highlighted, the name of the selected resource is highlighted in the resource list.

What you must know about the information area

In the information area, you find detailed information about the element that is currently selected in the topology tree or in the resource table.

On the pages in the information area, controls are available that let you perform actions on the selected element. Which pages are displayed and what they contain depends on the type of element that is currently selected:

| When you select ... | ...these pages are available |
|---|--|
| the end-to-end automation domain in the topology tree | <ul style="list-style-type: none"> • General • Policy |
| a first-level automation domain in the topology tree | <ul style="list-style-type: none"> • General • Policy • Additional Info |
| a resource or group in the resource table | <ul style="list-style-type: none"> • General • Relationships (available only if the resource has relationships) • Additional Info (available only if additional information exists) |

For detailed information about the pages in the information area, refer to the operations console online help. For detailed information about the internal states that are displayed on the Additional Info page for an end-to-end automation resource, refer to Appendix D, “Additional state-related information about end-to-end automation resources,” on page 287.

What you must know about the Smart refresh bar

The smart refresh function of the operations console checks at short intervals whether new information is available for any of the displayed elements. If new information is available, for example, when the state of a resource has changed, the operations console is updated accordingly.

On the smart refresh bar, which appears at the bottom of the operations console, you can force an immediate smart refresh, suspend the smart refresh, and reactivate it again. For more information, see “Refreshing the operations console” on page 245.

What you must know about the main menu

The main menu is available on the menu bar of the operations console.

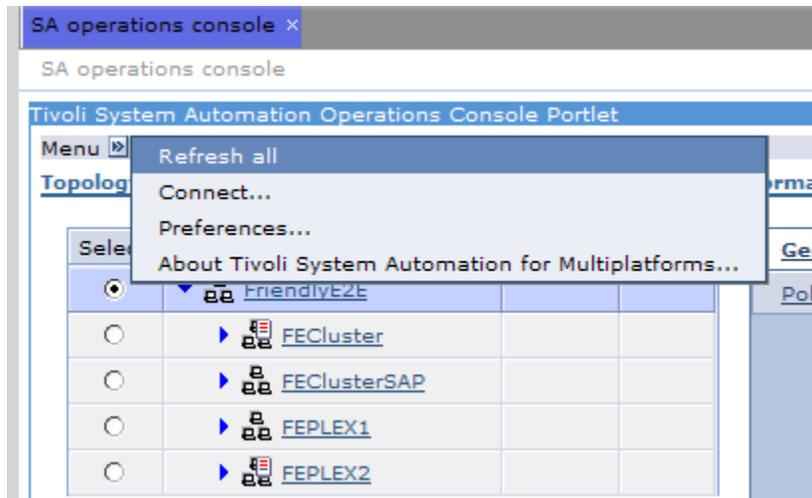


Figure 18. Main menu

You use the entries in the menu to perform these actions:

Refresh all

Retrieves the available information for all elements that are displayed on the operations console from the automation managers. The information on the operations console is updated. You will rarely need to use this function. Typically, the smart refresh function guarantees that the information on the operations console is up-to-date.

Connect

Lets you switch to a different end-to-end automation manager.

Preferences

Displays the Preferences panel. You use the pages that are available on the Preferences panel to customize your operations console:

- On the User IDs page, you can manage your user IDs for first-level automation domains (for more information, refer to “Managing your user credentials for first-level automation domains” on page 247).
- On the Name filters page, you can define and manage the name filters you use for limiting the scope of the resource table (for more information, refer to “Searching for resources” on page 239).

Additionally, you can define which resources are to be used as domain health indicators (for more information, refer to “Using non-top-level resources as domain health indicators” on page 244).

- On the Visible automation domains page you can limit the scope of the topology tree by defining which domains should be hidden from view (for more information, refer to “Hiding domains” on page 243).

About Displays information about the version of the component you are using.

Using links to quickly jump to a specific element

In many places on the operations console, the names of elements are implemented as links that allow you to quickly jump to the element. When you click such a link, the current contents of the operations console change to display the information for the selected element.

Chapter 33. Monitoring resources

This section describes how you can use the operations console of SA for Multiplatforms to monitor the states of resources, and to identify and analyze problems.

State information provided on the operations console

Observing the states of resources is the most important aspect of monitoring. The topics in this section describe the state-related information that is provided on the operations console for domains, nodes, and resources.

Compound state and operational state

The compound state plays an important role in monitoring and problem analysis. It informs you of the health status of a domain, a group, or a resource.

On the operations console, information about the compound state is provided for domains, groups, and individual resources. The compound state is complemented by the operational state, which provides additional information about the compound state.

The compound state is displayed as an icon that appears in several places on the operations console:

- In the topology tree, a warning or error icon appears in the Status column when a resource that you are using as domain health indicator for the domain has encountered a problem. When no compound state icon is displayed in the topology tree, this indicates that the domain is healthy.
- In the resource table, the resource icon in the resource name column is highlighted with a warning or error icon when a resource has encountered a problem.
- The compound state icon also appears on the General page of a domain, group, or resource. To the right of the compound state icon on the General page, the operational state description is displayed providing additional information about the compound state.

The fact that the health status of a resource is indicated for the resource itself, for the group it belongs to, and for the domain which hosts it, allows you to monitor resources simply by observing the compound state of the domains in the topology tree. When no problem is indicated there, this usually means that all resources are working as desired.

Compound state values

The compound state has the following possible values:

OK The resource works as desired.

Warning

A problem has occurred. Operator intervention is not yet required, but careful monitoring is recommended.

Error A severe problem has occurred. Operator intervention is required.

Fatal An unrecoverable error has occurred. Operator intervention is required.

Compound state icons

The following table lists the compound state icons that appear on the operations console when a problem has occurred.

Table 40. Compound state icons

| Icon | Example | Description |
|---|---|---|
|  |  | <p>Compound state: Warning</p> <p>The yellow icon indicates that the resource may require your attention. However, the problem may still be resolved by automation management. Check the operational state description on the General page for more information on the problem.</p> <p>When the resource for which the warning is indicated is used as domain health indicator, the warning icon is also displayed in the status column of the topology for the domain that hosts the resource.</p> |
|  |  | <p>Compound state: Error</p> <p>The red icon indicates that the resource may require operator intervention. Check the operational state description on the General page for more information on the problem.</p> |
|  |  | <p>Compound state: Fatal</p> <p>The black icon indicates that an unrecoverable error has occurred. Operator intervention is required to resolve the problem. Check the operational state description on the General page for more information.</p> <p>Note: When an unrecoverable error has occurred and the problem has been resolved, the resource will not be automated again automatically. To include the resource in automation again, the function Reset from unrecoverable error must be used (see "Resetting a resource from an unrecoverable error" on page 254).</p> |

State information provided for domains

This section describes the states that are displayed on the operations console for a domain:

- Operational state
- Domain state
- Communication state

In the topology tree, icons inform you of the compound state, the domain state, and the communication state of a domain. Additional information about these states is available in the status section on the General page for the domain that is selected in the topology tree.

The following figure shows the status section on the General page for a domain:

Information area

General **Automation domain**

Policy

Name: FriendlyE2E

Status

Domain state: Online

Communication state: No communication problems

The automation domain's top level resources work as desired

Log

[View log...](#)

Operational state descriptions provided on the General page

The following table lists some of the operational state descriptions that are displayed on the General page when a domain is selected in the topology tree, and provides some basic information on how you can proceed when a problem has occurred.

The operational state description is displayed to the right of the compound state icon on the General page. For general information about the compound state, see “Compound state and operational state” on page 225.

Table 41. Operational state descriptions provided on the General page for a domain

| Description on the General page | Troubleshooting |
|---|---|
| The domain's top-level resources work as desired. | None. |
| The domain contains top-level resources with warnings. At least one resource matching the name filter <current domain filter> has a warning. | What it means: At least one of the resources you are using as domain health indicators has encountered a problem. What you can do: Find out which resource is affected and monitor it carefully. Usually, the resource will recover automatically. |

Table 41. Operational state descriptions provided on the General page for a domain (continued)

| Description on the General page | Troubleshooting |
|---|---|
| <p>The domain contains top-level resources with errors.</p> <p>At least one resource matching the name filter <current domain filter> has an error.</p> | <p>What it means: At least one of the resources you are using as domain health indicators has encountered a serious problem. Operator intervention may be required.</p> <p>What you can do:</p> <p>Find out which resource is affected and analyze the problem, for example:</p> <ul style="list-style-type: none"> • View the domain log file and check for error messages. • Drill down to the affected first-level automation resource and check its compound state. • Check the relationships of the affected resource. • View the requests and votes that have been issued against the resource. • Consult the information pages for the resource. The information pages are available in the information area when you select the resource in the resource table. • Contact the owner of the application. |

Table 41. Operational state descriptions provided on the General page for a domain (continued)

| Description on the General page | Troubleshooting |
|---|--|
| <p>The domain contains top-level resources with unrecoverable errors.</p> <p>At least one resource matching the name filter <current domain filter> has an unrecoverable error.</p> | <p>What it means:</p> <p>At least one of the resources you are using as domain health indicators has encountered an unrecoverable problem.</p> <p>What you can do:</p> <p>Find out which resource is affected and analyze the problem, for example:</p> <ul style="list-style-type: none"> • View the domain log file and check for error messages. • Identify the location of the resource and check the system and application logs for error messages. • Drill down to the affected first-level automation resource and check its compound state. • Consult the information pages for the resource. The information pages are available in the information area when you select the resource in the resource table. • Contact the owner of the application. <p>If the message is displayed for the end-to-end automation domain, ensure that the automation engine's user credentials for the first-level automation domains are specified correctly in the configuration dialog.</p> <p>After resolving the problem, you must use the Reset function to include the resource in automation again.</p> |

Domain state

The domain state indicates whether the domain is currently online, offline, or whether the state is unknown. The domain state value is displayed on the General page. Possible values are:

- Online
- Offline
- Unknown

In the topology tree, the appearance of the domain icon shows the state of the domain:

Table 42. Domain state icons

| Icon | State | Description |
|---|--------------------|--|
|  | Online | The active icon indicates that the domain is online. |
|  | Offline or Unknown | The grayed out icon indicates that the domain is offline or that its state is unknown. |

Communication state

The communication state provides you with the following information:

- Adapter-related information: whether the adapter to the first-level automation domain is operational
- Connectivity-related information:
 - whether events can be received from the automation adapter
 - whether requests or queries can be submitted to the automation adapter
- When events were lost

On the operations console, the communication state is indicated in two places:

- In the topology tree, the appearance of the domain icon changes when a problem has occurred.
- On the General page of a domain, a description of the communication state is provided.

The following table gives you an overview of how a problem is indicated in the topology tree and on the General page. The complete list of domain icons that may appear in the topology tree is available in Appendix F, “Operations console icons,” on page 309.

Table 43. Communication state

| Communication state icons and state descriptions | What it means |
|--|--|
|  <p>No communication problems.</p> | No action is required. |
|  <p>Commands and queries can currently be issued against this domain, but at least one resource event was lost.</p> | <p>The state information provided for the domain’s resources may be outdated. Perform a Refresh all to update the information.</p> <p>If the domain stays in this state for a longer period of time, the configuration properties of the domain may need to be changed. Inform the system administrator of the domain.</p> |
|  <p>No commands or queries can currently be issued against this domain, but resource events are still received.</p> | <p>One of the following problems may have occurred:</p> <ul style="list-style-type: none"> • The adapter has failed. Try to start the adapter. • The network is down. Call the network administrator. • A firewall has been activated which commands or queries cannot pass. Call the responsible administrator. • For the end-to-end automation domain: Contact the system administrator. The administrator should check whether the automation engine is still active. |

Table 43. Communication state (continued)

| Communication state icons and state descriptions | What it means |
|---|--|
|  <p>Commands and queries can currently be issued against this domain, but no resource events are received.</p> | <p>The state information provided for the resources that are hosted by the domain may be outdated.</p> <p>Perform a Refresh all to update the information.</p> <p>The configuration properties of the domain may need to be changed. Inform the system administrator of the domain.</p> |
|  <p>None of the communication paths to this domain are currently working.</p> | <p>No queries can be submitted, no events can be received. The resource state information may be outdated. No refresh is possible.</p> <p>Check if the adapter has failed.</p> |
|  <p>No commands or queries can currently be issued against this domain and at least one resource event was lost.</p> | <p>The state information provided for the resources of the domain may be outdated.</p> <p>View the log files manually for further information.</p> |
| <p>The automation adapter is currently not running.</p> | <p>The adapter may have been stopped intentionally by an administrator.</p> |

State information provided for nodes

The observed state of a node indicates whether a node is currently

- online or offline
- included in automation or excluded from automation

The observed state of a node is visible in the topology tree and in the state section on the General page. The following table gives you an overview of how the observed state is displayed. The complete list of icons is available in Appendix F, “Operations console icons,” on page 309.

Table 44. Observed state of a node

| Icon | State | Description |
|---|---------|--|
|  | Online | The active icon indicates that the node is online. |
|  | Offline | The grayed out icon indicates that the node is offline. |
|  | Online | The node is online and has been excluded from automation. |
|  | Offline | The node is offline and has been excluded from automation. |

State information provided for resources

On the operations console, you find the following state-related information about a resource or group:

Compound state

The compound state icon indicates whether a resource works as desired or has encountered an error.

Operational state

The operational state provides additional information about the compound state. The operational state description is displayed to the right of the compound state icon on the General page.

Observed state

The observed state represents the current state of the resource as reported by the automation manager of the domain by which it is hosted.

Desired state

The desired state reflects the automation goal of the resource.

Information about these states is available on the General page in the resource status section. The observed state and the compound state are also visible in the resource table.

The following figure shows the resource status section on the General page for a resource reference. When you select a different type of resource in the resource table, the section header on the General page changes accordingly, but the appearance of the section itself and the way in which the state information is provided are identical for all types of resources.

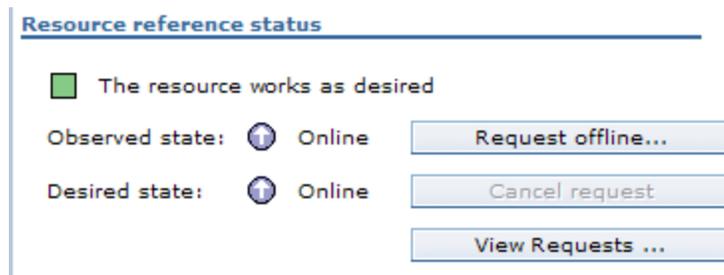


Figure 19. State information on the General page

The following sections describe the states and their possible values, and explain how and where the states are displayed.

Operational state descriptions provided on the General page

The possible values of the compound state and how the compound state of a resource is indicated in the resource table and in the topology tree is described in “Compound state and operational state” on page 225.

The following table lists some of the operational state descriptions that are displayed on the General page when a resource is selected in the resource table. Most of the descriptions that appear there are self-explanatory. In some cases, the table provides additional information about what may have caused a problem.

Table 45. Operational state descriptions on the General page for a resource

| Operational state description on the General page | Possible causes and actions |
|---|-----------------------------|
| The resource works as desired. | |

Table 45. Operational state descriptions on the General page for a resource (continued)

| Operational state description on the General page | Possible causes and actions |
|--|---|
| The resource works as desired but is dormant. | No action required. |
| Warning: The resource is performing poorly. | |
| Warning: The resource has stopped but not completed its job. | |
| Warning: No contact to resource. | This message is displayed for end-to-end automation resources only. Usually, the message is transient and requires no action. It is displayed after the automation engine is started, indicating that the initial event for the resource has not yet been received from the first-level automation domain. The message usually disappears as soon as the initial event has been received. |
| Warning: The communication has been interrupted. | |
| Warning: The resource has been forced down. | This message is displayed for first-level automation resources only. It usually means that the resource was forced down by a first-level operator. |
| Error: The hosting domain is gone. | This message is displayed for resource references only. It indicates that the first-level automation domain which hosts the referenced resource is not available. |
| Error: The hosting node is gone. | This message is displayed for first-level automation resources only. It indicates that the node on which the resource is located is offline. |
| Error: The resource has been excluded from automation. | |
| Error: The resource reference references a resource that does not exist. | This message indicates that the policy contains an incorrect reference or that the adapter cannot send the names of the resources of the domain. |
| Error: The start processing did not finish successfully. | |
| Error: The stop processing did not finish successfully. | |
| Error: The referenced resource is in an error state. | This message indicates that the end-to-end automation manager cannot bring the resource reference into the desired state because the referenced resource has encountered an error. To correct the error, the problem that was encountered by the referenced resource must be resolved. |
| The resource has an unrecoverable problem. | |
| The resource has an unrecoverable problem: The start processing did not finish successfully. | |
| The resource has an unrecoverable problem: The stop processing did not finish successfully. | |

Table 45. Operational state descriptions on the General page for a resource (continued)

| Operational state description on the General page | Possible causes and actions |
|---|--|
| The resource has an unrecoverable problem: Unable to contact the referenced resource. | This message indicates that the end-to-end automation manager cannot establish contact with the referenced resource. This problem occurs when the end-to-end automation manager caught some exceptions when it tried to access the referenced resource. To analyze the problem, look in the end-to-end automation domain log file for additional information about the exception. After resolving the problem you must use the Reset function to include the resource in automation again. |
| The following messages are displayed when a start or stop request has been submitted. | |
| Warning: Online request pending. | An operator has submitted a start request against the resource. |
| Warning: Offline request pending. | An operator has submitted a stop request against the resource. |
| Warning: Operation in progress. | A temporary state. The message is displayed while a resource is starting or stopping. |
| Error: The resource cannot be started because the online request did not win at this moment. | The start request did not win. However, the request stays in the request list and may be processed at a later time. You can check the request list of the resource to find out why the request did not win. |
| Error: The resource cannot be stopped because the offline request did not win at this moment. | The stop request did not win. However, the request stays in the request list and may be processed at a later time. You can check the request list of the resource to find out why the request did not win. |
| Error: The resource cannot be started because of unfulfilled dependencies. | The resource could not be started because a resource that had to be started first could not be started. Check the relationships of the resource to find out which target resource could not be started. |
| Error: The resource cannot be stopped because of unfulfilled dependencies. | The resource could not be stopped because a resource that had to be stopped first could not be stopped. Check the relationships of the resource to find out which target resource could not be stopped. |

Observed state

The observed state represents the current state of the resource as reported by the automation manager.

Possible values are:

Online

The resource is online.

Offline

The resource is offline.

Starting

The resource is starting.

Stopping

The resource is stopping.

Unknown

The automation manager has no information about the current state of the resource. When displayed for an end-to-end automation resource, this state indicates that the resource has not been contacted yet.

On the General page, the state value is provided in the resource state section (see Figure 19 on page 232). In the resource table, the resource icon indicates the observed state of the resource:

- When the icon is active, the resource is online or stopping.
- When the icon is grayed out, the resource is not online. This is the case when the resource is offline or starting, or when the current state of the resource is unknown.

Desired state

The desired state reflects the automation goal of a resource. The automation manager tries to keep the resource in this state. The default desired state is specified in the automation policy. At runtime, the desired state is influenced by operator actions (start and stop requests) and by a resource's relationships (StartAfter, StopAfter, and ForcedDownBy relationships). (For more information on automations goals and relationships, see Chapter 5, "Automation concepts," on page 27.)

Possible values are:

Online

The automation goal is set to online. The automation manager tries to keep the resource online.

Offline

The automation goal is set to offline. The automation manager tries to keep the resource offline.

Not changeable

This value is displayed for monitor resources, which can be monitored on the operations console but whose desired state cannot be changed through start or stop requests.

Monitoring tasks

The following sections describe tasks you will perform to obtain information about resources and for analyzing problems.

Locating a resource

To find out where a resource or the members of a group are located, select the resource or group in the resource table. One or more check marks appear in the **Located here** column of the topology tree. The check marks indicate by which automation domain the selected resource or the members of the selected group are hosted and on which nodes they are located (see Figure 16 on page 215).

Switching between resource references and referenced resources

In many places on the operations console, the names of elements are implemented as links that allow you to quickly jump to the element. Typically, when you click such a link, the current contents of the operations console change to display the information for the selected element. You can use the links, for example, to perform the following tasks:

Identify which first-level automation resource is referenced by a resource reference

This is helpful when you are monitoring the resources of the end-to-end automation domain and you see that a problem is indicated for a resource reference.

Identifying the resource reference that references a first-level automation resource

These tasks are described in the following sections.

Identifying which first-level automation resource is referenced by a resource reference

Perform the following steps:

1. Select the resource reference in the resource table.

This is what is displayed on the operations console:

- In the **Located here** column of the topology tree, a check mark indicates which first-level domain hosts the resource.
- In the information area, the information pages for the resource reference are displayed. The **Referenced resource** section on the General page shows the name of the referenced resource.

-
2. Click the name of the referenced resource in the **Referenced Resource** section.
-

Results:

This is what is displayed on the operations console:

- In the topology tree, the first-level automation domain that hosts the referenced resource is selected.
- The resources section header displays the name of the first-level automation domain.
- In the resource table, the referenced resource is selected.

- In the information area, the information pages for the referenced resource are displayed.

Identifying the resource reference that references a first-level automation resource

Perform the following steps:

1. Select the first-level automation resource in the resource table to display the information pages for the resource in the information area. In the **Used by** section on the General page, the name of the corresponding resource reference is displayed.

2. Click the name of the resource reference in the **Used by** section.

Results:

This is what is displayed on the operations console:

- In the topology tree, the end-to-end automation domain is selected.
- In the header of the resources section the name of the end-to-end automation domain is displayed.
- In the resource table, the resource reference is selected.
- In the information area, the information pages for the resource reference are displayed.

Finding out to which groups a resource belongs

To find out of which groups a resource is a member, select the resource in the resource table. The groups to which the resource belongs are listed in the **Used by** section on the General page in the information area.

Displaying relationships

You use the Relationships page in the information area to display the forward and backward relationships for a resource. For each resource that participates in a relationship, a hyperlink lets you jump to the resource.

Before you begin:

- The Relationships page is only available for resources for which relationships have been defined.
- For first-level automation resources, the Relationships page may contain first-level automation-specific relationships.

Perform the following steps to display the relationships of a resource:

1. Select the resource in the resource table.

2. In the information area, click the Relationships tab to open the Relationships page.
To jump to a resource, click the name of the resource in the relationship table.

Viewing log files

Much information about a domain, its nodes, and the resources that are hosted by the domain is written to the log file of the domain. You can display the domain log

file from the operations console. Checking a log file for messages always is an important step in problem analysis. Viewing a log file is especially important when the domain icon indicates that there are new severe errors in the log file ().

You can display a domain log file by performing the following steps:

1. Select the domain in the topology tree.

-
2. On the General page, click **View log**.

Result: The log file is displayed in the **Log viewer** panel.

For information about displaying the log file of the end-to-end automation domain when the file is not accessible from the operations console, for example, because the automation engine is not running, refer to “Viewing the XML log file of the automation engine” on page 294.

Displaying operator instructions using the info link

Instructions that have been specifically provided for a resource can be helpful when a problem occurs and you need additional information about the resource.

To display the operator instructions for a resource, perform the following steps:

1. Select the resource in the resource table and open the General page in the information area.

-
2. On the General page, click **Info link**.

Result: The operator instructions for the resource are displayed.

Displaying owner contact information

Information about the owner of a resource is available on the General page for a resource. To display the General page, select the resource in the resource table and click the **General** tab in the information area.

Limiting the scope of the resource table

This section describes how you use the View and Search functions to limit the scope of resources that are displayed in the resource table.

Displaying only resources that are in an error or warning state

The item **Errors and warnings** that is available in the **View** field allows you to list only resources in the resource table that are in an error or warning state.

To activate the **Errors and warnings** view, select the corresponding item in the **View** field. To deactivate it, select the item **All resources** from the **View** list.

The **Errors and warnings** view is always applied to the list of resources that is currently displayed in the resource table:

- The top-level resources of a domain or node are displayed in the resource table: When you activate the **Errors and warnings** view, the resource table lists all resources of the domain or node that are in an error or warning state.

- A group is selected in the resource table:
When you activate the **Errors and warnings** view, the resource table displays only the group members that are in an error or warning state.
- You are displaying the results of a search:
When you activate the **Errors and warnings** view, the resource table displays only the resources that match the search criteria and are in an error or warning state.

Searching for resources

You can search for resources using the following search arguments:

- You can search for resources by name to display only resources whose name contains a filter phrase you specify.
- You can search for resources against which operator requests were submitted.
- You can combine both search methods.

Defining a name filter

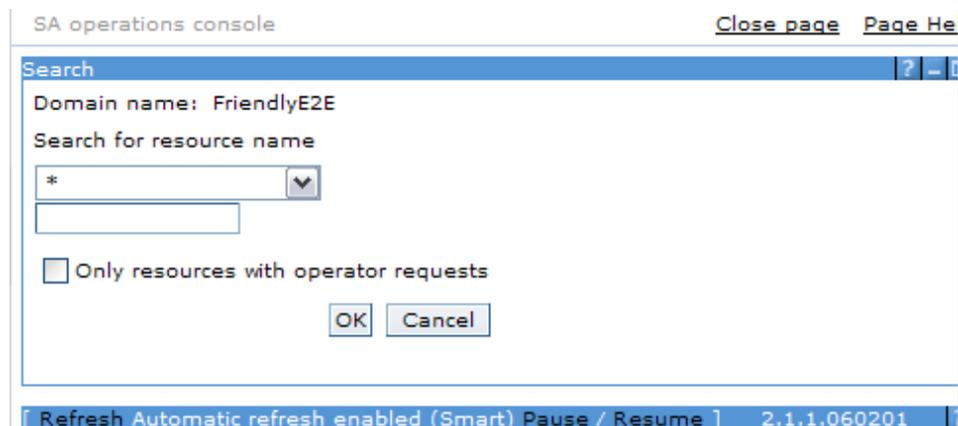
Before you begin:

You can define name filters on the following pages:

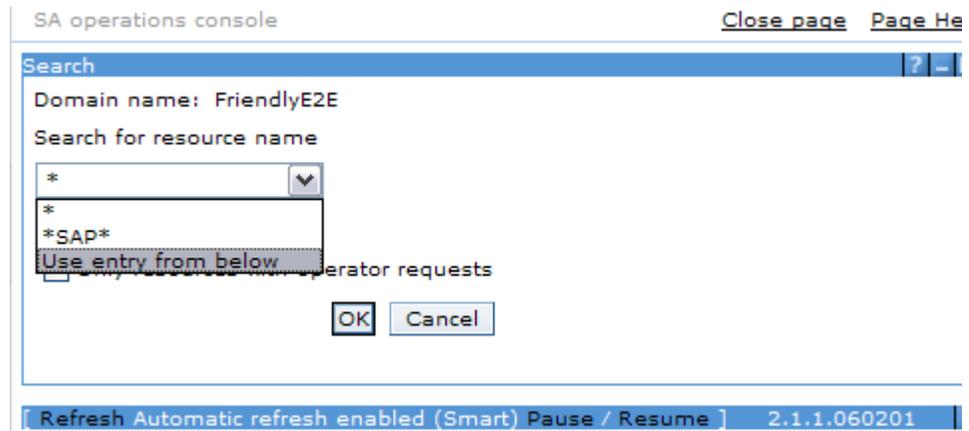
- On the Name filters page of the **Preferences** panel.
To open the **Preferences** panel, click **Menu** —> **Preferences**
- On the Search page, which appears when you click the **Search** button that is displayed above the resource table. This is described in the following procedure.

Perform the following steps to define a name filter:

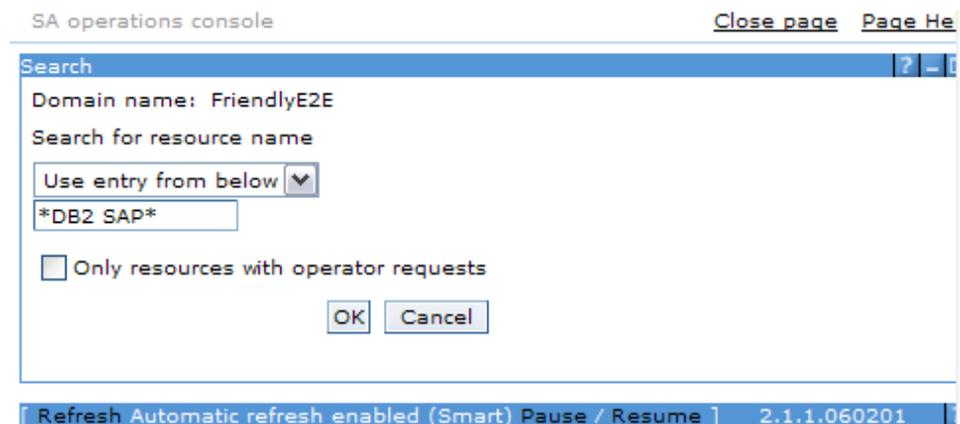
1. In the topology tree, select the domain or the node for which you want to define the filter.
-
2. Click **Search**. The Search page is displayed.



3. From the **Search for resource name** drop-down list, select **Use entry from below**.



4. In the field below, type the phrase by which the domain or node should be filtered. You have the following options:
 - To display only one specific resource in the resource table, type the exact resource name.
 - Use the asterisk * as wildcard character to display all resources whose names contain the filter phrase. The wildcard character can appear in any position and, if necessary, more than once (for example, *DB2*), and can stand for 0..n characters.
 - To display all resources whose names contain at least one of several search phrases, type both phrases and separate them by a blank; the wildcard character can be used in one or both phrases (for example, *DB2* SAP*).
 - For resource names that may contain blanks, type the complete filter phrase, including the blank, and enclose the filter phrase in single or double quotation marks, for example, "SAP *Server". This ensures that it will be recognized as a single phrase.



5. Click **OK** to submit the query.

Results:

- The search results view of the resource table is displayed. Depending on whether you selected a domain or a node in the topology tree, the resource table lists only the resources of the selected domain or node whose names match the filter criteria you specified.

Resources of FriendlyE2E

View

Search results for ***DB2 SAP***

| Resource name ^ | Comp... ^ | Member of ^ |
|--|-----------|---------------------------|
|  Enterprise DB2 | OK | Stock Trading Application |
|  DB2 | OK | Friendly Computer Shop |
|  SAP App Server | OK | mySAP Solutions |
|  SAP ENQ Server | OK | mySAP Solutions |

- The filter remains active until you deactivate it by clicking **Clear results**.
- The filter is added to your user preferences for the domain, which means that it is from now on available in the drop-down list on the Search page for the domain and all of its nodes, regardless of whether you defined the name filter for the domain or only for one of its nodes.

Applying an existing name filter

Name filters that you have defined for a domain or a node are available in the list of name filters on the Search page when you have selected the domain or one of its nodes in the topology tree. Because name filters are domain-specific, a name filter is only available in the filter list for the domain for which it was specified, and for the nodes that belong to that domain. If you want to use a name filter with the same filter criteria for any other domain and its nodes, you must define a new filter with the same filter criteria.

Perform the following steps to apply an existing filter:

1. Select the domain or the node to which you want to apply the filter.

2. Click **Search**. The Search page is displayed.

3. From the **Search for resource name** drop-down list, select the filter you want to apply.

4. Click **OK** to apply the filter.

Results:

- The search results view of the resource table is displayed. Depending on whether you selected a domain or a node in the topology tree, the table lists only the resources of the selected domain or node whose names match the filter criteria.
- The filter remains active until you deactivate it by clicking **Clear results**.

Administering name filters

On the Name filters page, on the Preferences panel, you can perform the following tasks:

- Define a new filter
- Edit a filter

- Delete filters

Perform the following steps to administer your name filters:

1. Open the Preferences panel (**Menu** → **Preferences**).

2. Open the Name filters page.

3. Select the domain whose filters you want to work with. The list of name filters that have been defined for the domain is displayed. Depending on whether name filters have already been defined for the domain, buttons are enabled that allow you to work with the name filters.

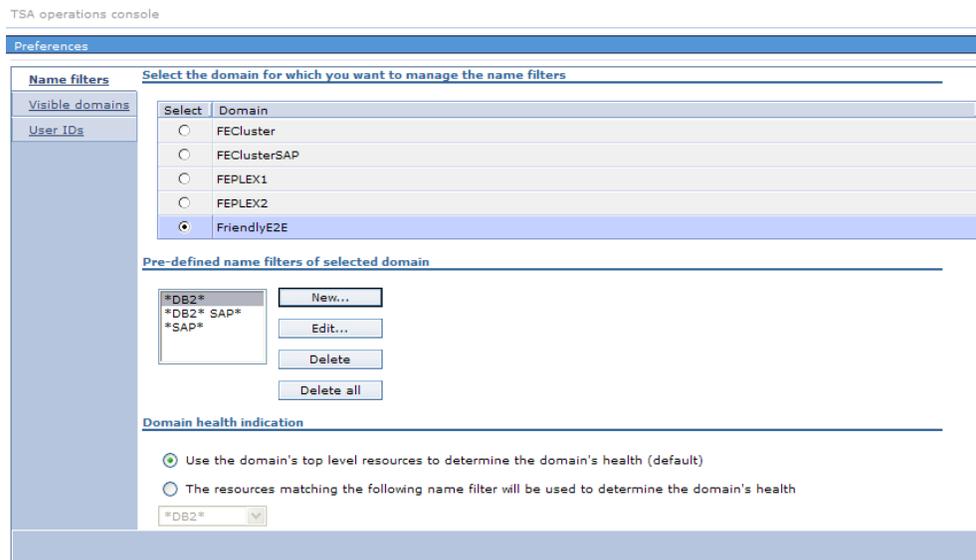


Figure 20. Name filters page on the Preferences panel

4. You use the buttons to perform the following tasks:
 - New** Opens the Name filters page that allows you to specify a new name filter.
 - Edit** Opens the Name filters page on which you can edit the name filter you have selected. The button is only enabled when you have already defined a name filter for the selected domain.
 - Delete** Deletes the name filter you have selected. The button is only enabled when you have already defined a name filter for the selected domain.
 - Delete all** Deletes all name filters that have been defined for the selected domain. The button is only enabled when you have already defined a name filter for the selected domain.

Displaying only resources against which operator requests were submitted

You can limit the scope of the resource table to resources against which operator requests were submitted. You use this option separately or combine it with a name filter.

Perform the following steps use the option:

1. Select the domain or the node.

2. Click **Search**. The Search page is displayed.

3. Select the check box **Only resources with operator requests**.

4. Click **OK**.

Results:

- The search results view of the resource table is displayed. Depending on whether you selected a domain or a node in the topology tree, the table lists only the resources of the selected domain or node against which operator requests have been submitted.
- You return to the group hierarchy view by clicking **Clear results**.

Hiding domains

By default, all domains are displayed in the topology tree. You can limit the scope of the topology tree by hiding domains from view, for example, domains in which you are not interested or for which you are not authorized. This has the advantage that you will no longer be prompted for your user credentials for these domains.

Perform the following steps:

1. Open the Preferences panel (**Menu** —>**Preferences**)

2. Click the Visible domains tab to open the Visible automation domains page. The page shows a hierarchical view of the available domains.

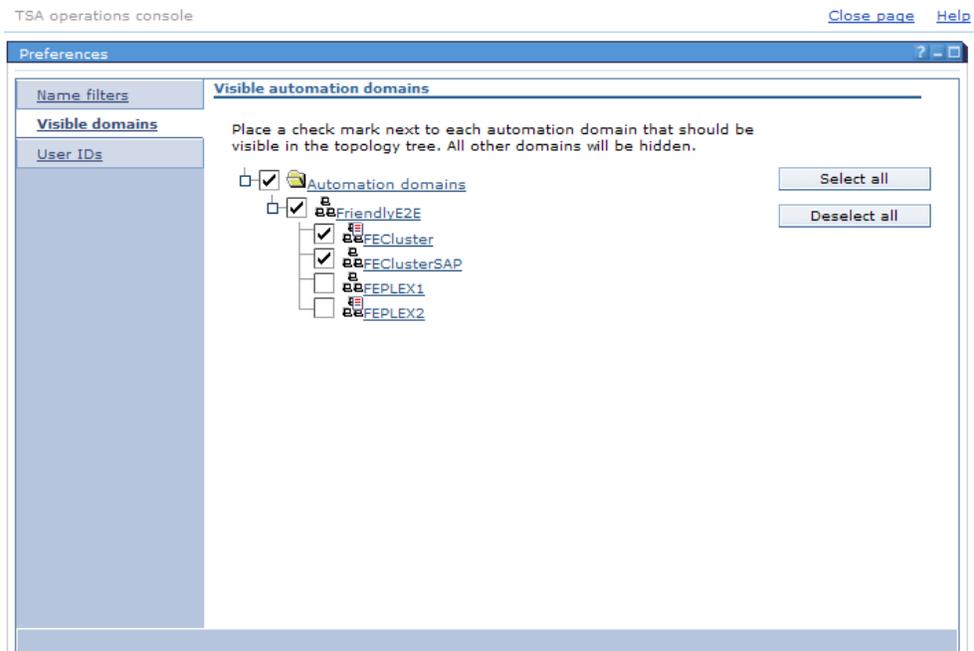


Figure 21. Visible automation domains page

3. Deselect the domains that should not appear in the topology tree and click **OK**.

Result: The topology tree only shows the selected domains and you will receive events for these domains only.

Using non-top-level resources as domain health indicators

Domain health indicators are resources whose state is used to indicate whether a domain is healthy. When such a resource goes into a warning or error state, a warning or error icon appears in the Status column of the topology tree for the domain that hosts the affected resource.

By default, the top-level resources of a domain are used as domain health indicators, but you can specify that other resources are to be used as domain health indicators by performing the steps below.

To specify which resources are to be used as domain health indicators, you use a name filter, either an existing one or one that you create specifically for the purpose.

Perform the following steps:

1. Open the Preferences panel (**Menu** → **Preferences**).
2. Open the Name filters page.
3. Select the domain from the list of domains.
4. If the filter you want to use is already available, proceed with step 5.

If you want to use new filter, click **New** and define the name filter on the panel that appears.

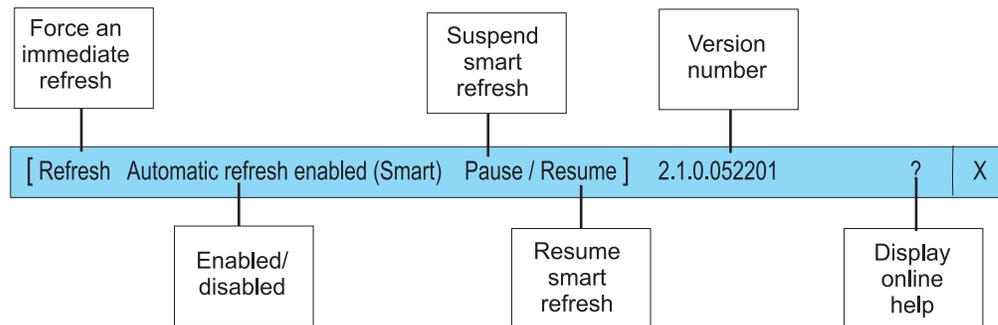
- At the bottom of the Name filters page, select the check box *The resources matching the following name filter will be used to determine the domain's health*. The list of available filters below the radio button is now active.
- Select a filter from the list and click **OK**.

Result: The resources that match the criteria defined in the selected filter will be used as domain health indicators.

Refreshing the operations console

The smart refresh function of the operations console checks at short intervals whether new information is available for any of the displayed elements. If new information is available, for example, when the state of a resource has changed, the operations console is updated with the new information.

On the smart refresh bar, you can force an immediate smart refresh, suspend and resume the smart refresh, and you can see whether smart refresh is enabled or suspended. The smart refresh bar is displayed at the bottom of the operations console:



Note: A smart refresh only updates the information on the operations console that has changed since the last smart refresh. This usually guarantees that the information displayed on the console reflects the actual current state of all elements.

In rare cases, you may want to use **Refresh all (Menu → Refresh all)** to update the operations console. **Refresh all** retrieves the latest information for all elements that are displayed on the operations console from the automation managers and updates the complete contents of the operations console regardless of whether or not the information has changed.

The following controls and fields are available on the smart refresh bar:

Refresh

Click **Refresh** to force an immediate smart refresh.

Pause Click **Pause** to suspend the smart refresh.

The smart refresh will resume automatically when you click a button or link on the operations console, or select, expand or collapse an element in the topology tree or in the resource table. To resume the refresh manually, click **Resume**.

Resume

To reactivate the smart refresh, click **Resume**.

Automatic refresh enabled/Automatic refresh disabled

This field shows whether smart refresh is enabled or disabled.

Version number

This field shows the version number of the operations console you are using.

? Click ? to display the online help for the smart refresh bar.

Switching to a different end-to-end automation manager

When you want to switch to a different end-to-end automation manager, for example, to connect to a first-level automation domain that the automation manager you are currently connected to does not know, you have the following options:

- You can launch an *additional* operations console within the same Integrated Solutions Console. This allows you to quickly switch between automation managers.

To launch an additional operations console, click **SA operations console** in the navigation tree of Integrated Solution Console. From the Connect panel that appears, connect to the server.

The new instance of the operations console is displayed in the work area of Integrated Solutions Console and an entry for that instance is added to the page bar. You use the entries on the page bar to switch between the instances of the operations console.

- You can connect to a different end-to-end automation manager from the **Menu** of the operations console you are currently working with. In this case, the operations console you are currently displaying will be closed. This is described below.

Steps for connecting to a different end-to-end automation manager from the operations console

1. On the operations console, click **Menu** → **Connect**. The connection panel is displayed.

2. Select the name of the server you want to connect to from the **Server name** drop-down list or enter the name and port number of the server and click **OK**.

Result: The operations console displaying the resources of the selected end-to-end automation manager replaces the currently displayed operations console.

Managing your user credentials for first-level automation domains

You can store the user IDs and passwords you need for logging on to a first-level automation domain in the Integrated Solutions Console credential vault. From there, the information is retrieved automatically when needed. This saves you from having to enter your user ID and password in these cases.

Note: The automation engine of the end-to-end automation management component does not use the user IDs and passwords that are stored in the credential vault to authenticate itself to first-level automation domains. The credentials of the automation engine are stored in a properties file and managed on the User credentials page of the configuration dialog (see Chapter 29, “Configuring the end-to-end automation manager,” on page 193).

Perform the following steps to manage your user credentials for first-level automation domains:

1. Open the Preferences panel (**Menu** → **Preferences**).

-
2. Open the User IDs page.

-
3. Select a domain from the list of domains.

-
4. You have the following options:

- To add or change your credentials, click **Edit**. The Credentials panel is displayed.

On the Credentials panel, create or change your user credentials. To store the credentials in the credential vault, ensure that the check box on the panel is selected.

- To delete a specific set or all of your credentials, click **Delete** or **Delete all**.
-

Chapter 34. Managing resources

In end-to-end automation management, managing resources comprises the following tasks:

- Activating a policy
- Starting or stopping a resource or a group of resources
- Excluding a node that is managed by a first-level automation manager from automation and including it in automation again
- Resetting a resource from an unrecoverable error
- Starting and stopping a choice groups or changing its preferred member

These management tasks are performed from the operations console. This chapter provides the background information you need to manage the resources and describes how the tasks are performed on the operations console.

Working with policies

The tasks described in this section can only be performed for the end-to-end automation domain.

Activating a policy

Steps for checking the validity of a policy from the operations console

You can invoke the policy checking tool from the operations console by proceeding as if you wanted to activate the policy, but without actually activating it. This ensures that the policy is ready for use when you actually want to activate it.

Before you begin:

- The policy file must be available in the policy pool directory.
- The domain name specified in the element `<AutomationDomainName>` in the policy file must be identical to the domain name specified in the field **Domain name** on the Domain page of the configuration dialog.

Perform the following steps:

1. Log in to the operations console.

-
2. In the topology tree, select the end-to-end automation domain.

-
3. On the Policy page in the information area, click **Activate policy**.

The policy list is displayed. If errors or warnings were issued during the validity check that was performed automatically when you opened the policy list, you see a warning or an error icon in the rightmost column of the policy table.

-
4. Select the policy in the policy list.

When problems were detected in the XML file, a button becomes available that lets you display the list of messages that were issued during the check.

5. Click **Cancel** to close the policy list. If errors were found in the file, you must correct them before the policy can be activated. Although warnings do not prevent the policy from being activated, you should check if they cannot be avoided.
-

Result: Repeat the procedure until all problems in the file are resolved.

Steps for activating a policy

Before you begin:

- The policy must be available in the policy pool directory.
- The validity of the policy has been checked and all errors that would prevent the activation of the policy have been corrected.

Perform the following steps:

1. Log in to the operations console.

2. In the topology tree, select the end-to-end automation domain.

3. On the Policy page, click **Activate policy**. The policy list is displayed.

4. Select the policy you want to activate.

5. Click **Activate** to activate the policy.

Note: If you try to activate a policy that is already active, you receive a warning.

Results:

- A confirmation message is displayed on the information bar, indicating that the command to activate the policy has been submitted for processing.
- The policy is activated and the policy definitions are used for automation management.

Deactivating a policy

You need to deactivate a policy if an active policy causes severe problems that cannot be resolved in any other way.

Perform the following steps:

1. Log in to the operations console.

2. In the topology tree, select the end-to-end automation domain.

3. On the Policy page of the end-to-end automation domain, click **Deactivate policy**.

Result:

- A confirmation message is displayed on the information bar, indicating that the command to deactivate the policy has been submitted for processing.

- The automation policy is deactivated.
- All automation requests that were propagated to the first-level automation domains are canceled.
- End-to-end automation management is suspended until a new policy is activated.

Modifying a policy

Modified policies are treated like new policies. Before you activate a modified policy:

- Make sure that you have updated the version information in the PolicyToken tag in the XML policy file.
- Check the validity of the policy as described in “Steps for checking the validity of a policy from the operations console” on page 249 and correct any errors.

To activate the policy, proceed as described in “Activating a policy” on page 249.

Working with requests

In automation management, you start and stop resources by changing their desired state. This you achieve by submitting start or stop requests that ask the automation manager to bring a resource online or offline. The automation manager will only change the desired state of a resource when your request wins. When your request wins, the actual resource will only be started or stopped after all relationships have been fulfilled. (For a detailed description of how start and stop requests are processed by the automation manager, refer to Chapter 5, “Automation concepts,” on page 27)

For submitting requests, the following rules apply:

- Start requests can only be submitted against resources in desired state Offline.
- Stop requests can only be submitted against resources in desired state Online.
- Requests cannot be submitted if another operator request has already been submitted against the resource. In this case, the operator request must be canceled to change the desired state of the resource.
- Requests cannot be submitted against members of a choice group but must be submitted against the group. This will bring the preferred member online or offline.
- Requests should not be submitted against first-level automation resources that are referenced by a resource reference. Only when you submit the request against the resource reference it is ensured that all relationships are fulfilled before the resource is started or stopped.
- Requests cannot be submitted against monitor resources. For such resources, the buttons for submitting requests are not available on the operations console.

Submitting start requests

Perform the following steps to submit a start request:

1. In the resource table, select the resource you want to start.

-
2. On the **General** page, click **Request Online**.

The Request Online panel is displayed.

3. Use the entry field on the Request Online panel to provide a short description of why you want to change the automation goal of the resource to Online.

4. Click **Submit** to submit the request.

Results:

- A confirmation message is displayed on the information bar, indicating that the request has been submitted for processing.
- After the next refresh, resource icon is highlighted with the yellow operator icon, indicating that a request was issued against the resource.
- The request is processed. Processing of the request is complete when the resource has been started.

Submitting stop requests

Perform the following steps to submit a stop request:

1. In the resource table, select the resource you want to stop.

2. On the General page, click **Request Offline**.
The Request Offline panel is displayed.

3. Use the entry field to provide a short description of why you want to change the automation goal of the resource to Offline.

4. Click **Submit** to submit the request.

Results:

- A confirmation message is displayed on the information bar, indicating that the request has been submitted for processing.
- After the next refresh, resource icon is highlighted with the yellow operator icon, indicating that a request was issued against the resource.
- The request is processed. Processing of the request is complete when the resource has been stopped.

Displaying information about an operator request

When an operator has submitted a start or stop request against a resource, an operator request icon appears on the General page for the resource. The icon indicates the status of the request:

Table 46. Operator request icons in the information area

| Operator request icon | Description |
|---|---|
|  | A stop request has been submitted. The yellow operator icon indicates that the observed state of the resource is not Offline yet. |
|  | A start request has been submitted. The yellow operator icon indicates that the observed state of the resource is not Online yet. |
|  | The green operator icon indicates that the stop request has been completed successfully. The observed state of the resource is Offline. |

Table 46. Operator request icons in the information area (continued)

| Operator request icon | Description |
|---|---|
|  | The green operator icon indicates that the start request has been completed successfully. The observed state of the resource is Online. |

This is how you can display more information about the request:

- Move the mouse over the operator request icon to display the user ID of the operator who submitted the request.
- Click the operator request icon to bring up the Request details panel.

Displaying request lists

All requests and votes (internal requests that were propagated due to relationships) that have been submitted against a resource are added to the resource's request list. You can display the list to find out which requests and votes have been issued and which of the requests wins. The list is sorted by priority with the winning request listed at the top.

The list contains information about each request or vote, for example:

- its source (for example, the name of the operator who submitted the request)
- its priority
- the creation date and time

From the Request list panel, you can display detailed information about each of the requests or votes, including the comments that were added by operators when they submitted the request.

Steps for viewing a request list and request details

Perform the following steps:

1. In the resource table, select the resource whose request list or request details you want to view.

2. On the General page, click **View requests**.
The Request list is displayed. The list is sorted by priority. The first entry is the winning request.

3. To display the details for a request, select the resource in the list and click **More info**.
The Request details panel is displayed.

Canceling requests

You can cancel operator requests that have been submitted against resources. Votes and requests generated by automation managers cannot be canceled.

This is what happens when you cancel a request:

- When you cancel a request that did not win, you prevent it from being completed at a later time.

- When you cancel the request that is responsible for the current desired state of the resource, you change the desired state of the resource to the opposite if there are no other requests or votes in the request list that will win when the canceled request is removed.
- When you cancel a request, votes that were generated against other resources because of StartAfter or StopAfter relationships are canceled as well.

Steps for canceling requests

Perform the following steps to cancel a request:

1. Select the resource in the resource table.
-

2. On the **General** page, click **Cancel request**.

The button is only enabled if there is an operator request in the request list of the resource.

The text to the left of the **Cancel request** button describes the resource's expected desired state after the request has been canceled. The expected desired state is calculated in this way:

- If there are other requests or votes in the request list, the winning request determines the expected desired state.
- If there are no other request or votes in the list, the desired state that is defined in the policy becomes the automation goal.

The desired state that is actually set after cancelation can differ from the expected state, for example, when a new request or vote is generated at the same time or immediately after you canceled the request.

Resetting a resource from an unrecoverable error

When a resource becomes available for automation management again after an unrecoverable error was resolved by an operator, the automation manager will not start automating the resource again without your intervention. When the resource is available again, you must inform the automation manager that the resource can be included in automation management again. You do this by using the Reset function on the operations console. The Reset function is only available for first-level automation resources and resource references that are in state Unrecoverable error.

Steps for resetting a resource

Perform the following steps:

1. Select the resource in the resource table.
-

2. On the General page, click **Reset** to include the resource in automation management again.
-

Results:

- A confirmation message is displayed on the information bar, indicating that the command to reset the resource has been submitted for processing.
- Automation management for the resource will resume:
 - When you have reset a first-level automation resource, the resource will be managed by the first-level automation manager again.

- When you have reset a resource reference, the end-to-end automation manager will take over again. If the referenced first-level automation resource also was in state Unrecoverable error, the reset will be propagated to the referenced resource.

Including a node in automation and excluding a node from automation

From the operations console, you can exclude a node from first-level automation, for example, for maintenance purposes, and include it again when you want the automation manager to take over again:

- When you exclude a node, the corresponding command is sent directly to the first-level automation manager. The first-level automation manager stops all resources that are running on the node and moves them to a different node if possible.

As the command is sent directly to the first-level automation manager, the end-to-end automation manager is not informed of the fact that the resources were stopped deliberately by an operator. However, as most of the first-level automation resources will be moved to a different node and run there, the automation manager will not even realize that these resources were stopped at their original location.

For the resources that could not be moved, however, end-to-end automation management may not be successful while they are down. For resources for which a resource reference exists and that have the desired state Online, the end-to-end automation manager will unsuccessfully issue start requests, and the resource references pointing to these resources will go into warning state. The start requests sent by the end-to-end automation manager will be retained and, if they win, be completed when the node is included again.

- When you include a node in automation again, the first-level automation manager will start the resources whose automation goal is Online. All resources that are located on the node will automatically be included in first-level and end-to-end automation again.

Steps for excluding a node from automation

To exclude a node from automation, perform the following steps:

1. Select the node in the topology tree.

-
2. On the General page, click **Exclude node**.

Before the exclude command is sent to the first-level automation manager, you will be asked to confirm the action. Click **OK** to send the exclude command to the first-level automation manager.

Results:

- A confirmation message is displayed on the information bar, indicating that the exclude node command has been submitted for processing.
- The first-level resource manager will stop all resources that are running on the node, moving them to a different node if possible.

Steps for including a node in automation

Perform the following steps:

1. Select the node in the topology tree.
-

2. On the General page, click **Include node**.

Note: The button is only available if the node is currently excluded from automation.

Results:

- A confirmation message is displayed on the information bar, indicating that the include node command has been submitted for processing.
 - The first-level automation manager will start all resources on the node whose automation goal is Online. First-level and end-to-end automation for the resources will commence.
-

Working with choice groups

Choice groups are end-to-end automation resources. They have the following characteristics:

- The members are configuration alternatives that provide the same functionality (for example, two database instances where one is used as the production database and the other serves as backup).
- Only one of the members can be online at a time.
- Members can be either resource groups or resource references. The first-level automation resources which are referenced by the members of a choice group can be located on different nodes or hosted by different domains.
- One member of the choice group is defined as the so-called preferred member. When the desired state of the choice group is Online, the preferred member is kept online by the automation manager while the other members are kept offline.
- When a member other than the preferred member is to be brought online, the preferred member must be changed.

When you want to change the desired state of a choice group or bring a member other than the currently preferred member online, the following rules apply:

- Start or stop requests must be submitted against the choice group, not against an individual member (see “Steps for starting the preferred member of a choice group”).
- To bring a member other than the currently preferred member online, you change the preferred member of the choice group by using a simple function on the operations console. Changing the preferred member for a choice group whose desired state is online, leads to the following results:
 - the old preferred member is brought offline if it is still online
 - the new preferred member of the group is brought online and kept online by the automation manager.

This is described in “Steps for starting a different member of a choice group” on page 257.

Steps for starting the preferred member of a choice group

Perform the following steps to start the preferred member of a choice group whose current state is Offline:

1. In the resource table, select choice group whose preferred member you want to start.
-

2. On the General page, click **Request online**.
-

Results:

- A confirmation message is displayed on the information bar, indicating that the request to start the resource has been submitted for processing.
- When the request has been completed:
 - the preferred member is online
 - the automation manager will try to keep the preferred member online and the other members offline

Steps for starting a different member of a choice group

Use the procedure described below:

- for choice groups whose desired state is Online
- and the preferred member of the choice group has failed or needs to be stopped
- and a different member of the choice group is to be started

Note: You can also use this procedure for choice groups whose desired state is Offline, for example, because you want to be sure that a member other than the currently preferred member is started when a start request is issued for the group. In such a case, only the preferred member setting is changed. The automation manager will continue to try to keep all members of the group offline.

Perform the following steps:

1. Select the choice group in the resource table.
-

2. In the Possible Choices table on the General page, select the choice group member that you want to start. Below the table, the button **Set as preferred** appears.
-

3. Click **Set to preferred**.

If the desired state of the choice group is Online, this will trigger the following actions:

- If the old preferred member is online, it is stopped.
- The new preferred member is started.
- The automation manager will try to keep the new preferred member online and the other members offline.

If the desired state of the choice group is Offline, just the setting for the preferred member is changed, the automation manager will continue to try to keep all members of the choice group offline.

Part 6. Appendixes

Appendix A. Default directories

The following table lists the default directory paths for which variables are used in this guide. The paths in your environment may differ, for example, if you changed the default path during the installation of the application or component.

Table 47. Default directories

| Variable used in this guide | Default path |
|-----------------------------|--|
| <DB2_install_location> | Windows: C:\Program Files\IBM\SQLLIB AIX: /usr/opt/db2_08_01 Linux: /opt/IBM/db2/V8.1 |
| <EEZ_CONF_ROOT> | Windows: C:\Program Files\IBM\tsamp\eez\cfg AIX/Linux: /etc/opt/IBM/tsamp/eez/cfg |
| <EEZ_INSTALL_ROOT> | Windows: C:\Program Files\IBM\tsamp\eez AIX/Linux: /opt/IBM/tsamp/eez The configuration properties files are located in the directory <EEZ_CONF_ROOT>. |
| <isc_runtime_root> | Windows: C:\Program Files\IBM\ISC AIX/Linux: /opt/IBM/ISC |
| <Tivoli_Common_Directory> | Windows: C:\Program Files\IBM\tivoli\common AIX/Linux: /var/ibm/tivoli/common The path to the Tivoli Common Directory is specified in the properties file log.properties. The file log.properties is located in the following directory: <ul style="list-style-type: none">• Windows: C:\Program Files\IBM\tivoli\common\cfg• AIX/Linux: /etc/ibm/tivoli/common/cfg |
| <was_root> | Windows: C:\Program Files\IBM\WebSphere\AppServer AIX: /usr/IBM/WebSphere/AppServer Linux: /opt/IBM/WebSphere/AppServer |

Appendix B. Configuration properties files

This chapter provides a short description of the configuration properties files of the end-to-end automation management component. During the installation of the component, the properties in the files are set to values that allow you to work with the sample end-to-end automation domain as described in Part 3, “First steps,” on page 125.

To change the values of the properties, you use the configuration dialog of the end-to-end automation management component. This ensures that the files are not corrupted during manual editing and that the change history in the files is updated whenever a property is changed.

For information about the configuration dialog, refer to Chapter 29, “Configuring the end-to-end automation manager,” on page 193. For a detailed description of the configuration properties, refer to the configuration dialog help.

The configuration properties files are located in the following directory:

- **Windows:** <EEZ_INSTALL_ROOT>/cfg
- **AIX/Linux:** /etc/opt/IBM/tsamp/eez/cfg

These are the properties files that are changed when you modify a property value in the configuration dialog:

eez.automation.engine.properties

The properties in this file are used to configure exactly one instance of the automation engine. The configuration properties specify, for example, the name of the end-to-end automation domain and the location of the policy pool directory.

You browse and edit the properties that are in the file on the Domain page of the configuration dialog.

eez.automation.engine.dif.properties

The domain identification file of the automation engine contains the user IDs and the corresponding passwords that the end-to-end automation engine uses to authenticate itself to first-level automation domains and to the WebSphere Application Server JMS Provider.

The file must be protected by means of operating system mechanisms in such a way that the automation engine is still able to read the contents of the file.

You browse and edit the properties that are in the file on the User credentials page of the configuration dialog.

eez.fl.a.ssl.properties

This file contains the configuration properties for the SSL connection to the first-level automation domains.

You browse and edit the properties that are in the file on the Security page of the configuration dialog.

eez.jlog.properties

The properties in this file determine which information is written to the log and trace files of the automation engine.

You browse and edit the properties that are in the file on the Logger page of the configuration dialog.

Appendix C. Automation policy reference

Worksheet for gathering data for defining a policy

Use this worksheet for collecting the information about the first-level automation resources for which you want to define resource references in the end-to-end automation policy.

Table 48. Worksheet for defining an end-to-end automation policy

| | | | |
|-------|-------------------------------|----------------------------------|--|
| 1.1 | First-level automation domain | Domain name | |
| 1.2 | | Host name | |
| 1.3 | | Owner | |
| 1.4 | | User ID for accessing the domain | |
| 2.1.1 | Resource information | Name | |
| 2.1.2 | | Class | |
| 2.1.3 | | Node | |
| 2.1.4 | | Owner | |
| 2.1.5 | | Description | |
| 2.1.6 | | URL for InfoLink | |
| 2.1.7 | | Relationship(s) to | |
| 2.2.1 | Resource information | Name | |

Table 48. Worksheet for defining an end-to-end automation policy (continued)

| | | | |
|-------|--|----------------------|------|
| 2.2.2 | | Class | |
| 2.2.3 | | Node | |
| 2.2.4 | | Owner | |
| 2.2.5 | | Description | |
| 2.2.6 | | URL for InfoLink | |
| 2.2.7 | | Relationship(s) to | |
| 2.3.1 | | Resource information | Name |
| 2.3.2 | | Class | |
| 2.3.3 | | Node | |
| 2.3.4 | | Owner | |
| 2.3.5 | | Description | |
| 2.3.6 | | URL for InfoLink | |
| 2.3.7 | | Relationship(s) to | |
| 2.4.1 | | Resource information | Name |
| 2.4.2 | | Class | |
| 2.4.3 | | Node | |
| 2.4.4 | | Owner | |
| 2.4.5 | | Description | |
| 2.4.6 | | URL for InfoLink | |
| 2.4.7 | | Relationship(s) to | |

Policy XML reference

This section lists the XML elements, sub-elements, and attributes of end-to-end automation policies in alphabetical order. For each element, the following information is provided:

Name The name and a short description of the element or sub-element.

Restrictions

Possible restrictions that apply.

Examples

One or more examples showing how the element is used.

Attributes

Descriptions of the element's attributes.

Contexts

A list of elements that can directly contain the element (its "Parents") and a list of elements the element can contain (its "Children").

AutomationDomain

The AutomationDomain element is used within a ResourceReference specification. It is used to specify the name of the first-level automation domain that hosts the ReferencedResource.

Restrictions

Maximum number of characters supported: 64

Example

```
<ResourceReference name="WebSphere EE">
  <DesiredState>Offline</DesiredState>
  <Owner>Harry Smith</Owner>
  <InfoLink>http://www.example.com/WebSphere/help</InfoLink>
  <ReferencedResource>
    <AutomationDomain>FEPLEX2</AutomationDomain>
    <Name>WAS</Name>
    <Class>ResourceGroup</Class>
    <Node>node1</Node>
  </ReferencedResource>
</ResourceReference>
```

Attributes

None

Contexts

Children: None

Parents: ReferencedResource

AutomationDomainName

The AutomationDomainName element specifies the name of the end-to-end automation domain for which the policy will be used. A policy can only be activated if the name specified in its AutomationDomainName element is identical to the name specified on the Domain page of the configuration dialog (field Domain name).

Restrictions

Maximum number of characters supported: 64

Example

```
<PolicyInformation>
  <PolicyName>Sample E2E Policy</PolicyName>
  <AutomationDomainName>FriendlyE2E</AutomationDomainName>
  <PolicyToken> 1.0.1</PolicyToken>
  <PolicyAuthor>Michael Atkins</PolicyAuthor>
  <PolicyDescription>Sample policy for domain
    FriendlyE2E</PolicyDescription>
</PolicyInformation>
```

Attributes

None

Contexts

Children: None

Parents: PolicyInformation

AutomationPolicy

The AutomationPolicy element is the top-level element of the end-to-end automation policy. It serves as a container for all other elements of the policy.

Restrictions

The element is required and must occur only once in a policy.

Example

```
<AutomationPolicy version="1.0"
  xmlns="http://www.ibm.com/TSA/Policy.xsd"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.ibm.com/TSA/Policy.xsd EEZPolicy.xsd">
  ...
</AutomationPolicy>
```

Attributes

version

Identifies the minimum version of the end-to-end automation component. The attribute is required. Maximum number of characters supported: 32

xmlns Name space declaration for this policy. The attribute is required and must be specified as presented in the example above.

xmlns:xsi

Identifies the XML schema format for this policy XML. The attribute is required and must be specified as presented in the example above.

xsi:schemaLocation

Identifies the XML schema to which this XML policy must conform. The attribute is required and must be specified as presented in the example above.

Contexts

Must be the top-level element of a policy and must contain all other elements.

Children: PolicyInformation, ResourceReference ...

Parents: None

ChoiceGroup

The ChoiceGroup element is used as a logical container for a collection of redundant, equivalent end-to-end automation domain resources of which one is

defined as the *preferred* member. If the group's desired state is Online, the *preferred* member is active and the other members of the group serve as standbys. The automation manager will try to keep the *preferred* member online and will restart it in place if it fails. If the *preferred* member fails or is to be stopped, an operator can start one of the remaining members of the ChoiceGroup.

Restrictions

- A ChoiceGroup must not be a member of a ResourceGroup.

Example

```
<ChoiceGroup name="HTTP Server">
  <DesiredState> Offline </DesiredState>
  <Description>Choice group for choosing one HTTP Server</Description>
  <Owner>Jenny Parker</Owner>
  <InfoLink>http://www.friendly.com/choice</InfoLink>
  <Members>
    <ResourceReference name="HTTP Server Prim" preferred="true"/>
    <ResourceReference name="HTTP Server Backup"/>
  </Members>
</ChoiceGroup>
```

Attributes

name="ChoiceGroupName"

The attribute is used to specify the name of the ChoiceGroup. The name must be unique within an XML policy document, that is, the name that is used for a ChoiceGroup must not be used for any other ChoiceGroup, ResourceGroup or ResourceReference in the same XML policy. The attribute is required.

The name can have a maximum of 64 characters. Do not use more than one blank to separate strings within the name. Duplicate blanks will be ignored.

Contexts

Children: Members, DesiredState, Description, Owner, InfoLink
Parents: AutomationPolicy

Class

The Class element is used within a ResourceReference definition. It specifies the class of the first-level automation ReferencedResource (for example, IBM.ResourceGroup). The element is optional, but must be specified for ReferencedResources that are managed by SA for Multiplatforms or SA z/OS.

Restrictions

- Maximum number of characters supported: 64
- SA for Multiplatforms: Only create resource references to first-level automation resources of Class type IBM.ResourceGroup. Resources references that refer to first-level automation resources of class types IBM.Application, IBM.ServiceIP, and IBM.Test may cause undesired effects or severe problems in first-level automation management.

Example

```
<ResourceReference name="WebSphere EE">
  <DesiredState>Offline</DesiredState>
  <Description>Resource reference WebSphere EE </Description>
  <Owner>Harry Smith</Owner>
  <InfoLink>http://www.example.com/WebSphere/help</InfoLink>
  <ReferencedResource>
```

```

        <AutomationDomain>FEPLEX2</AutomationDomain>
        <Name>WAS</Name>
        <Class>ResourceGroup</Class>
        <Node>node1</Node>
    </ReferencedResource>
</ResourceReference>

```

Attributes

None

Contexts

Children: None

Parents: ReferencedResource

Description

The Description element is used to add a free text that provides further information about the parent element. The element is optional.

Restrictions

Maximum number of characters supported: 1024

Maximum number of occurrences per parent element: 1

Attributes

None

Contexts

Children: None

Parents: ChoiceGroup, ResourceGroup, ResourceReference

DesiredState

The DesiredState element is used to define the default desired state (the automation goal) for its parent element. Possible values are Online and Offline. The end-to-end automation manager will try to keep the resource in the DesiredState.

The element is optional.

Restrictions

Maximum number of occurrences per parent element: 1

As the element is optional, it can be omitted. In such a case, the following rules apply:

- If the desired state for a top-level resource reference is not defined in the policy, the default desired state is set to Online.
- If the desired state for a top-level resource group is not defined in the policy, the default desired state is set to Online.

The end-to-end automation manager will try to keep *all* members of the resource group (including nested resource groups and their members) online. Desired state settings for individual resource references or nested resource groups are ignored.

- If the desired state for a choice group is not defined in the policy, the default desired state is set to Online. The end-to-end automation manager will try to keep the preferred member online and the other members offline.

Example

```

<ResourceReference name="WebSphere EE">
    <DesiredState>Offline</DesiredState>
    <Description>Resource reference WebSphere EE</Description>

```

```

    <Owner>Harry Smith</Owner>
    <InfoLink>http://www.example.com/WebSphere/help</InfoLink>
    <ReferencedResource>
      <AutomationDomain>FEPLEX2</AutomationDomain>
      <Name>WAS</Name>
      <Class>ResourceGroup</Class>
      <Node>node1</Node>
    </ReferencedResource>
  </ResourceReference>

```

Attributes

None

Contexts

Children: None

Parents: ChoiceGroup, ResourceGroup, ResourceReference

InfoLink

The element InfoLink is used to specify the URL of an HTML page where the operator can find additional information about the resource.

Restrictions

Maximum number of characters supported: 1024

Example

```

<ResourceReference name="WebSphere EE">
  <DesiredState>Offline</DesiredState>
  <Description>Resource reference WebSphere EE</Description>
  <Owner>Harry Smith</Owner>
  <InfoLink>http://www.example.com/WebSphere/help</InfoLink>
  <ReferencedResource>
    <AutomationDomain>FEPLEX2</AutomationDomain>
    <Name>WAS</Name>
    <Class>ResourceGroup</Class>
    <Node>node1</Node>
  </ReferencedResource>
</ResourceReference>

```

Attributes

None

Contexts

Children: None

Parents: ChoiceGroup, ResourceGroup, ResourceReference

Node

The element Node is used within a ResourceReference specification. It is used to specify the host name (SA for Multiplatforms) or system name (SA z/OS) where the first-level automation resource (the ReferencedResource) is located. The element is optional.

Restrictions

- Maximum number of characters supported: 256
- Host name or system name must be specified in first-level domain syntax.
- The Node element is optional. Only specify the Node element when you create a resource reference for a fixed resource. Do not specify the node for any other type of first-level automation resource.

Note that creating resource references for fixed resources is **not** recommended.

Example

```
<ResourceReference name="NFS Client">
  <DesiredState>Offline</DesiredState>
  <Description>Resource reference NFS Client </Description>
  <Owner>Henry Smith</Owner>
  <InfoLink>http://www.example.com/help/NFS/?help+NFS+Client</InfoLink>
  <ReferencedResource>
    <AutomationDomain>cm1cm2</AutomationDomain>
    <Name>nfs-client</Name>
    <Class>IBM.ResourceGroup</Class>
    <Node>1nxcm2</Node>
  </ReferencedResource>
</ResourceReference>
```

Attributes

None

Contexts

Children: None

Parents: ReferencedResource

Members

The Members element is a container element. It contains the list of resources that are members of an end-to-end automation ResourceGroup or ChoiceGroup.

Example

Example of a Members element in a ResourceGroup specification:

```
<ResourceGroup name="Friendly Computer Shop" >
  <DesiredState>Offline</DesiredState>
  <Description>Resource group Friendly Computer Shop</Description>
  <Owner>Jerr Owens</Owner>
  <InfoLink>http://www.example.com/help/policy/compshop.html</InfoLink>
  <Members>
    <ResourceGroup name="mySAP Solutions"/>
    <ResourceReference name="WebSphere AE"/>
  </Members>
</ResourceGroup>
```

Example of a Members element in a ChoiceGroup specification:

```
<ChoiceGroup name="HTTP Server">
  <DesiredState> Offline </DesiredState>
  <Description>Choice group for choosing one HTTP Server</Description>
  <Owner>Jenny Parker</Owner>
  <InfoLink>http://www.example.com/choice</InfoLink>
  <Members>
    <ResourceReference name="HTTP Server Prim" preferred="true"/>
    <ResourceReference name="HTTP Server Backup"/>
  </Members>
</ChoiceGroup>
```

Attributes

None

Contexts

Children: ResourceReference, ResourceGroup, ChoiceGroup

Parents: ChoiceGroup, ResourceGroup

Name

The Name element is used within a ResourceReference specification. It is used to specify the actual name the ReferencedResource in its first-level domain. This name may be different from the name specified for the ResourceReference that refers to this resource. The element is required.

Restrictions

Maximum number of characters supported: 64

Example

```
<ResourceReference name="SAP AppServer">
  <DesiredState>Offline</DesiredState>
  <Description>Resource reference SAP AppServer </Description>
  <Owner>Jenny Parker</Owner>
  <InfoLink>http://www.example.com/help/policy/SAPAPP.html</InfoLink>
  <ReferencedResource>
    <AutomationDomain>FEClusterSAP</AutomationDomain>
    <Name>SAP AppServer</Name>
    <Class>ResourceGroup</Class>
    <Node>node4</Node>
  </ReferencedResource>
</ResourceReference>
```

Attributes

None

Contexts

Children: None

Parents: ReferencedResource

Owner

The element Owner is used to provide information about the owner of the parent element, for example, the name and telephone number of the owner of a referenced resource. The element is optional.

Restrictions

The element can only occur once per parent element.
Maximum number of characters supported: 1024

Example

Attributes

None

Contexts

Children: None

Parents: ResourceReference, ChoiceGroup, ResourceGroup

PolicyAuthor

The PolicyAuthor element is used to identify the author of the policy.

Restrictions

Maximum number of characters supported: 64

Example

```
<PolicyInformation>
  <PolicyName>Sample E2E Policy</PolicyName>
  <AutomationDomainName>FriendlyE2E</AutomationDomainName>
```

```
<PolicyToken> 1.0.1</PolicyToken>
<PolicyAuthor>Michael Atkins</PolicyAuthor>
<PolicyDescription>Sample policy for domain FriendlyE2E</PolicyDescription>
</PolicyInformation>
```

Attributes

None

Contexts

Children: None

Parents: PolicyInformation

PolicyDescription

The PolicyDescription element is available for adding a free text to the XML document. It can be used, for example, to maintain a change history within the file.

Restrictions

Maximum number of characters supported: 1024

Example

```
<PolicyInformation>
  <PolicyName>Sample E2E Policy</PolicyName>
  <AutomationDomainName>FriendlyE2E</AutomationDomainName>
  <PolicyToken> 1.0.1</PolicyToken>
  <PolicyAuthor>Michael Atkins</PolicyAuthor>
  <PolicyDescription>Sample policy for domain FriendlyE2E
    Last Update: 08/16/04
    Last Editor: Michael Atkins
    Change History:
    -----
    Date      Name      Description
    -----
    08/16/05  Michal Atkins  Initial policy definition
    -----
  </PolicyDescription>
</PolicyInformation>
```

Attributes

None

Contexts

Children: None

Parents: PolicyInformation

PolicyInformation

The PolicyInformation element contains subelements that uniquely identify this policy.

Restrictions

The element is required and must occur only once in an XML policy file.

Example

```
<PolicyInformation>
  <PolicyName>Sample E2E Policy</PolicyName>
  <AutomationDomainName>FriendlyE2E</AutomationDomainName>
  <PolicyToken> 1.0.1</PolicyToken>
  <PolicyAuthor>Michael Atkins</PolicyAuthor>
  <PolicyDescription>Sample policy for domain FriendlyE2E</PolicyDescription>
</PolicyInformation>
```

Attributes

None

Contexts

Children: PolicyName, AutomationDomainName, PolicyToken, PolicyAuthor, PolicyDescription

Parents: AutomationPolicy

PolicyName

The PolicyName element is used to specify a meaningful name for the specific policy.

Restrictions

The element is required and must occur only once in an XML policy file.
Maximum number of characters supported: 64

Example

```
<PolicyInformation>
  <PolicyName>Sample E2E Policy</PolicyName>
  <AutomationDomainName>FriendlyE2E</AutomationDomainName>
  <PolicyToken> 1.0.1</PolicyToken>
  <PolicyAuthor>Michael Atkins</PolicyAuthor>
  <PolicyDescription>Sample policy for domain FriendlyE2E</PolicyDescription>
</PolicyInformation>
```

Attributes

None

Contexts

Children: None

Parents: PolicyInformation

PolicyToken

The PolicyToken element specifies the version of the policy. It should be updated whenever the policy is changed.

Restrictions

The element is required.

Maximum number of characters supported: 64

Example

```
<PolicyInformation>
  <PolicyName>Sample E2E Policy</PolicyName>
  <AutomationDomainName>FriendlyE2E</AutomationDomainName>
  <PolicyToken> 1.0.1</PolicyToken>
  <PolicyAuthor>Michael Atkins</PolicyAuthor>
  <PolicyDescription>Sample policy for domain FriendlyE2E</PolicyDescription>
</PolicyInformation>
```

Attributes

None

Contexts

Children: None

Parents: PolicyInformation

ReferencedResource

The ReferencedResource element is a child of the ResourceReference element. The children of the ReferencedResource element are used to describe the actual first-level automation resource to which the ResourceReference refers.

Restrictions

- The element must only occur once per parent element.
- The ReferencedResource as defined by the values specified for its children (AutomationDomain, Name, Class, and Node) must be unique within the policy, that is, the same ReferencedResource must not be referenced by more than one ResourceReference.
- SA for Multiplatforms: Resource references should always point to top-level resources (resource groups of resource class type IBM.Resource Group). Creating resource references for first-level automation resources of resource class types IBM.Application, IBM.ServiceIP, and IBM.Test is not recommended and may cause problems in first-level automation management.

Example

```
<ResourceReference name="Enterprise DB2">
  <DesiredState>Offline</DesiredState>
  <Description>Resource reference Enterprise DB2 </Description>
  <Owner>Bob Owens</Owner>
  <InfoLink>http://www.example.com/help/DB2</InfoLink>
  <ReferencedResource>
    <AutomationDomain>FEPLEX2</AutomationDomain>
    <Name>DB2</Name>
    <Class>ResourceGroup</Class>
    <Node>node1</Node>
  </ReferencedResource>
</ResourceReference>
```

Attributes

None

Contexts

Children: AutomationDomain, Name, Class, Node

Parents: ResourceReference

Relationship

The sub-elements of the Relationship element are used to define the relationship (StartAfter, StopAfter, ForcedDownBy) between two end-to-end automation resources (ResourceReferences or end-to-end automation groups). The element is optional.

Restrictions

Example

```
<Relationship>
  <Source>
    <ResourceReference name="DB2 Connect Client"/>
  </Source>
  <Type> StartAfter</Type>
  <Target>
    <ResourceReference name="Enterprise DB2"/>
  </Target>
</Relationship>
```

Attributes

None

Contexts

Children: Source, Target, Type
Parents: AutomationPolicy

ResourceGroup

The ResourceGroup element is used to define end-to-end automation resource groups. The element is optional.

Restrictions

ResourceGroups can be nested, but a ResourceGroup must not be a member of more than one ResourceGroup.
A ChoiceGroup must not be a member of a ResourceGroup.

Example

```
<ResourceGroup name="Friendly Computer Shop" >
  <DesiredState>Offline</DesiredState>
  <Description>Resource group Friendly Computer Shop</Description>
  <Owner>Jerr Owens</Owner>
  <InfoLink>http://www.example.com/help/policy/compshop.html</InfoLink>
  <Members>
    <ResourceGroup name="mySAP Solutions"/>
    <ResourceReference name="WebSphere AE"/>
  </Members>
</ResourceGroup>
```

Attributes

name="*ResourceGroupName*"

The attribute is used to specify the name of the ResourceGroup. The name must be unique within an XML policy document, that is, the name that is used for a ResourceGroup must not be used for another ResourceGroup, ChoiceGroup or ResourceReference within the same XML policy. The attribute is required.

The name can have a maximum of 64 characters. Do not use more than one blank to separate strings within the name. Duplicate blanks will be ignored.

Contexts

Children: Description, Owner, InfoLink, Members
Parents: AutomationPolicy

ResourceReference

This element is used in three contexts:

- It is used to define an end-to-end automation resource. Each XML policy must contain at least one ResourceReference element. The element has a number of required sub-elements.
- In ResourceGroup and ChoiceGroup definitions, it is used to specify which resources references are members of a group.
- In Relationship definitions, it is used to specify the source and target of a relationship.

Restrictions

The value of the Name attribute must be unique within the policy, that is, the name cannot be used for more than one ResourceReference.

The same name must not be used for a ChoiceGroup or ResourceGroup.

Examples

Example of a ResourceReference element in a ResourceReference specification:

```
<ResourceReference name="WebSphere AE">
  <DesiredState>Offline</DesiredState>
  <Description>Resource reference WebSphere AE </Description>
  <Owner>Harry Smith</Owner>
  <InfoLink>http://www.example.com/help/WebSphere.html</InfoLink>
  <ReferencedResource>
    <AutomationDomain>FECluster</AutomationDomain>
    <Name>WebSphere AppServer</Name>
    <Class>ResourceGroup</Class>
    <Node>node2</Node>
  </ReferencedResource>
</ResourceReference>
```

Example of a ResourceReference element in a ResourceGroup specification:

```
<ResourceGroup name="Friendly Computer Shop" >
  <DesiredState>Offline</DesiredState>
  <Description>Resource group Friendly Computer Shop</Description>
  <Owner>Jerr Owens</Owner>
  <InfoLink>http://www.example.com/help/policy/compshop.html</InfoLink>
  <Members>
    <ResourceGroup name="mySAP Solutions"/>
    <ResourceReference name="WebSphere AE"/>
  </Members>
</ResourceGroup>
```

Example of ResourceReference elements in a Relationship specification:

```
<Relationship>
  <Source>
    <ResourceReference name="WebSphere EE"/>
  </Source>
  <Type> StartAfter</Type>
  <Target>
    <ResourceReference name="DNS Server"/>
  </Target>
</Relationship>
```

Attributes

name="ResourceReferenceName"

The name by which the end-to-end automation domain resource will be known in the end-to-end automation domain. The name specified here can differ from the name used for the ReferencedResource in its first-level automation domain.

The name can have a maximum of 64 characters. Do not use more than one blank to separate strings within the name. Duplicate blanks will be ignored.

Contexts

In ResourceReference specifications:

Children: DesiredState, Description, Owner, InfoLink, ReferencedResource

Parents: AutomationPolicy

In ResourceGroup or ChoiceGroup specifications:

Children: None

Parents: Members

In Relationship specifications:

Children: None

Parents: Source, Target

Source

The Source element is a container element. Its subelement ResourceReference or ResourceGroup specifies the dependent resource (ResourceReference or ResourceGroup) in a relationship (element Relationship).

Restrictions

Example

```
<Relationship>
  <Source>
    <ResourceReference name="WebSphere EE"/>
  </Source>
  <Type> StartAfter</Type>
  <Target>
    <ResourceReference name="DNS Server"/>
  </Target>
</Relationship>
```

Attributes

None

Contexts

Children: ResourceReference, ResourceGroup, ChoiceGroup

Parents: Relationship

Target

The Target element is used to specify the supporting resource (ResourceReference or ResourceGroup) in a relationship (element Relationship).

Restrictions

Example

```
<Relationship>
  <Source>
    <ResourceReference name="WebSphere EE"/>
  </Source>
  <Type> StartAfter</Type>
  <Target>
    <ResourceReference name="DNS Server"/>
  </Target>
</Relationship>
```

Attributes

None

Contexts

Children: ResourceReference, ResourceGroup, ChoiceGroup

Parents: Relationship

Type

The Type element is used to specify the relationship between two resources (the Source and the Target). The relationship defines in which sequence the resources must be started or stopped. The automation manager will interpret a relationship definition as automation rule and will ensure that the resources are started or stopped in the correct sequence.

Possible values:

StartAfter

Defines in which sequence two dependent resources must be started. For a relationship *Resource A StartAfter Resource B* this means that when resource A is to be started, Resource B is started first. Only after Resource B has become online will Resource A be started.

StopAfter

Defines in which sequence two resources must be stopped. For a relationship *Resource A StopAfter Resource B* this means that when resource A has to be stopped, Resource B is stopped first. Only after Resource B has stopped will Resource A be stopped.

ForcedDownBy

Defines that a resource must be forced down when another resource is stopped.

For a relationship *Resource A ForcedDownBy Resource B* this means that resource A must be forced offline if Resource B goes offline unexpectedly or is forced down itself. Resource A and Resource B may stop in parallel. The force down of Resource A will be triggered in the following situations:

- When Resource B enters in any of the regular down states (Offline) after having previously been in an online state.
- When Resource B enters any of the terminal down states (Failed offline), regardless of its former state.

Restrictions

The element must occur once per Relationship element.

Example

```
<Relationship>
  <Source>
    <ResourceReference name="WebSphere EE"/>
  </Source>
  <Type> StartAfter</Type>
  <Target>
    <ResourceReference name="DNS Server"/>
  </Target>
</Relationship>
```

Attributes

None

Contexts

Children: None

Parents: Relationship

Policy rule reference

When you define an end-to-end automation policy, you must ensure that a number of rules are observed. The compliance with these rules is checked automatically when you activate a policy. If one of the rules is violated, the policy checking tool issues a warning or an error message. Policies for which only warnings are generated can be activated, errors prohibit the activation of a policy.

Rules for resource groups and choice groups

An error message is issued if one of the following rules is not observed:

- Only resource references that are defined in the policy can be added as members to a group.
- A choice group or a resource group cannot belong to more than one group.
- A resource reference must not be a member of more than one choice group.
- A resource reference must not be a member of both a resource group and a choice group.
- A member must not occur in a group more than once.

A warning message is issued if one of the following rules is not observed:

- A choice group must not be a member of a resource group.
- A member of a choice group must not be the source or the target of a relationship.

Relationship rules

An error message is issued if one of the following rules is not observed:

- Source and target resources of a relationship must have been defined as resource references.
- A resource reference must not be the first source and the last target of a relationship of the same type.

A warning message is issued if one of the following rules is not observed:

- A relationship of type StartAfter and a relationship of StopAfter must not have the same source and target.
- A member of a choice group must not be the source or the target of a relationship.

Rules that ensure that resource names are unique

An error message is issued if one of the following rules is not observed:

- Each resource reference in a policy must have a unique name.
- The name of a resource reference must not be used for a resource group or choice group.
- Each resource group in a policy must have a unique name.
- The name of a resource group must not be used for a resource reference or choice group.
- Each choice group in a policy must have a unique name.
- The name of a choice group must not be used for a resource reference or resource group.

Additional rules

An error message is issued if one of the following rules is not observed:

- A first-level automation resource can only be referenced by one resource reference.
- Each choice group or resource group must have at least one member.
- At least one resource reference must be defined in a policy.
- For one member of a choice group the attribute preferred must be set to true.

Sample automation policy

This is the sample automation policy `sample.xml`. You find it in your policy pool directory.

```
<?xml version="1.0" encoding="UTF-8"?>

<AutomationPolicy version="1.0"
  xmlns="http://www.ibm.com/TSA/Policy.xsd"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.ibm.com/TSA/Policy.xsd EEZPolicy.xsd ">
  <PolicyInformation>
    <PolicyName>Sample E2E Policy</PolicyName>
    <AutomationDomainName>FriendlyE2E</AutomationDomainName>
    <PolicyToken> 1.0.1 </PolicyToken>
    <PolicyAuthor>Michael Atkins</PolicyAuthor>
    <PolicyDescription>Sample policy for the domain FriendlyE2E</PolicyDescription>
  </PolicyInformation>

  <!-- ***** -->
  <!-- Define all resource references first. -->
  <!-- Gather information about referenced -->
  <!-- resources first. See User's Guide for -->
  <!-- instructions. -->
  <!-- ***** -->

  <ResourceReference name="Enterprise DB2">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name Enterprise DB2 </Description>
    <Owner>Bob Owens</Owner>
    <InfoLink>http://www.example.com/help/DB2</InfoLink>
    <ReferencedResource>
      <AutomationDomain>FEPLEX2</AutomationDomain>
      <Name>DB2</Name>
      <Class>ResourceGroup</Class>
      <Node>node1</Node>
    </ReferencedResource>
  </ResourceReference>

  <ResourceReference name="DB2 Connect Client">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name DB2 Connect Client </Description>
    <Owner>Bob Owens</Owner>
    <InfoLink>http://www.example.com/help/DB/Client</InfoLink>
    <ReferencedResource>
      <AutomationDomain>FECluster</AutomationDomain>
      <Name>DB2 Connect</Name>
      <Class>ResourceGroup</Class>
      <Node>node2</Node>
    </ReferencedResource>
  </ResourceReference>

  <ResourceReference name="WebSphere EE">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name WebSphere EE </Description>
    <Owner>Harry Smith</Owner>
    <InfoLink>http://www.example.com/WebSphere/help</InfoLink>
    <ReferencedResource>
```

```

        <AutomationDomain>FEPLEX2</AutomationDomain>
        <Name>WAS</Name>
        <Class>ResourceGroup</Class>
        <Node>node1</Node>
    </ReferencedResource>
</ResourceReference>

<ResourceReference name="DNS Server">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name DNS Server </Description>
    <Owner>Jenny McDowell</Owner>
    <InfoLink>http://www.example.com/help/DNS Server</InfoLink>
    <ReferencedResource>
        <AutomationDomain>FECluster</AutomationDomain>
        <Name>DNS Server</Name>
        <Class>ResourceGroup</Class>
        <Node>node2</Node>
    </ReferencedResource>
</ResourceReference>

<ResourceReference name="NFS Client">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name NFS Client </Description>
    <Owner>Henry Smith</Owner>
    <InfoLink>http://www.example.com/help/NFS/client/minor_01/?help+NFS+Client</InfoLink>
    <ReferencedResource>
        <AutomationDomain>cm1cm2</AutomationDomain>
        <Name>nfs-client</Name>
        <Class>IBM.ResourceGroup</Class>
        <Node>1nxc2</Node>
    </ReferencedResource>
</ResourceReference>

<ResourceReference name="NFS Server">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name "NFS Server </Description>
    <Owner>Bob Owens</Owner>
    <InfoLink>file://X:/help/NFS.pdf</InfoLink>
    <ReferencedResource>
        <AutomationDomain>FEPLEX1</AutomationDomain>
        <Name>NFS Server</Name>
        <Class>ResourceGroup</Class>
        <Node>node3</Node>
    </ReferencedResource>
</ResourceReference>

<ResourceReference name="WebSphere AE">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name WebSphere AE </Description>
    <Owner>Harry Smith</Owner>
    <InfoLink>http://www.example.com/help/WebSphere.html</InfoLink>
    <ReferencedResource>
        <AutomationDomain>FECluster</AutomationDomain>
        <Name>WebSphere AppServer</Name>
        <Class>ResourceGroup</Class>
        <Node>node2</Node>
    </ReferencedResource>
</ResourceReference>

<ResourceReference name="eMail Hosting">
    <DesiredState>Online</DesiredState>
    <Description>This is the reference with the name eMail Hosting </Description>
    <Owner>Sarah O'Neill</Owner>
    <InfoLink>http://www.example.com/help/policy/email.html</InfoLink>
    <ReferencedResource>
        <AutomationDomain>FECluster</AutomationDomain>
        <Name>Mail Server</Name>

```

```

        <Class>IBM.Application</Class>
        <Node>node2</Node>
    </ReferencedResource>
</ResourceReference>

<ResourceReference name="HTTP Server Prim">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name HTTP Server </Description>
    <Owner>Owens, Bob</Owner>
    <InfoLink>http://www.example.com/help/policy/HTTPServer.html</InfoLink>
    <ReferencedResource>
        <AutomationDomain>FEClusterSAP</AutomationDomain>
        <Name>HTTP Server</Name>
        <Class>ResourceGroup</Class>
        <Node>node4</Node>
    </ReferencedResource>
</ResourceReference>

<ResourceReference name="HTTP Server Backup">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name HTTP Server </Description>
    <Owner>Harry O'Neill</Owner>
    <InfoLink>--- No help ---</InfoLink>
    <ReferencedResource>
        <AutomationDomain>FECluster</AutomationDomain>
        <Name>HTTP Server Back</Name>
        <Class>ResourceGroup</Class>
        <Node>nodebackup</Node>
    </ReferencedResource>
</ResourceReference>

<ResourceReference name="SAP ENQ Server">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name SAP ENQ Server </Description>
    <Owner>Bob Owens</Owner>
    <InfoLink>http://www.example.com/help/policy/enq.html</InfoLink>
    <ReferencedResource>
        <AutomationDomain>FEPLEX1</AutomationDomain>
        <Name>SAP ENQ Env.</Name>
        <Class>IBM.Application</Class>
        <Node>node3</Node>
    </ReferencedResource>
</ResourceReference>

<ResourceReference name="SAP AppServer">
    <DesiredState>Offline</DesiredState>
    <Description>This is the reference with the name SAP AppServer </Description>
    <Owner>Jenny Parker</Owner>
    <InfoLink>http://www.example.com/help/policy/SAPAPP.html</InfoLink>
    <ReferencedResource>
        <AutomationDomain>FEClusterSAP</AutomationDomain>
        <Name>SAP AppServer</Name>
        <Class>ResourceGroup</Class>
        <Node>node4</Node>
    </ReferencedResource>
</ResourceReference>

<!-- ***** -->
<!-- Resource references can be grouped -->
<!-- together. -->
<!-- See User's Guide for instructions -->
<!-- ***** -->

<ChoiceGroup name="HTTP Server">
    <DesiredState> Offline </DesiredState>
    <Description>This is the choice-Group for choosing one HTTP Server</Description>
    <Owner>Jenny Parker</Owner>

```

```

    <InfoLink>http://www.example.com/choice</InfoLink>
    <Members>
    <ResourceReference name="HTTP Server Prim" preferred="true"/>
    <ResourceReference name="HTTP Server Backup"/>
  </Members>
</ChoiceGroup>

  <ResourceGroup name="WebSphere E2E" >
    <DesiredState> Offline </DesiredState>
    <Description>This is collection group with the name WebSphere E2E</Description>
    <Owner>Harry Smith</Owner>
    <InfoLink>http://www.example.com/help/policy/WebSphere.html</InfoLink>
    <Members>
      <ResourceReference name="Enterprise DB2"/>
      <ResourceReference name="DB2 Connect Client"/>
      <ResourceReference name="WebSphere EE"/>
      <ResourceReference name="DNS Server"/>
      <ResourceReference name="NFS Client"/>
      <ResourceReference name="NFS Server"/>
      <ResourceReference name="WebSphere AE"/>
    </Members>
  </ResourceGroup>

  <ResourceGroup name="mySAP Solutions" >
    <DesiredState> Offline </DesiredState>
    <Description>This is collection group with the name mySAP Solutions</Description>
    <Owner>Jerry Owens</Owner>
    <InfoLink>http://www.example.com/help/policy/mysap/Solution.html</InfoLink>
    <Members>
      <ResourceReference name="DNS Server"/>
      <ResourceReference name="SAP ENQ Server"/>
      <ResourceReference name="SAP AppServer"/>
    </Members>
  </ResourceGroup>

  <ResourceGroup name="Friendly Computer Shop" >
    <DesiredState>Offline</DesiredState>
    <Description>This is collection group with the name Friendly Computer Shop</Description>
    <Owner>Jerr Owens</Owner>
    <InfoLink>http://www.example.com/help/policy/compshop.html</InfoLink>
    <Members>
      <ResourceGroup name="mySAP Solutions"/>
      <ResourceReference name="WebSphere AE"/>
    </Members>
  </ResourceGroup>

<!-- ***** -->
<!-- Define relationships of resources -->
<!-- to each other -->
<!-- See User's Guide for instructions -->
<!-- ***** -->

  <Relationship>
    <Source>
      <ResourceReference name="DB2 Connect Client"/>
    </Source>
    <Type> StartAfter</Type>
    <Target>
      <ResourceReference name="Enterprise DB2"/>
    </Target>
  </Relationship>

  <Relationship>
    <Source>
      <ResourceReference name="WebSphere EE"/>
    </Source>
    <Type> StartAfter</Type>

```

```
        <Target>
          <ResourceReference name="DNS Server"/>
        </Target>
      </Relationship>

      <Relationship>
        <Source>
          <ResourceReference name="NFS Client"/>
        </Source>
        <Type> StartAfter</Type>
        <Target>
          <ResourceReference name="NFS Server"/>
        </Target>
      </Relationship>

      <Relationship>
        <Source>
          <ResourceReference name="SAP AppServer"/>
        </Source>
        <Type> StartAfter</Type>
        <Target>
          <ResourceReference name="SAP ENQ Server"/>
        </Target>
      </Relationship>
    </AutomationPolicy>
```

Appendix D. Additional state-related information about end-to-end automation resources

When you select an end-to-end automation resource in the resource table on the operations console, the Additional info page provides information about the internal states of the automated resource. The automation engine uses these internal states to calculate the common states that are displayed on the General page.

The internal states on the Additional Info page can help an experienced operator and IBM support to understand why a specific common state is shown on the General page. The information can be useful for finding the root cause of problems that the selected resource has encountered.

The internal states that appear on the Additional info page are described in the following sections. Note that these internal states are subject to change and are listed here for reference purposes only.

This is an example of an Additional info page:

Information area

| Name | Value |
|----------------------|----------------|
| sa_status_compound | Satisfactory |
| sa_status_desired | Offline |
| sa_status_observed | Offline |
| sa_status_binding | Not Applicable |
| sa_status_automation | Idle |
| sa_status_control | Startable |
| sa_status_health | Not Applicable |

Seite 1 von 1 Summe: 7 Gefiltert: 7 Angezeigt: 7

Figure 22. Additional info page for a resource reference

The following tables lists the internal automation states.

Table 49. Internal automation states on the Additional info page

| State | Description |
|----------------------|--|
| sa_status_compound | Indicates if automation is being performed successfully or if a problem has occurred. For details, refer to “Values for sa_status_compound.” |
| sa_status_desired | Indicates whether the desired state of the resource is Online or Offline. For details, refer to “Values for sa_status_desired” on page 289. |
| sa_status_observed | Indicates if the resource is currently running or if a problem has occurred. For details, refer to “Values for sa_status_observed” on page 289. |
| sa_status_binding | Indicates if the resource has been selected to run on a specific system. For end-to-end automation resources, this value is not used. It is always set to Not Applicable. |
| sa_status_automation | Indicates if the end-to-end automation manager is automating successfully or has encountered problems. For details, refer to “Values for sa_status_automation” on page 290. |
| sa_status_control | Provides additional information about sa_status_automation. For details, refer to “Values for sa_status_control” on page 290. |
| sa_status_health | Indicates if the resource has a performance problem. For end-to-end automation resources, this value is not used. It is always set to Not Applicable. |

Values for sa_status_compound

This internal automation engine state can have the values that are listed in the following table.

Table 50. Values for sa_status_compound

| sa_status_compound | Description |
|--------------------|--|
| Null | The state is currently undefined. |
| Problem | Some problem occurred calculating this value. |
| Denied | This state appears for SA z/OS resources only. It indicates that the resource cannot be automated. |
| Inhibited | A resource on which this resource has dependencies is not in a satisfactory state. |

Table 50. Values for `sa_status_compound` (continued)

| <code>sa_status_compound</code> | Description |
|---------------------------------|--|
| Awaiting Automation | Automation is in progress. An automation action has been triggered. The automation engine is waiting for a response. |
| InAutomation | Automation is in progress. |
| Degraded | In end-to-end automation: An interim state for groups that are being automated. |
| Satisfactory | Requested state matches the current state. |

Values for `sa_status_desired`

This internal automation engine state can have the values that are listed in the following table.

Table 51. Values for `sa_status_desired`

| <code>sa_status_desired</code> | Description |
|--------------------------------|---|
| Null | The state is currently undefined. |
| Online | The current desired state of the resource is Online. |
| Offline | The current desired state of the resource is Offline. |

Values for `sa_status_observed`

This internal automation engine state can have the values that are listed in the following table.

Table 52. Values for `sa_status_observed`

| <code>sa_status_observed</code> | Description |
|---------------------------------|---|
| Null | The state is currently undefined. |
| Unknown | The state has not been detected yet. The targeted first-level automation domain may not be online or the initial event that supplies the state information for the referenced resource has not yet been received. |
| SystemGone | The connection to the first-level automation domain was lost. In end-to-end automation, SystemGone is interpreted to mean that the first-level automation domain is gone. |
| Offline | The resource is currently offline. |
| Failed | The resource cannot be started. |
| PendingOnline | The resource is starting. |
| Online | The resource is currently online. |
| StuckOnline | The resource cannot be stopped. |
| PendingOffline | The resource is stopping. |
| WasOnline | The connection to the resource was lost. |

Table 52. Values for sa_status_observed (continued)

| sa_status_observed | Description |
|--------------------|--|
| Ineligible | The resource is in standby mode. |
| InvalidReference | The referenced resource cannot be found in the targeted first-level automation domain. |

Values for sa_status_automation

This internal automation engine state can have the values that are listed in the following table.

Table 53. Values for sa_status_automation

| sa_status_automation | Description |
|----------------------|---|
| Null | The state is currently undefined. |
| Unknown | The state has not been detected yet. The targeted first-level automation domain may not be online or the initial event that supplies the state information for the referenced resource has not yet been received. |
| SystemGone | The connection to the first-level automation domain was lost. In end-to-end automation, SystemGone is interpreted as First-level automation domain is gone. |
| Idle | No automation action is currently in progress for the selected resource. |
| Ordered | A request has been generated. If the selected resource is a resource reference, the request is being dispatched to the first-level automation domain. |
| Busy | The request has been accepted by the targeted resource. Automation is in progress. |
| Denied | This state appears for SA z/OS resources only. It indicates that the resource cannot be automated. |
| Problem | Automation has encountered a problem. |
| Internal | Automation is in progress. |

Values for sa_status_control

This internal automation engine state can have the values that are listed in the following table.

Table 54. Values for sa_status_control

| sa_status_control | Description |
|-------------------|--|
| Null | The state is currently undefined. |
| Startable | The resource can be started. |
| NotStartable | Currently, the resource cannot be started. |
| StartInhibited | A dependency on another resource prevents the resource from being started. |

Table 54. Values for *sa_status_control* (continued)

| sa_status_control | Description |
|--------------------------------|--|
| Denied | This state appears for SA z/OS resources only. It indicates that the resource cannot be automated. |
| StopInhibited | A dependency on another resource prevents the resource from being stopped. |
| NotControllable | The automation engine cannot control the resource. |
| NodeExcluded | The system or first-level automation domain that hosts this resource or referenced resource is currently excluded from automation. |
| NodeExcludedAndNotControllable | See the descriptions for NodeExcluded and NotControllable above. |
| BrokenReference | The end-to-end automation manager cannot contact the referenced resource in the first-level automation domain because an exception, for example, a security exception, has occurred. |
| ReferenceInError | The end-to-end automation resource is in an error state because the referenced resource has encountered an error. The error is not due to a problem within the end-to-end automation engine. |
| NonRecoverableError | The resource has encountered a problem that must be resolved by an operator. When the problem is resolved, a manual reset may be required. |

Appendix E. Troubleshooting

Where to find the log and trace files

This section describes where you find the log and trace files that are relevant for end-to-end automation management.

Where to find the Tivoli Common Directory

Message and trace logs for Tivoli products are located under a common parent called the Tivoli Common Directory. The log and trace files of all subcomponents of SA for Multiplatforms that are not running within WebSphere Application Server, for example, the log and trace files of the end-to-end automation engine and of the automation adapters, are written to the product-specific subdirectory of the Tivoli Common Directory.

The path to the Tivoli Common Directory is specified in the properties file `log.properties`. The file `log.properties` is located in the following directory:

- **Windows:** `C:\Program Files\IBM\tivoli\common\cfg`
- **AIX/Linux:** `/etc/ibm/tivoli/common/cfg`

In the `log.properties` file, the path to the Tivoli Common Directory is defined in the property `tivoli_common_dir=<path_to_Tivoli_Common_Directory>`.

These are the default values:

- For **Windows** systems: `C:/Program Files/IBM/tivoli/common`
Note that forward slashes are used as path delimiters in this properties file.
- For **AIX** and **Linux** systems:
`/var/ibm/tivoli/common`

These are the relevant subdirectories for end-to-end automation management:

| Subdirectory | Description |
|---|--------------------------------|
| <code><Tivoli_Common_Directory>/eez/logs</code> | message log files, trace files |
| <code><Tivoli_Common_Directory>/eez/ffdc</code> | FFDC files |

For additional information on where to find the log and trace files of the automation engine, see below. For information about the log and trace files of the automation adapters, refer to the adapter-specific documentation.

Log and trace files of the automation engine

The log files and trace files of the automation engine are available in the directory `<Tivoli_Common_Directory>/eez/logs`.

Message log file: `<Tivoli_Common_Directory>/eez/logs/msgengine.log`

This is the domain log file of the end-to-end automation domain that can be displayed from the operations console.

Trace log file: `<Tivoli_Common_Directory>/eez/logs/traceengine.log`

Which messages and traces are written to the files is specified on the Logger page of the configuration dialog. For information about the configuration dialog, refer to Chapter 29, “Configuring the end-to-end automation manager,” on page 193. For a detailed description of the properties that can be configured on the page, refer to the configuration dialog help.

Viewing the XML log file of the automation engine

The log and trace files are written in XML format. Because the XML files may be difficult to read, you can use a tool that converts the XML file to HTML format and view the HTML file instead of the XML source file. This section describes how to use the tool.

Notes:

1. Typically, you will display and browse the log file of the end-to-end automation engine by selecting the end-to-end automation domain in the topology tree on the operations console and clicking **View log** on the General page. Only when you cannot access the log file from the operations console, for example, because the automation engine does not start, should you proceed as described in this section.
2. The trace files are intended for use by IBM support only.

You find the tool in the directory <EEZ_INSTALL_ROOT>/install. There, look for the file logviewer214_basics.zip.

Prerequisites for using the tool:

- A tool for unzipping the file (not included in the Tivoli System Automation for Multiplatform 2.1 package)
- J2SE 1.4.x (included in the WebSphere Application Server 6 installation)

After unzipping the file, refer to the file readme.html for further installation instructions and for information about the features of the formatting tool.

After you have installed the tool, you can use the following scripts to convert the log and trace files to HTML and display them in a Web browser:

- **Windows:** viewer.bat
- **AIX/Linux:** viewer.sh

As described in the readme.html, the viewer script takes a so-called query string to format the HTML output. This is an example of such a query string:

```
select Time,SourceFile, SourceMethod,MessageId,LogText,Exception,Thread
       where (ProductId=SAMP)
```

It is recommended that you save the query string in a plain text file (for example, with the name stdtrace). To invoke the viewer script, use the following command:

```
viewer -f stdtrace traceengine.xml > traceengine.html
```

Log and trace files of the automation J2EE framework and the resource adapters

The automation J2EE framework of the end-to-end automation management component, the automation engine resource adapter, and the first-level automation manager resource adapter use the log files and the tracing function of WebSphere Application Server.

By default, the information is written to these log and trace files:

- SystemOut.log

- SystemErr.log
- trace.log

The files are located in the following directory:

<was_root>/profiles/<profile_name>/logs/<server_name>

where <profile_name> is the name of the profile of the server where the automation J2EE framework is installed. The default profile name is "default".

You use the WebSphere Application Server administrative console to set the parameters for logging and tracing:

- To specify log file parameters, for example, the log file names, the maximum size, and the number of history log files to be preserved, open the administrative console and navigate to **Troubleshooting** → **Logs and Trace** → <server_name> → **JVM Logs**
- To set the parameters for tracing, for example, to switch tracing on or off or to define for which components traces should be recorded, open the administrative console and navigate to **Troubleshooting** → **Logs and Trace** → <server_name> → **Diagnostic trace**.

For more information, refer to the information center for WebSphere Application Server, Version 6.0. You find the information center at the following location:

<http://publib.boulder.ibm.com/infocenter/wasinfo/v6r0/>

Trace files of the operations console

The operation console utilizes the tracing function of WebSphere Portal Server.

Two different types of traces are recorded:

- Integrated Solutions Console trace

The trace files can be found in the following directory:

<isc_runtime_root>/PortalServer/log

The following trace files are relevant:

```
wps_<date_time_stamp>.log
SystemOut.log
SystemErr.log
```

The trace level can be changed permanently by editing the following file:

<isc_runtime_root>/PortalServer/shared/app/config/log.properties

To switch tracing on, replace the line #traceString==all=disabled with traceString=org.apache.jetspeed.portlet.PortletLog=all=enabled.

Temporary tracing can be changed while Integrated Solutions Console is running. To do so, log in to Integrated Solutions Console and click **Settings** → **Enable Tracing** and append these trace settings:

```
org.apache.jetspeed.portlet.PortletLog=all=enabled
```

The temporary setting will be lost when Integrated Solutions Console is stopped or restarted.

- The following trace file is important for analyzing runtime problems:

<was_root>/profiles/<profile_name>/csa/logs/WnServlet_0.log

How to determine the server port number for connecting to the operations console

When WebSphere Application Server is set up with ports other than the default ports, you may have to specify a port number other than the default port number 2809 on the Connect panel that appears when you have logged on to Integrated Solutions Console and want to connect to the operations console of Tivoli SA for Multiplatforms.

To find the correct port number, open the administrative console of WebSphere Application Server and navigate to **Servers** —> **Application Servers**—> server1.

You find the correct port number in the Communications section:
Ports —> variable `BOOTSTRAP_ADDRESS`.

Problems occur when multiple browser windows are used to connect to the same Integrated Solutions Console from the same client system

If you are using a browser other than Microsoft Internet Explorer, opening multiple browser windows on the same client machine to connect to the same Integrated Solutions Console will cause unexpected results. This is because only Microsoft Internet Explorer establishes a separate HTTP session for each browser instance. Other browser types will share a single session between multiple browser instances on the same system if these instances connect to the same Integrated Solutions Console.

The same situation occurs if you open multiple Microsoft Internet Explorer browser windows using **File** —> **New Window** (or Ctrl+N) from an existing Integrated Solutions Console session, because in this case the new browser window and the one from which it was opened will also share the same session.

The end-to-end automation domain is not displayed on the operations console

If the end-to-end automation domain is not displayed on the operations console although the automation J2EE framework is running and the automation engine is started, perform the following steps:

1. In the end-to-end automation management configuration dialog, verify that all parameters are set correctly.
2. Restart the automation engine.

For information on the configuration dialog, refer to Chapter 29, “Configuring the end-to-end automation manager,” on page 193. For information on starting the automation engine, refer to Chapter 27, “Using the command-line interface of the automation engine,” on page 183.

A System Automation for Multiplatforms domain is not displayed in the topology tree

If a first-level automation domain does not appear in the topology tree on the operations console, perform the following steps to analyze and resolve the problem:

1. Check if the adapter is running by issuing the following command on one of the nodes of the domain:

```
samadapter status
```

If the adapter is running, a message like in the following example comes up:
samadapter is running on sapb13

If the adapter is automated, a message like in the following example comes up:

```
Automated ResourceGroup 'samadapter-rg' runs on sapb13
```

Make a note of the name of the node on which the adapter runs (in the example this is sapb13) and proceed with step 4.

2. If the adapter is not running, issue the following command to check if the domain is online:

```
lsrpdomain
```

A message like in the following example comes up:

| Name | OpState | RSCTActiveVersion | MixedVersions | TSPort | GSPort |
|---------|---------|-------------------|---------------|--------|--------|
| domain1 | Online | 2.4.4.2 | No | 12347 | 12348 |

If OpState is not Online, start the domain.

3. If the domain is online, start the adapter with the following command:

```
samadapter start
```

After the start message has appeared, reissue the following command:

```
samadapter status
```

4. If the adapter is running, check again on the operations console if the domain now appears in the topology tree. Note that it may take time until the contact to the end-to-end automation manager is established after the adapter is started.
-

5. If the domain still does not appear in the topology tree, you need the connection information that you specified in the adapter configuration dialog to resolve the problem.

Perform the following steps:

- a. Launch the adapter configuration dialog of SA for Multiplatforms by issuing the following command on a node in the domain:

```
cfgsamadapter
```

- b. On the entry panel of the configuration dialog, click **Configure**.
-

- c. Open the Adapter page on the Configure panel and write down the values that appear in the following fields:

- **Host name or IP Address**
- **Request port number**

This is the connection information the end-to-end automation management host uses to reach the adapter on any of the nodes in the domain.

- d. Open the page Host using adapter and write down the values that appear in the following fields:
- **Host name or IP Address**
 - **Event port number**

This is the connection information the adapter on any of the nodes in the domain uses to reach the end-to-end automation management host.

-
6. Check if end-to-end automation management can be reached from each node in the domain. A simple test is `ping <end-to-end management host>`. If there is a firewall between the nodes of the domain and the end-to-end automation management host, check with the network administrator if the firewall permits a connection between the node (page Adapter: **Host name or IP Address**) and the end-to-end management host (page Host using adapter: **Host name or IP Address** and **Event port number**).

-
7. The adapter determines whether SSL must be used for the communication with the end-to-end automation manager. To check the SSL settings of the adapter, launch the adapter configuration dialog using the command `cfgsamadapter`. On the Security page, verify that the SSL settings are correct.

Note: If the end-to-end automation manager is configured for using SSL, the adapter must be configured for SSL as well. The SSL configuration of the end-to-end automation manager is done on the WebSphere Application Server administrative console.

-
8. On the end-to-end automation management host, use `netstat` to find out if it is listening for events on the event port defined in **Event port number**. When the event port number is set to 2002 on a Windows host, `netstat` brings up a message like in the following example:

```
C:\>netstat
Active Connections
  Proto Local Address           Foreign Address         State
  ...
  TCP    E2EHOST:2002          sapb13.boeblingen.de.ibm.com:45688 ESTABLISHED
  ...
```

If `netstat` does not display any information about the event port defined in **Event port number**, open the file `/etc/hosts` (on Windows the file is located in `C:\WINDOWS\system32\drivers\etc\hosts`) and verify that the loopback address (`127.0.0.1`) is not related to the actual host name. The loopback address should be related to `localhost` only.

For example, the entry in `/etc/hosts` may look like the following:

```
127.0.0.1          localhost.localdomain localhost
```

-
9. Check if each node in the domain can be reached from end-to-end automation management. A simple test is `ping <hostname or IP Address>`. If there is a firewall between the end-to-end automation management host and the nodes of the domain, check with the network administrator if the firewall permits a connection between the end-to-end automation management host (page Host using adapter: **Host name or IP Address** and **Request port number**) and the node (page Adapter: **Host name or IP Address**).

-
10. On the node on which the adapter is running, use `netstat` to find out if it is listening on the port defined in **Request port number**.

For example, when the Request port number is set to 2001, **netstat** brings up a message like this on AIX and Linux hosts:

```
sapb13:~ # netstat -atn |grep 2001
tcp      0      0 9.152.20.113:2001      :::*          LISTEN
```

11. When the communication between all ports has been established correctly (see the descriptions above), check whether the EEZ Publisher is running. The EEZ Publisher must be running on the master node of the Base component of SA for Multiplatforms.

To check if the Publisher is running, perform the following steps:

- a. Issue the following command on one of the nodes of the first-level automation domain:

```
- issue lssamctrl
```

If the Publisher is enabled, you will receive output like in the following example:

```
safli03:~ # lssamctrl | grep Publisher
EnablePublisher      = EEZ
```

- b. Issue the following command on the master node of the Base component of SA for Multiplatforms:

```
ps ax
```

You should receive output like in the following example:

```
safli04:~ # ps ax | grep Publisher
25756 ?    S  0:00
                TECPublisher /etc/opt/IBM/tsamp/sam/cfg/EEZPublisher.conf EEZ
25757 ?    S  0:00
                TECPublisher /etc/opt/IBM/tsamp/sam/cfg/EEZPublisher.conf EEZ
25758 ?    S  0:00
                TECPublisher /etc/opt/IBM/tsamp/sam/cfg/EEZPublisher.conf EEZ
25759 ?    S  0:00
                TECPublisher /etc/opt/IBM/tsamp/sam/cfg/EEZPublisher.conf EEZ
```

- c. Issue the following command on the SA for Multiplatforms node on which the adapter is running:

```
netstat
```

You should receive output like in the following example:

```
afli03:~ # netstat -atn | grep 5539
tcp      0      0 :::5539      :::*          LISTEN
tcp      0      0 9.152.21.82:5539  9.152.20.92:32793 ESTABLISHED
```

If the Publisher is not running or communication on port 5539 cannot be established, perform the following steps:

- a. Check that the file `/etc/Tivoli/tec/samPublisher.conf` contains the following entry:

```
#--SAMP-EEZ:
Publisher=EEZ
LibraryPath=libTECPublisher.so
ConfigPath=/etc/opt/IBM/tsamp/sam/cfg/EEZPublisher.conf
```

- b. Check that the file `/etc/opt/IBM/tsamp/sam/cfg/EEZPublisher.conf` contains the following entries:

```
ServerLocation=adapter_ip_address
ServerPort=5539
```

The value specified for `adapter_ip_address` in the file must match the value provided on the Adapter page of the SA for Multiplatforms adapter configuration dialog.

12. If the domain still does not appear on the operations console, contact IBM support and provide diagnostic information:
 - a. On each node in the domain, find out where the trace files are located. The trace files can be found in the /eez/logs subdirectory of the Tivoli Common Directory. To find the path to the Tivoli Common Directory, issue the following command:


```
cat /etc/ibm/tivoli/common/cfg/log.properties
```

 The command returns the path to the Tivoli Common Directory, for example:


```
Tivoli_common_dir=/var/ibm/tivoli/common
```

 This means that the trace files can be found in the following directory:


```
/var/ibm/tivoli/common/eez/logs
```
 - b. Use tar to package all files in the directory and provide the archive to IBM support.

Critical exceptions in the WebSphere Application Server log file

If the end-to-end automation management component cannot be accessed from the operations console although the WebSphere Application Server is running, or if the domain topology in the operations console does not look like expected, check the WebSphere Application Server log file for one or multiple of the following exceptions or stack trace fragments:

```
java.lang.IllegalMonitorStateException: JVMLK002: current thread not owner
```

```
CNTR0019E: EJB threw an unexpected (non-declared) exception during invocation
of method "findByPrimaryKey". Exception data: java.lang.NullPointerException
at
com.ibm.ejs.container.activator.UncachedActivationStrategy.atActivate(
UncachedActivationStrategy.java(Compiled Code))
[...]
at com.ibm.eez.aab.subscription.EJSLocalCMPEEZDomainSubscriptionHome_25634d48.findByPrimaryKey(
EJSLocalCMPEEZDomainSubscriptionHome_25634d48.java(Compiled Code))
at com.ibm.eez.aab.EEZDomainSessionBean.unsubscribeAll(EEZDomainSessionBean.java(Compiled Code))
```

```
CNTR0019E: EJB threw an unexpected (non-declared) exception during invocation
of method "findByPrimaryKey". Exception data:
com.ibm.websphere.cpi.CPIException: ; nested exception is:
java.lang.ClassCastException: com.ibm.eez.aab.EEZDomainSessionBean
[...]
at com.ibm.eez.aab.subscription.EJSCMPEEZDomainSubscriptionHomeBean_25634d48.findByPrimaryKey_Local(
EJSCMPEEZDomainSubscriptionHomeBean_25634d48.java(Inlined Compiled Code))
[...]
at com.ibm.eez.aab.EEZDomainSessionBean.unsubscribeAll(EEZDomainSessionBean.java(Compiled Code))
```

To resolve the problem, do this:

1. Disable the just-in-time compiler (JIT) of the WebSphere Java Virtual Machine (JVM)
2. Restart WebSphere Application Server

If the domain topology still does not look like expected, deactivate the end-to-end automation policy and activate it again.

Security exception when trying to subscribe to resources that are hosted on a first-level automation domain

If you see the following error messages in the domain log file of the end-to-end automation domain, verify that the credentials for the first-level automation domain have been specified:

EEZD0069E

A Security Exception was caught trying to subscribe to resources hosted on automation domain with name first-level domain. Following is a list of resources the automation engine tried to subscribe to: (resource_group/IBM.ResourceGroup/).

EEZD0072E

An EEZUserSecurityException was caught trying to contact another automation domain. Original message text is: EEZA0009E Invocation of adapter plug-in failed:
plug-in=com.ibm.sam.eezplugin.SAMFLA, method=SUBSCRIBE_RESOURCE, internalRetcode=41, taskRetcode=0.

To check that the user credentials for the first level automation domain have been specified correctly, check the relevant settings on the User credentials page of the configuration dialog.

For information on how to launch the configuration dialog, refer to “Launching the configuration dialog” on page 193. For detailed information about the User credentials page, refer to the online help of the configuration dialog.

Resolving timeout problems

If you experience timeout problems when accessing first-level automation domains, this may mean that the default values of some optional J2EE framework environment variables are not appropriate for your environment.

The following table lists the environment variables that you may need to change to resolve the problems.

More information about the environment variables is provided in the following sections. Section “Modifying the environment variables for the automation J2EE framework” on page 303 describes how you change the environment variables on the administrative console of WebSphere Application Server.

Table 55. Environment variables of the automation J2EE framework

| Variable name | Minimum value | Default value | Maximum value |
|--|---------------|---------------|---------------|
| com.ibm.eez.aab.watchdog-interval-seconds | 60 | 300 | 86400 |
| com.ibm.eez.aab.watchdog-timeout-seconds | 2 | 10 | 60 |
| com.ibm.eez.aab.domain-removal-hours | 1 | 48 | 1000 |
| com.ibm.eez.aab.invocation-timeout-seconds | 30 | 60 | 3600 |

Rules:

- If the value of an environment variable is below the minimum value for that variable, the minimum value is used.
- If the value of an environment variable is above the maximum value for that variable, the maximum value is used.
- Cross-dependency: To ensure that domains are removed only after the health state has moved to some timeout or failed state, the value of the variable **com.ibm.eez.aab.domain-removal-hours** must be greater than the value of **com.ibm.eez.aab.watchdog-interval-seconds/3600**.

If you specify values that violate this rule, the user-specified value for **com.ibm.eez.aab.domain-removal-hours** is ignored and the value of **com.ibm.eez.aab.domain-removal-hours** is set to **com.ibm.eez.aab.watchdog-interval-seconds/3600 +1**.

Watchdog - A mechanism for monitoring the domain communication states

The automation J2EE framework includes a watchdog mechanism to determine the health state of the communication with each domain (either the end-to-end automation domain or a first-level domain). If the automation J2EE framework and the domain in question have not communicated successfully during the time interval defined by the environment variable **com.ibm.eez.aab.watchdog-interval-seconds** (default value: 300), the automation J2EE framework invokes a test operation on the domain. This test operation may only take a limited amount of time, as defined by the environment variable **com.ibm.eez.aab.watchdog-timeout-seconds**. Depending on the outcome of this test operation, the domain communication health state is updated and reflected in the operations console accordingly.

Database clean-up timeout for automation domains

The automation J2EE framework contains a mechanism for removing automation domains from the database after a period of inactivity. The domains themselves are not removed, just the representation of the domains in the automation J2EE framework is removed.

When the automation J2EE framework detects that no communication with a particular domain has occurred for a time interval that is longer than the clean-up timeout interval defined in the environment variable **com.ibm.eez.aab.domain-removal-hours**, it removes the related domain information from the database.

If the automation J2EE framework had been stopped for a time, such domains will be removed only after attempts to contact them have failed.

Whenever the automation J2EE framework removes a domain, the operations console is notified about the change and refreshed accordingly.

Method invocation timeout between the automation J2EE framework and the automation adapters

A timeout value can be set in order to control how long an operation between the automation J2EE framework and the automation adapters may take. The environment variable **com.ibm.eez.aab.invocation-timeout-seconds** is used to define this timeout value.

The value of this environment variable should be at least 15 seconds less than the value of the WebSphere ORB request timeout property. Otherwise,

"CORBA.NO_RESPONSE: Request timed out" errors may be encountered by the operations console or the automation engine if an operation takes longer than the time interval specified by the ORB request timeout. The default value for the WebSphere ORB request timeout is 180 seconds. The ORB request timeout property can be changed on the administrative console of WebSphere Application Server. To view or change the property, open the administrative console and navigate to **Servers** —> **Application Servers** —> **server1** —> **Container Services** —> **ORB service**. See the WebSphere documentation for more information about the ORB request timeout property.

The **com.ibm.eez.aab.invocation-timeout-seconds** variable is used for the communication with all automation adapters. There is no individual timeout value per automation adapter.

Note: The communication with the end-to-end automation engine does not support method invocation timeout. This means that either the connection cannot be established, in which case the operation returns with an exception immediately, or the operation will continue until a connection is established.

Modifying the environment variables for the automation J2EE framework

The current value of each variable is displayed when the application EEZEAR is started. Look for messages EEZJ1004I, EEZJ1005I, EEZJ1006I in the WebSphere Application Server log (SystemOut.log).

If the default values are not appropriate for your environment, you can change the environment variables on the administrative console of WebSphere Application Server.

Perform the following steps:

1. Connect to the administrative console.

2. Click **Servers** —> **Application Servers** —> **server1** —> **Server Infrastructure** —> **Java and Process Management** —> **Process Definition** —> **Additional Properties** —> **Java Virtual Machine Additional Properties** —> **Custom Properties**

3. Click **New** to change the setting of a variable.

4. Enter values for **Name** (com.ibm.eez.aab.<variable_name>) and **Value** (<new_value>). You can also enter a description.

5. Save your changes.

WebSphere Application Server must be restarted for the changes to take effect.

Modifying the time zone settings for the operations console

The times stamps that are displayed on the operations console are derived from the time zone settings of the operating system on the system on which the Integrated Solutions Console server is installed. If the times in the time stamps differ from the local time at your location, check the time zone settings on your Integrated Solutions Console server.

The time settings can usually be set with the configuration tools that are provided with the operating system:

- On AIX, you can configure time settings with the `smit` or `smitty` system configuration tool. Use the menu entries **System environments** → **Change/Show Date and Time** to adjust the time settings.
- On SuSE Linux, you can use the `yast2` or `yast` system configuration tools. Use the menu entries **System** -> **Date and Time** (SLES-9) or **System** → **Set Time Zone** (SLES-8).
- On Red Hat Linux distributions, you can use the configuration tools `redhat-config-time` or `system-config-time`.
- On Windows, you can adjust the time settings on the Control Panel.

You may have to restart your operating system for the changes to take effect.

Note:

AIX, Linux:

If you have modified the time zone settings as described above but the times displayed in the time stamps on the operations console are still inappropriate, you can set the environment variable `TZ` to resolve the problem.

Examples:

- To set the time zone for Berlin, Germany, use the following command:

```
export TZ="Europe/Berlin"
```
- To set the time zone to US Eastern Standard Time, use the following command:

```
export TZ="US/Eastern"
```

Unrecoverable error state displayed for first-level automation resources is incorrect

When the connectivity between the nodes of a cluster is reestablished after a connectivity failure, the operations console may incorrectly indicate that the resources on the nodes of the cluster are in state Unrecoverable error.

This behavior is the result of a cluster split in cases where both subclusters do not terminate themselves (for more information on cluster split situations, refer to *IBM Tivoli System Automation for Multiplatforms Base Component User's Guide*).

To resolve the problem, that is, to display the correct state of the resources, event caching must be switched off in the event publisher.

To do this, perform the following steps:

1. Open the file `/etc/Tivoli/tec/EEZpublisher.conf`.

2. Locate the entry for the affected node.
3. In the relevant entry, change the setting for BufferEvents to NO.

Example:

This is the entry for the node "sapb04" in the file EEZpublisher.conf. The setting for BufferEvents has been changed to NO:

```
ServerLocation=sapb04
ServerPort=5529
ConnectionMode=connection_less
BufferEvents=NO
BufEvtPath=/etc/Tivoli/tec/EEZPublisher.cache
NO_UTF8_CONVERSION=YES
```

WebSphere Application Server cannot be started - DB2 is used as the user registry

This may indicate a problem with the DB2 instance account for the end-to-end automation management databases. To eliminate this as the cause of the problem, check whether the password for the DB2 instance account has expired or is incorrect.

WebSphere Application Server cannot connect to DB2

When you receive an error message indicating that WebSphere Application Server could not establish a connection with the DB2 database EAUTODBDS, this may indicate that the DB2 port number is not specified correctly on the WebSphere Application Server administrative console.

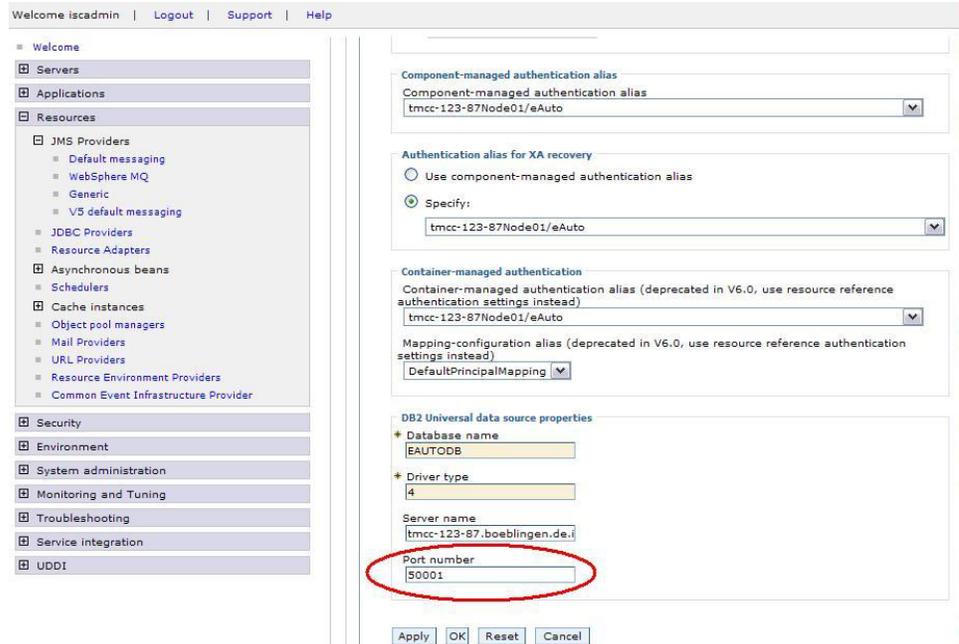
To check if this is the case, perform these steps:

1. On the DB2 server system, check which port number DB2 is using. On Linux, for example, use the **netstat** command to obtain the following information:

```
tmcc-123-87:~ # netstat -atnp | grep db2
tcp    0  0  0.0.0.0:50001      0.0.0.0:*        LISTEN      622/db2tcpm 0
tcp    0  0  9.152.123.87:50001  9.152.123.87:33090 ESTABLISHED 1362/db2agent (EAUT
tcp    0  0  9.152.123.87:50001  9.152.123.87:32954 ESTABLISHED 1379/db2agent (OPCO
```

In the example, the correct DB2 port number is 50001.

2. On the WebSphere Application Server administrative console, navigate to **Resources** → **JDBC providers** → **DB2 Universal JDBC Driver (XA)** → **Data sources** → **EAUTODBDS** and check whether the port number is specified correctly in the field **Port number**:



"Unable to set up the event path..." error message is displayed in Integrated Solutions Console

When you try to connect the operations console, the following error message is displayed in Integrated Solutions Console:

Unable to set up the event path between the operations console
and the management server:

CWSIA024E: An exception was received during the call to the method
JmsManagedConnectionFactoryImpl.createConnection:

com.ibm.websphere.sib exception SIRExourceException:

CWSIT0006E: It is not possible to contact a messaging engine in bus EEZBus

Regardless of whether you are using DB2 or LDAP as the user registry, this may indicate a problem with the DB2 instance account for the end-to-end automation management databases. To check if this is the case, check whether the password for the DB2 instance account has expired or is incorrect.

EEZBus is not started

The EEZBus is a component running within WebSphere Application Server that contains the automation J2EE framework. There are several potential reasons why the EEZBus cannot be started. The reasons and proposed actions are described in the following sections.

EEZBus is not started due to a security problem

If the EEZBus cannot be started, this may indicate a problem with the DB2 instance account for the end-to-end automation management databases, regardless of whether you are using DB2 or LDAP as the user registry.

In such a case, one or more of the following symptoms may occur:

- On the Messages engine panel of the WebSphere Application Server administrative console (Service integration —> Buses —> Messages engine) you can see that the EEZBus is not started. When you try to start the bus, the following error message is displayed:

The message engine <node_name.server_name> EEZBus cannot be started.

- Message "EEZD0010E" appears in the automation engine log file msgengine.log.
- If you are using DB2 as the user registry, the following exception appears in the WebSphere Application Server log file:

```
00000f1d FreePool      E   J2CA0046E:
Method createManagedConnectionWithMCWrapper caught an exception
during creation of the ManagedConnection for resource jms/
EEZTopicConnectionFactory,
throwing ResourceAllocationException.
Original exception: javax.resource.ResourceException:
CWSJR1028E: An internal error has occurred.
The exception com.ibm.websphere.sib.exception.SIResourceException:
CWSIT0006E: It is not possible to contact a messaging engine in bus EEZBus.
was received in method createManagedConnection.
```

- If you are using LDAP as the user registry, the following exception appears in the WebSphere Application Server log file:

```
000000a2 FreePool      E   J2CA0046E:
Method createManagedConnectionWithMCWrapper caught an exception
during creation of the ManagedConnection for resource jdbc/EAUTODBDS,
throwing ResourceAllocationException.
Original exception: com.ibm.ws.exception.WsException:
DSRA8100E: Unable to get a XAConnection from the DataSource.
with SQL State : null SQL Code : -99999
```

To eliminate a problem with the DB2 instance account as the cause, check the database connection from the administrative console:

1. Select the data source.
2. Click **Test connection**.

If the DB2 instance account for the end-to-end automation management databases causes the problem, you receive the following message:

```
Test connection failed for data source EAUTODBDS
on server <serverName> at node <nodeName> with the following exception:
java.lang.Exception: java.sql.SQLException:
Connection authorization failure occurred.
Reason: password invalid. DSRA0010E: SQL State = null, Error Code = -99,999.
```

EEZBus is not started because an internal database is in an inconsistent state

Check if the message log file of WebSphere Application Server contains the following message (where sapb11Node01.server1-EEZBus must be replaced with the messaging engine name based on the node name of your WebSphere Application Server installation):

```
[3/1/06 11:52:37:847 CET] 00000019 SibMessage
E   [EEZBus:sapb11Node01.server1-EEZBus]
CWSIS0002E:
The messaging engine encountered an exception while starting.
Exception: com.ibm.ws.sib.msgstore.PersistenceException:
CWSIS1501E:
The data source has produced an unexpected exception:
java.sql.SQLException: Failed to create database
'/opt/IBM/WebSphere/AppServer/profiles/default/databases/com.ibm.ws.sib/sapb11Node01.server1-EEZBus',
see the next exception for details.
DSRA0010E: SQL State = XJ041, Error Code = 40,000DSRA0010E: SQL State = XJ041, Error Code = 40,000
```

If this message exists, check if the directory described in the message exists in the file system. If it does, complete the following steps:

- Stop the WebSphere Application Server.
- Rename (or remove) the directory described in the message.
- Start the WebSphere Application Server.
- Verify in the WebSphere Application Server message log that the error message shown above does no longer appear and that the EEZBus was started successfully:

```
CWSID0016I: Messaging engine sapb11Node01.server1-EEZBus is in state Started.
```

Note: Similarly, if the CommonEventInfrastructure_Bus cannot be started and an analogous message appears in the WebSphere Application Server message log, remove the directory described in the message, and restart the WebSphere Application Server.

Appendix F. Operations console icons

The tables in this appendix list the icons that are displayed on the operations console.

Icons in the topology tree

In the tables that are presented in this section you find the icons that are used to indicate the state of domains and nodes in the topology tree.

The following table lists the domain icons that are displayed in the topology tree.

Table 56. Domain icons on the operations console

| Icon | Description |
|---|---|
|  | The domain is online. |
|  | The domain is online. A new severe error message has been added to the log file of the domain. |
|  | The domain is online. Commands and queries can be issued against this automation domain, but at least one resource event was lost. |
|  | The domain is online. Commands and queries can be issued against this automation domain, but at least one resource event was lost. A new severe error message was added to the log file of the domain. |
|  | The domain is online. Commands and queries can be issued against the domain, but resource events cannot be received. |
|  | The domain is online. Commands and queries can be issued against the domain, but resource events cannot be received. A new severe error message has been added to the log file of the domain. |
|  | The domain is online. No commands or queries can be issued against the domain, but resource events can be received. |
|  | The domain is online. No commands or queries can be issued against the domain, but resource events can be received. A new severe error message has been added to the log file of the domain. |
|  | The domain is online. No commands or queries can be issued against the domain, but resource events can be received. At least one resource event was lost. |
|  | The domain is online. No commands or queries can be issued against the domain but resource events can be received. At least one resource event was lost. A new severe error message has been added to the log file of the domain. |
|  | The domain is offline. |
|  | The domain is offline. A new severe error message has been added to the log file of the domain. |
|  | The domain is offline. At least one resource event was lost. |
|  | The domain is offline. At least one resource event was lost. A new severe error message has been added to the log file of the domain. |

Table 56. Domain icons on the operations console (continued)

| Icon | Description |
|---|--|
|  | The domain is offline. Commands and queries can be issued against the domain, but resource events cannot be received. |
|  | The domain is offline. Commands and queries can be issued against the domain, but resource events cannot be received. A new severe error message has been added to the log file of the domain. |
|  | The domain is offline. No commands or queries can be issued against the domain, but resource events can be received. |
|  | The domain is offline. No commands or queries can be issued against the domain, but resource events can be received. A new severe error message has been added to the log file of the domain. |
|  | The domain is offline. No commands or queries can be issued against the domain, but resource events can be received. At least one resource event was lost. |
|  | The domain is offline. No commands or queries can be issued against the domain but resource events can be received. At least one resource event was lost. A new severe error message has been added to the log file of the domain. |
|  | None of the communication paths to this domain are currently available. No queries can be submitted and no events can be received. |
|  | None of the communication paths to this domain are currently available. No queries can be submitted and no events can be received. A new severe error message has been added to the log file of the domain. |

The following table lists the node icons that are displayed in the topology tree.

Table 57. Node icons on the operations console

| Icon | Description |
|---|--|
|  | The node is online. |
|  | The node is online. It is excluded from automation. |
|  | The node is offline. |
|  | The node is offline. It is excluded from automation. |

Resource icons

The following tables list the icons that are displayed in the resource table and in the information area.

First-level automation resources

Table 58. Icons for first-level automation resources

| Icon | Description |
|---|---|
|  | A first-level automation resource. The resource is online and works as desired. |

Table 58. Icons for first-level automation resources (continued)

| Icon | Description |
|---|---|
|  | The first-level automation resource is online. The warning icon indicates that it is in state Warning. |
|  | The first-level automation resource is online. The red error icon indicates that the resource is in state Error. |
|  | The first-level automation resource is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A first-level automation resource against which an operator request has been submitted. The resource is online. |
|  | A first-level automation resource against which an operator request has been submitted. The resource is online and in state Warning. |
|  | A first-level automation resource against which an operator request has been submitted. The resource is online. The red error icon indicates that the resource is in state Error. |
|  | A first-level automation resource against which an operator request has been submitted. The resource is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A first-level automation resource against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource is online. |
|  | A first-level automation resource against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource is online. The warning icon indicates that the resource is in state Warning. |
|  | A first-level automation resource against which an operator request has been submitted. The resource is online. The green operator icon indicates that the request has been completed successfully. The red error icon indicates that the resource is in state Error. |
|  | A first-level automation resource against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The first-level automation resource is offline. |
|  | The first-level automation resource is offline. The warning icon indicates that the resource is in state Warning. |
|  | The first-level automation resource is offline. The red error icon indicates that the resource is in state Error. |
|  | The first-level automation resource is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A first-level automation resource against which an operator request has been submitted. The resource is offline. |
|  | A first-level automation resource against which an operator request has been submitted. The resource is offline. The warning icon indicates that the resource is in state Warning. |

Table 58. Icons for first-level automation resources (continued)

| Icon | Description |
|------|--|
| | A first-level automation resource against which an operator request has been submitted. The resource is offline. The red error icon indicates that the resource is in state Error. |
| | A first-level automation resource against which an operator request has been submitted. The resource is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
| | A first-level automation resource against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource is offline. |
| | A first-level automation resource against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource is offline. The warning icon indicates that the resource is in state Warning. |
| | A first-level automation resource against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource is offline. The red error icon indicates that the resource is in state Error. |
| | A first-level automation resource against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Resource references

Table 59. Icons for resource references

| Icon | Description |
|------|--|
| | A resource reference in the end-to-end automation domain. The resource reference is online and works as desired. |
| | The resource reference is online. The warning icon indicates that the resource is in state Warning. |
| | The resource reference is online. The red error icon indicates that the resource is in state Error. |
| | The resource reference is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
| | A resource reference against which an operator request has been submitted. The resource reference is online. |
| | A resource reference against which an operator request has been submitted. The resource reference is online. The warning icon indicates that the resource is in state Warning. |
| | A resource reference against which an operator request has been submitted. The resource reference is online. The red error icon indicates that the resource is in state Error. |
| | A resource reference against which an operator request has been submitted. The resource reference is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Table 59. Icons for resource references (continued)

| Icon | Description |
|---|--|
|  | A resource reference against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource reference is online. |
|  | A resource reference against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource reference is online. The warning icon indicates that the resource is in state Warning. |
|  | A resource reference against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource reference is online. The red error icon indicates that the resource is in state Error. |
|  | A resource reference against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource reference is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A resource reference in the end-to-end automation domain. The resource reference is offline. |
|  | The resource reference is offline. The warning icon indicates that the resource is in state Warning. |
|  | The resource reference is offline. The red error icon indicates that the resource is in state Error. |
|  | The resource reference is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A resource reference against which an operator request has been submitted. The resource reference is offline. |
|  | A resource reference against which an operator request has been submitted. The resource is offline. The warning icon indicates that the resource is in state Warning. |
|  | A resource reference against which an operator request has been submitted. The resource is offline. The red error icon indicates that the resource is in state Error. |
|  | A resource reference against which an operator request has been submitted. The resource reference is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A resource reference against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource reference is offline. |
|  | A resource reference against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource reference is offline. The warning icon indicates that the resource is in state Warning. |
|  | A resource reference against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource reference is offline. The red error icon indicates that the resource is in state Error. |

Table 59. Icons for resource references (continued)

| Icon | Description |
|---|---|
|  | A resource reference against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The resource reference is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Referenced resources

Table 60. Icons for referenced resources

| Icon | Description |
|---|---|
|  | A first-level automation resource that is referenced by a resource reference in the active end-to-end automation policy (referenced resource). The resource is online and works as desired. |
|  | The referenced resource is online. The warning icon indicates that the resource is in state Warning. |
|  | The referenced resource is online. The red error icon indicates that the resource is in state Error. |
|  | The referenced resource is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A referenced resource against which an operator request has been submitted. The resource is online. |
|  | A referenced resource against which an operator request has been submitted. The resource is online. The warning icon indicates that the resource is in state Warning. |
|  | A referenced resource against which an operator request has been submitted. The resource is online. The red error icon indicates that the resource is in state Error. |
|  | A referenced resource against which an operator request has been submitted. The resource is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A referenced resource against which an operator request has been submitted. The resource is online. The green operator icon indicates that the request has been completed successfully. |
|  | A referenced resource against which an operator request has been submitted. The resource is online. The green operator icon indicates that the request has been completed successfully. The warning icon indicates that the resource is in state Warning. |
|  | A referenced resource against which an operator request has been submitted. The resource is online. The green operator icon indicates that the request has been completed successfully. The red error icon indicates that the resource is in state Error. |
|  | A referenced resource against which an operator request has been submitted. The resource is online. The green operator icon indicates that the request has been completed successfully. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Table 60. Icons for referenced resources (continued)

| Icon | Description |
|---|--|
|  | A first-level automation resource that is referenced by a resource reference in the active end-to-end automation policy (referenced resource). The resource is offline. |
|  | The referenced resource is offline. The warning icon indicates that it is in state Warning. |
|  | The referenced resource is offline. The red error icon indicates that the resource is in state Error. |
|  | The referenced resource is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A referenced resource against which an operator request has been issued. The resource is offline. |
|  | A referenced resource against which an operator request has been issued. The resource is offline. The warning icon indicates that the resource is in state Warning. |
|  | A referenced resource against which an operator request has been issued. The resource is offline. The red error icon indicates that the resource is in state Error. |
|  | A referenced resource against which an operator request has been issued. The resource is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A referenced resource against which an operator request has been submitted. The resource is offline. The green operator icon indicates that the request has been completed successfully. |
|  | A referenced resource against which an operator request has been submitted. The resource is offline. The green operator icon indicates that the request has been completed successfully. The warning icon indicates that the resource is in state Warning. |
|  | A referenced resource against which an operator request has been submitted. The resource is offline. The green operator icon indicates that the request has been completed successfully. The red error icon indicates that the resource is in state Error. |
|  | A referenced resource against which an operator request has been submitted. The resource is offline. The green operator icon indicates that the request has been completed successfully. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Resource group icons

Table 61. Resource group icons

| Icon | Description |
|---|---|
|  | A resource group. The group is online. |
|  | The resource group is online. The warning icon indicates that the resource group is in state Warning. |
|  | The resource group is online. The red error icon indicates that the resource group is in state Error. |

Table 61. Resource group icons (continued)

| Icon | Description |
|---|---|
|  | The resource group is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The resource group is online. An operator request has been submitted against the group. |
|  | The resource group is online. An operator request has been submitted against the group. The warning icon indicates that the resource group is in state Warning. |
|  | The resource group is online. An operator request has been submitted against the group. The red error icon indicates that the resource group is in state Error. |
|  | The resource group is online. An operator request has been submitted against the group. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The resource group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. |
|  | The resource group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The warning icon indicates that the resource group is in state Warning. |
|  | The resource group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The red error icon indicates that the resource group is in state Error. |
|  | The resource group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A resource group. The group is offline. |
|  | The group is offline. The warning icon indicates that the resource group is in state Warning. |
|  | The group is offline. The red error icon indicates that the resource group is in state Error. |
|  | The resource group is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The resource group is offline. An operator request has been submitted against the group. |
|  | The resource group is offline. An operator request has been submitted against the group. The warning icon indicates that the resource group is in state Warning. |
|  | The resource group is offline. An operator request has been submitted against the group. The red error icon indicates that the resource group is in state Error. |
|  | The resource group is offline. An operator request has been submitted against the group. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Table 61. Resource group icons (continued)

| Icon | Description |
|---|--|
|  | The resource group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. |
|  | The resource group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The warning icon indicates that the resource group is in state Warning. |
|  | The resource group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The red error icon indicates that the resource group is in state Error. |
|  | The resource group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Referenced first-level automation groups

Table 62. Icons for referenced first-level automation groups

| Icon | Description |
|---|--|
|  | A first-level automation group that is referenced in the active end-to-end automation policy (referenced group). The group is online. |
|  | The referenced group is online. The warning icon indicates that the group is in state Warning. |
|  | The referenced group is online. The red error icon indicates that the group is in state Error. |
|  | The referenced group is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The referenced group is online. An operator request has been submitted against the group. |
|  | The referenced group is online. An operator request has been submitted against the group. The warning icon indicates that the group is in state Warning. |
|  | The referenced group is online. An operator request has been submitted against the group. The red error icon indicates that the group is in state Error. |
|  | The referenced group is online. An operator request has been submitted against the group. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The referenced group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. |
|  | The referenced group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The warning icon indicates that the group is in state Warning. |

Table 62. Icons for referenced first-level automation groups (continued)

| Icon | Description |
|---|--|
|  | The referenced group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The red error icon indicates that the group is in state Error. |
|  | The referenced group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A first-level automation group that is referenced in the active end-to-end automation policy (referenced group). The group is offline. |
|  | The referenced group is offline. The warning icon indicates that the group is in state Warning. |
|  | The referenced group is offline. The red error icon indicates that the group is in state Error. |
|  | The referenced group is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The reference group is offline. An operator request has been submitted against the group. |
|  | The referenced group is offline. An operator request has been submitted against the group. The warning icon indicates that the group is in state Warning. |
|  | The referenced group is offline. An operator request has been submitted against the group. The red error icon indicates that the group is in state Error. |
|  | The referenced group is offline. An operator request has been submitted against the group. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The referenced group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. |
|  | The referenced group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The warning icon indicates that the group is in state Warning. |
|  | The referenced group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The red error icon indicates that the group is in state Error. |
|  | The referenced group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Choice groups and first-level automation move groups

Table 63. Icons for choice groups and first-level automation move groups

| Icon | Description |
|------|---|
| | A choice group of an end-to-end automation domain or a first-level automation move group. The group is online. For a choice group this means that the preferred member is online and the other members are offline. |
| | The choice group or move group is online. The warning icon indicates that the group is in state Warning. |
| | The choice group or move group is online. The red error icon indicates that the group is in state Error. |
| | The choice group or move group is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
| | The choice group or move group is online. An operator request has been submitted against the group. |
| | The choice group or move group is online. An operator request has been submitted against the group. The warning icon indicates that the group is in state Warning. |
| | The choice group or move group is online. An operator request has been submitted against the group. The red error icon indicates that the group is in state Error. |
| | The choice group or move group is online. An operator request has been submitted against the group. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
| | The choice group or move group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. |
| | The choice group or move group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The warning icon indicates that the group is in state Warning. |
| | The choice group or move group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The red error icon indicates that the group is in state Error. |
| | The choice group or move group is online. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
| | A choice group of an end-to-end automation domain or a first-level automation move group. The group is offline. |
| | The choice group or move group is offline. The warning icon indicated that the group is in state Warning. |
| | The choice group or move group is offline. The red error icon indicates that the group is in state Error. |
| | The choice group or move group is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Table 63. Icons for choice groups and first-level automation move groups (continued)

| Icon | Description |
|---|--|
|  | The choice group or move group is offline. An operator request has been submitted against the group. |
|  | The choice group or move group is offline. An operator request has been submitted against the group. The warning icon indicates that group is in state Warning. |
|  | The choice group or move group is offline. An operator request has been submitted against the group. The red error icon indicates that the group is in state Error. |
|  | The choice group or move group is offline. An operator request has been submitted against the group. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The choice group or move group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. |
|  | The choice group or move group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The warning icon indicates that the group is in state Warning. |
|  | The choice group or move group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The red error icon indicates that the group is in state Error. |
|  | The choice group or move group is offline. An operator request has been submitted against the group. The green operator icon indicates that the request has been completed successfully. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Referenced move groups

Table 64. Referenced move groups

| Icon | Description |
|---|---|
|  | A move group that is hosted by a first-level automation domain. The move group is referenced in the active end-to-end automation policy (referenced move group). The group is online. |
|  | The referenced move group is online. The warning icon indicates that the group is in state Warning. |
|  | The referenced move group is online. The red icon indicates that the group is in state Error. |
|  | The referenced move group is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A referenced move group against which an operator request has been submitted. The group is online. |
|  | A referenced move group against which an operator request has been submitted. The move group is online. The warning icon indicates that the group is in state Warning. |

Table 64. Referenced move groups (continued)

| Icon | Description |
|---|---|
|  | A referenced move group against which an operator request has been submitted. The move group is online. The red error icon indicates that the group is in state Error. |
|  | A referenced move group against which an operator request has been submitted. The move group is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A referenced move group against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The move group is online. |
|  | A referenced move group against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The move group is online. The warning icon indicates that the group is in state Warning. |
|  | A referenced move group against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The move group is online. The red error icon indicates that the group is in state Error. |
|  | A referenced move group against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The move group is online. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | The referenced move group is offline. |
|  | The referenced move group is offline. The warning icon indicates that the group is in state Warning. |
|  | The referenced move group is offline. The red error icon indicates that the group is in state Error. |
|  | The referenced move group is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A referenced move group against which an operator request has been submitted. |
|  | A referenced move group against which an operator request has been submitted. The move group is offline. The warning icon indicates that the group is in state Warning. |
|  | A referenced move group against which an operator request has been submitted. The move group is offline. The red error icon indicates that the group is in state Error. |
|  | A referenced move group against which an operator request has been submitted. The move group is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |
|  | A referenced move group against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The move group is offline. |
|  | A referenced move group against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The move group is offline. The warning icon indicates that the group is in state Warning. |

Table 64. Referenced move groups (continued)

| Icon | Description |
|---|--|
|  | A referenced move group against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The move group is offline. The red error icon indicates that the group is in state Error. |
|  | A referenced move group against which an operator request has been submitted. The green operator icon indicates that the request has been completed successfully. The move group is offline. The black error icon indicates that the resource is in an error state from which it cannot be recovered by system automation. |

Information about requests on the General page

When an operator has submitted a start or stop request against a resource, an operator request icon appears on the General page for the resource. The icon indicates the status of the request.

You can use the operator icons to display more information about the request:

- Move the mouse over the operator request icon to display the user ID of the operator who submitted the request.
- Click the operator request icon to bring up the Request details panel.

Table 65. Operator request icons in the information area

| Operator request icon | Description |
|---|---|
|  | A stop request has been submitted. The yellow operator icon indicates that the observed state of the resource is not Offline yet. |
|  | A start request has been submitted. The yellow operator icon indicates that the observed state of the resource is not Online yet. |
|  | The green operator icon indicates that the stop request has been completed successfully. The observed state of the resource is Offline. |
|  | The green operator icon indicates that the start request has been completed successfully. The observed state of the resource is Online. |

Appendix G. Messages

This appendix lists the messages that are generated by the end-to-end automation management component. For information about additional messages you might encounter while working with the component, for example, messages that are produced by the automation adapters, refer to the documentation for the relevant first-level automation product.

How to interpret messages

Each message begins with a unique message code as shown in the following general example:

EEZS1234E

Note: The messages that are generated by the Smart refresh function of the operations console have the component identifier **AZP**. These messages are listed separately in section “Smart refresh messages” on page 373.

- **EEZ** – component identifier of the end-to-end automation management component
- **S** – represents one of the following prefixes:
 - **A** - Messages issued by automation adapters

Note: SA z/OS adapter messages:

- Within NetView an additional * may be appended to the end of the message text.
- Because these messages are written to the syslog on z/OS, the message text must be in English.
- **B, J, L, T** – Automation J2EE framework messages
- **C** – Messages issued by various utilities
- **D** – Automation engine messages
- **F** – Automation engine resource adapter messages
- **I** – Automation manager resource adapter messages
- **K, X** – Automation Software Development Kit messages
- **N** – Installation-related messages
- **P** – Policy-related messages
- **U** – Operations console messages

In this appendix, the messages are sorted alphabetically by subcomponent prefix.

- **1234** – unique four-digit number
- **E** – one of the following severity code identifiers:
 - **I** for Information
 - **W** for Warning
 - **E** for Error

Messages

This appendix lists the messages of the end-to-end automation management component. The messages are sorted alphabetically by subcomponent prefix.

EEZA0001E Syntax error on line *line number*

Explanation: A syntax error has occurred in the configuration file, for example, a leading = on a line.

System action: The automation adapter stops.

Operator response: Analyze the configuration file for invalid syntax.

EEZA0002E Wrong datatype in key *the key*. Expected *the desired type*, found value "*the value that was found*"

Explanation: The value of the given key cannot be interpreted as the desired type. For example, the system expected a boolean value but found the string "hello".

System action: The automation adapter stops.

Operator response: Analyze the configuration file for invalid key/value pairs.

EEZA0003E The key "*the key that was not found*" was not found and no default value was given

Explanation: The system wanted to retrieve from the configuration file a value that did not exist and no default value was given.

System action: The automation adapter stops.

Operator response: Supply a value for the key in the configuration file.

EEZA0004E Integer out of bounds in key "*the key*". Expected value between *the lower bound expected* and *the upper bound expected*, found *the value parsed*

Explanation: The system expected an integer value between the given bounds (inclusive) for the given key, but found a value outside these bounds.

System action: The automation adapter stops.

Operator response: Supply a value within the given bounds for the key.

EEZA0006E Cannot create an instance of the class because class not found: *class name*

Explanation: The automation adapter cannot load the class.

System action: The automation adapter rejects the request.

Operator response: Check whether the class name is valid and is available in the corresponding classpath.

EEZA0007E Cannot create an instance of the class because method not found: *class name*

Explanation: The automation adapter can load the class but cannot create an instance.

System action: The automation adapter rejects the request.

Operator response: Check whether the class is valid.

EEZA0008E Cannot create an instance of the class because of an unknown error: *class name*

Explanation: The automation adapter cannot load the class or create an instance.

System action: The automation adapter rejects the request.

Operator response: Check whether the class is valid and analyze the attached original exception.

EEZA0009E Invocation of adapter plug-in failed: *plug-in=plug-in name, method=method name, internalRetcode=internal return code, taskRetcode=task return code*

Explanation: The automation adapter client API was called to execute a task on the remote automation adapter. There are three error categories:

- The client suffers an error on the connection
- The execution of the task within the automation adapter backend failed
- Execution failed in the automation adapter plug-in.

Automation adapter return codes:

- | | |
|----|--|
| 3 | Backend detects that the request expired before it was started. |
| 4 | Backend cannot be loaded. |
| 22 | Backend cannot load the plug-in Java class. |
| 23 | An exception occurred in the backend or the plug-in returns a non-zero task return code. |
| 24 | Backend rejects the execution of the task. See the automation adapter trace for details. |
| 30 | Backend detects invalid input arguments. Cannot execute task. This is an internal error. |
| 40 | Authentication failed. The user ID or password has not been specified or an internal error occurred. |
| 41 | Authentication failed. The user ID or password is invalid on the automation adapter host. |

- 42 Authentication failed. The user ID has been revoked. JAAS login modules may not provide this information or may also use it if the password has expired.
- 43 Authentication failed. The password has expired. JAAS login modules may not provide this information.

The automation adapter client connection return codes:

- 101 Connection timed out while open.
- 102 Connection failed during open/read/write because peer closed.
- 103 Not all data received. The data was read successfully but the number of bytes is too small. This might be an internal error.
- 104 Unable to connect to the automation adapter.
- 105 Unable to close connection. This might be an internal error.
- 106 Unable to send data. Connection has been aborted due to, for example, SSL mismatch or the automation adapter stopped.
- 107 Marshal error occurred while sending or receiving data. This is an internal error.
- 108 Connection is invalid, for example, connection has not been opened or previously failed.
- 109 Read response failed. Data is zero, for example, the connection has been aborted due to an SSL mismatch.
- 110 Connection timed out during read/write.

Predefined task return codes from plug-in:

- <0 Internally used task return codes have been mapped to one of the automation adapter return codes.
- 0 The plug-in was successful.
- 4004 The plug-in connection timed out.
- 4008 The plug-in authorization is missing to execute the request.
- 4012 The plug-in failed but might execute the next request successfully. This is the default return code if the plug-in throws an exception but the plug-in return code is zero.
- 4016 The plug-in failed and might not execute the next request successfully.
- 4020 The plug-in failed and forces the entire automation adapter to stop.

System action: Execution of the remote task fails.

Operator response: Analyze the return code description. Contact IBM Support if it is an internal error.

EEZA0010E Request expires before the adapter passes it to the adapter plug-in. Timeout period is *timeout value* seconds

Explanation: All requests have an associated expiration date. The request is scheduled to an execution thread that detected that the expiration time had expired.

System action: The automation adapter rejects the request.

Operator response: Analyze the reason (for example, high working load). Increase the timeout period if necessary.

EEZA0011E The backend program specification is invalid

Explanation: The backend program is not a Java program or the Java program name was not specified.

System action: The automation adapter rejects the request.

Operator response: Check the program that called the automation adapter client API or contact IBM Support.

EEZA0012E Invalid parameter list

Explanation: The automation adapter detected a request that is associated with an invalid parameter list.

System action: The automation adapter rejects the request.

Operator response: Check the program that called the automation adapter client API or contact IBM Support.

EEZA0013E Authentication for user ID *user name* failed

Explanation: The request is associated with a user ID and password that have been validated unsuccessfully.

System action: The automation adapter rejects the request.

Operator response: Check whether the user ID is authorized for the system and check the security policy.

EEZA0014E The original exception *original-class* needs to be transported to the remote caller

Explanation: An exception from an underlying component needs to be transported to the remote caller.

System action: None.

Operator response: Analyze the original exception attached with this message.

EEZA0015E Method not supported: *name of the missing method*

Explanation: The automation adapter detected an unknown method name. The list of all valid method names is defined in the EEZAdapterInteraction interface.

System action: The automation adapter rejects the request.

Operator response: Contact IBM Support.

EEZA0017E Request not supported: *name of the unsupported request*

Explanation: The automation adapter plug-in does not support the specified request.

System action: The request might be rejected depending on the behavior of the plug-in.

Operator response: None.

EEZA0022E Adapter client is unable to connect to the adapter at *host:port* due to exception: *the exception that was caught*

Explanation: The automation adapter client cannot connect to the server at the given host and port. The original exception text is provided.

System action: The connection is not established.

Operator response: Analyze the original exception.

EEZA0023E Cache directory is invalid

Explanation: The EIF cache directory is not a directory.

System action: The automation adapter stops.

Operator response: Correct the configuration file.

EEZA0024E EIF sender and receiver must not be equal

Explanation: The EIF configuration parameters are not allowed to point to each other.

System action: The automation adapter stops.

Operator response: Correct the configuration file.

EEZA0025E Cannot find the plug-in configuration file: *configuration file name*

Explanation: The master configuration file contains the name of a plug-in configuration file that cannot be found.

System action: The automation adapter stops.

Operator response: Correct the configuration file.

EEZA0026E No plug-in configuration file was specified

Explanation: The master configuration file must contain at least one plug-in configuration file.

System action: The automation adapter stops.

Operator response: Correct the configuration file.

EEZA0027E Cannot load configuration file: *configuration file name*

Explanation: The specified configuration file cannot be loaded.

System action: The automation adapter stops.

Operator response: Correct the configuration file.

EEZA0028E Plug-in configuration file does not contain all mandatory parameters: *configuration file name*

Explanation: The specified configuration file does not contain all mandatory parameters. The plug-in is not used.

System action: The automation adapter does not deploy the plug-in.

Operator response: Correct the configuration file.

EEZA0029E Cannot create the first instance of the plug-in class: *class name*

Explanation: An attempt was made to create the first instance of the plug-in during initialization. Creation failed.

System action: The automation adapter does not deploy the plug-in.

Operator response: Correct the configuration file.

EEZA0030E Cannot set up event subscription list for plug-in configuration file: *plug-in configuration file name*

Explanation: The specification of the EIF event classes in the plug-in configuration file is invalid.

System action: The automation adapter does not deploy the plug-in.

Operator response: Correct the configuration file.

EEZA0031E Cannot load configuration file from: *plug-in configuration file name*

Explanation: The automation adapter cannot load the specified configuration file because either no configuration file or an invalid one was specified.

System action: The automation adapter stops.

Operator response: Check whether the name of the configuration file is correct.

EEZA0032E Initialization of the adapter failed: *original exception*

Explanation: An error occurred in the initialization step of the automation adapter.

System action: The automation adapter stops.

Operator response: Analyze the associated exception. If there is no exception text for this message, try to find additional messages that were sent previously.

EEZA0033E Unable to create *type of factory* SocketFactory

Explanation: The automation adapter server or client cannot create a socket factory for the remote contact.

System action: The automation adapter client cannot create a connection or the automation adapter server cannot receive connections.

Operator response: Analyze the reason using previous messages.

EEZA0036E The adapter suffered an unexpected interruption: *original exception*

Explanation: The automation adapter waits for a termination command. An unexpected interruption occurred.

System action: The automation adapter stops.

Operator response: Analyze the original exception.

EEZA0037E The adapter stops because no plug-in has been successfully initialized

Explanation: At least one plug-in must have been successfully initialized otherwise the automation adapter stops.

System action: The automation adapter stops.

Operator response: Analyze previous messages and exceptions issued by the failing plug-in.

EEZA0038E A (SSL) socket configuration error occurred: *exception text*

Explanation: An error occurred during the loading or processing of (SSL) socket-related configuration data. An SSL handshake exception will only be reported during initial contact.

System action: The automation adapter client cannot create a connection or the automation adapter server cannot receive connections.

Operator response: Analyze the exceptions text. Check the SSL configuration file, `ing.adapter.ssl.properties`, if necessary.

EEZA0039E Not all data was read from socket: *number of bytes read bytes read, number of bytes expected bytes expected to be read*

Explanation: The incoming request has a length in bytes, but not all bytes can be read.

System action: The automation adapter rejects the request.

Operator response: Check why the socket connection was broken while transferring data.

EEZA0040E The adapter client cannot establish connection to the adapter: *string representation of the connection*

Explanation: Opening the connection failed. A request cannot be sent to the automation adapter. The string representation of the connection contains details about the connection.

System action: The automation adapter frontend failed.

Operator response: Analyze the connection information.

EEZA0041E The adapter client cannot invoke an adapter request: *InternalRC=internal return code, TaskRC=task return code*

Explanation: A connection to the automation adapter has been successfully established. The automation adapter frontend might have sent a request to the automation adapter but the request failed. If the internal or task return codes are not applicable (n/a), some other unexpected exception occurred.

System action: The automation adapter frontend failed.

Operator response: Analyze the internal and task return codes (see EEZA0009E for an explanation of the return codes).

EEZA0042E The adapter has thrown a remote exception: *InternalRC=internal return code, TaskRC=task return code. The original message was: message text*

Explanation: A connection to the automation adapter has been successfully established. The automation adapter frontend has sent a request to the automation adapter but the plug-in has thrown an exception.

System action: None.

Operator response: Analyze the internal and task return codes (see EEZA0009E for an explanation of the return codes).

EEZA0043E A required command line parameter is missing

Explanation: One of the required command line parameters is missing (such as `-start`, `-stop` or `-terminate`).

System action: The automation adapter frontend failed.

Operator response: Specify the required command-line parameters and try again.

EEZA0045E The adapter cannot establish a server socket due to illegal arguments: *exception text*

Explanation: The automation adapter cannot establish a receiver thread and cannot accept incoming connections.

System action: The automation adapter stops.

Operator response: Analyze the configuration file for an invalid IP address.

EEZA0047E The adapter is unable to accept connections due to socket exception "*exception*"

Explanation: An exception occurred as the automation adapter was about to accept an incoming connection.

System action: The automation adapter stops.

Operator response: Analyze the exception text.

EEZA0051W Termination of the adapter failed due to exception: *error message*

Explanation: The attempt to stop the receiver thread failed because an exception occurred.

System action: None.

Operator response: Analyze the exception text.

EEZA0052E Cannot create an in-storage EIF configuration file: *exception text*

Explanation: An instance of the Java class `ByteArrayInputStream` cannot be created or written.

System action: The automation adapter stops.

Operator response: This is probably an internal error. The exception text might give the reason for the problem.

EEZA0053E Missing argument for command line parameter "*the parameter*"

Explanation: A required argument for a command line parameter (such as `-start`) is missing. For example, `AdapterCmd -start` would be incorrect because `-start` requires an argument. A correct example would be: `AdapterCmd -start com.ibm.ing.saplugin.INGXPluginInvocation`

System action: Processing of this command ends.

Operator response: Check the documentation for information about valid command line arguments and their parameters.

EEZA0055E Remote Contact inactivity threshold exceeded: elapsed seconds=*elapsed seconds* threshold=*threshold*

Explanation: The automation adapter calculates the elapsed time since the last synchronous request was received. The automation adapter stops itself if this time exceeds the number specified in the parameter `eez-remote-contact-activity-interval-seconds`. Any incoming event is used as a trigger for the calculation.

System action: The automation adapter stops.

Operator response: You might want to increase the number of seconds specified by the parameter `eez-remote-contact-activity-interval-seconds`. Setting this parameter to 0 (zero) means it never expires.

EEZA0056I Initial contact was enabled and the connection to the management server has been established

Explanation: The parameter `eez-initial-contact` was set to true and the automation adapter attempted to connect the management server. The handshake to the management server was successful.

System action: None.

Operator response: None.

EEZA0057E The connection to the management server cannot be established

Explanation: The automation adapter stops attempting to connect to the management server because the timeout interval is over.

System action: The automation adapter stops.

Operator response: You might want to increase the number of minutes specified by the parameter `eez-initial-contact-retry-interval-minutes`. Specify the value 0 (zero) in order to retry forever.

EEZA0058E The plug-in has not been deployed or not yet started: *name of the Java plug-in class*

Explanation: An attempt was made to issue a request against a non-deployed plug-in or a plug-in that has not been started.

System action: The automation adapter rejects the request.

Operator response: Check the plug-in configuration file and deploy the missing plug-in class. Search for message EEZA0115I.

EEZA0059E An internal error occurred

Explanation: The automation adapter detected an internal error.

System action: None.

Operator response: Contact IBM Support. Check the trace file for the place that threw the internal exception.

EEZA0060I The termination of the adapter is delayed for *duration of the delay in seconds* seconds

Explanation: Stopping the automation adapter is delayed for a short while until it has sent the appropriate domain leave events. You can configure the duration of this delay with the `eez-stop-delay-seconds` parameter.

System action: The automation adapter attempts to send domain leave events.

Operator response: None.

EEZA0061E Unable to bind a socket to address *eez-remote-contact-hostname* at port *eez-remote-contact-port*. Reason: *message of the exception*

Explanation: The automation adapter was unable to use this address or port. Possible causes of the problem are:

- The port is already in use by another program.
- The address could not be assigned.

System action: The automation adapter stops.

Operator response: Make sure that no program is using this port (that is, an automation adapter that is already running). If another program needs this port, then configure the automation adapter to use another port (with the `eez-remote-contact-port` parameter in the master configuration file). Ensure that the address is valid.

EEZA0062I The start command of the automation plug-in *name of the Java plug-in class* was successful

Explanation: The selected automation plug-in was successfully started.

System action: The automation adapter has started the automation plug-in.

Operator response: None.

EEZA0063I The stop command of the automation plug-in *name of the Java plug-in class* was successful

Explanation: The selected automation plug-in was successfully stopped.

System action: The automation adapter has stopped the automation plug-in.

Operator response: None.

EEZA0064I The termination command for the adapter was successful

Explanation: The automation adapter was successfully stopped.

System action: The automation adapter stops.

Operator response: None.

EEZA0070E The host name *eez-remote-contact-hostname* is unknown

Explanation: The automation adapter was unable to resolve the host name.

System action: The automation adapter stops.

Operator response: Specify a valid host name.

EEZA0071E The domain name is either null or empty

Explanation: The plug-in returned an invalid domain name since it is either null or empty.

System action: The plug-in cannot be started.

Operator response: Specify a valid domain name in the plug-in configuration file.

EEZA0100I The adapter has been started

Explanation: This is the first of a sequence of three messages until the automation adapter is ready. The automation adapter starts initialization and will try to connect to the management server if `eez-initial-contact=true`.

System action: None.

Operator response: None.

EEZA0101I The adapter is active

Explanation: The automation adapter becomes "active" after a connection has been successfully established to the management server. The automation adapter continues initialization, finds and starts up all plug-ins.

System action: None.

Operator response: None.

EEZA0102I The adapter is ready

Explanation: The automation adapter startup sequence is complete.

System action: None.

Operator response: None.

EEZA0103I The adapter is stopping

Explanation: An internal or an external stop command has been received.

System action: The automation adapter is about to stop.

Operator response: None.

EEZA0104I The adapter has been stopped

Explanation: The automation adapter termination is complete. All possible stop delay periods are over. The process stops immediately.

System action: The automation adapter has stopped.

Operator response: None.

EEZA0105I The adapter has been stopped due to a failure, rc=return code

Explanation: The automation adapter stopped because an error occurred. All possible stop delay periods are over. The process stops immediately. The return code might be:

- 12 if initial contact failed
- 13 if the remote contact activity threshold is exceeded
- 16 if a plug-in forced termination of the automation adapter
- 20 if initialization failed
- 24 if an error occurred after initialization was successful
- 28 if an unsupported environment is detected

System action: The automation adapter stops.

Operator response: Search for error messages that were issued previously. On z/OS return code 28 might be caused by the 64-bit JVM. You should use the 32-bit JVM instead.

EEZA0111I The plug-in is starting: name of the Java plug-in class

Explanation: The automation adapter has already successfully created an instance of the plug-in class and will now call function INIT_DOMAIN.

System action: None.

Operator response: None.

EEZA0112I The plug-in has been started: name of the Java plug-in class

Explanation: The automation adapter plug-in has successfully initialized the domain (INIT_DOMAIN).

System action: None.

Operator response: None.

EEZA0113I The plug-in is stopping: name of the Java plug-in class

Explanation: The automation adapter will call plug-in function TERM_DOMAIN.

System action: None.

Operator response: None.

EEZA0114I The plug-in has been stopped: name of the Java plug-in class

Explanation: The automation adapter plug-in has successfully stopped the domain (TERM_DOMAIN).

System action: None.

Operator response: None.

EEZA0115I The plug-in startup failed: name of the Java plug-in class

Explanation: This message might follow after EEZA0111I, but the attempt to start the plug-in via function INIT_DOMAIN failed. The automation adapter plug-in will not be started automatically.

System action: The plug-in will be disabled. A join event was not sent.

Operator response: You might want to restart the plug-in using the automation adapter start command. Analyze further plug-in messages.

EEZA9991E The message file is not installed

Explanation: The English message file must be available.

System action: The automation adapter stops.

Operator response: Make sure that the message file is in the class path

EEZB0001E The following runtime exception occurred: exceptionText

Explanation: Processing was interrupted by a RuntimeException and cannot complete correctly.

System action: The current task ends.

Operator response: If you cannot resolve the problem

provide the trace file and this message text to IBM support.

EEZB0003E A critical error has occurred in class: *className*, method: *methodName*. The logger object could not be initialized.

Explanation: This component could not initialize and access a logger object.

This indicates either a configuration or programming error.

System action: The process cannot be completed. All parts of this component are affected. The system is not operational.

Operator response: Please check that the path settings are correct and all required libraries exist. If the problem persists, collect trace data and this message text and contact IBM support.

EEZB0015E An error has occurred in class: *className*. The variable 'locale' has not been set.

Explanation: This variable is mandatory for executing a request at the automation engine.

This indicates a programming error on the EJB exploiter side.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZB0025E An error has occurred in class: *className*, method: *methodName*. Illegal parameter at invocation of this method.

Explanation: The method has been invoked with a null parameter.

The method must be invoked with a parameter that is not null.

This indicates a programming error.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZB0030E A exception has occurred in class: *className*, method: *methodName*. The nested exception equals null.

Explanation: No exception object was linked to the ResourceException that has been caught.

This is an unexpected behavior and indicates a programming error on the J2C side.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZB0031E An error has occurred in class: *className*, method: *methodName*. Invalid nested exception: *nestedException*.

Explanation: An invalid exception object was linked to the ResourceException that has been caught. This is an unexpected behavior and indicates a programming error on the J2C side.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZB0033E An error has occurred in class: *className*, method: *methodName*. JNDI Name: *jndiName* did not return a ConnectionFactory object.

Explanation: The JNDI lookup of this J2C has encountered an internal error. The ConnectionFactory object could not be retrieved.

This indicates a JNDI configuration error.

System action: The current task ends. No connection to the automation engine will be possible until this problem is fixed.

Operator response: Please ensure that the JNDI settings for the J2C connection factories are correct and restart the server. Contact IBM support if the problem persists.

EEZC0001I Setting up Tivoli Common Directory at location where Tivoli Common Directory is being set up.

Explanation: The Tivoli Common Directory path was set to its default value, as shown in the message text.

System action: No system action required.

Operator response: No operator action required.

EEZC0002I Unable to determine Tivoli Common Directory. Diverting serviceability related output to *alternative location*.

Explanation: The system was not able to determine the Tivoli Common Directory.

System action: Processing continues. The application will attempt to divert serviceability related output to another location for this session.

Operator response: In order to manage its serviceability related output, the application should be granted read/write permission to the location /etc/ibm/tivoli/common (UNIX) or <Program_Files_Dir>\ibm\tivoli\common (Windows).

EEZC0003I Base output directory for serviceability related files (for example, message log files and trace files) has been set to *new output directory*.

Explanation: The output directory for serviceability related files was set to its default value, as shown in the message text.

System action: From now on the application will write serviceability related information to the directory that is contained in the message text.

Operator response: No action is required if the base output directory for serviceability related files is acceptable. Otherwise, if it is required to relocate the base output directory, modify the entry in log.properties which should be located at <Program_Files_dir>\ibm\tivoli\common\cfg (Windows) or /etc/ibm/tivoli/common/cfg/log.properties (UNIX). Changes to this file will take effect once the corresponding component is restarted.

EEZC0004I Changing base output directory for serviceability related files of *name of logger* from *old output directory* to *new output directory*.

Explanation: Due to changes in configuration settings the output directory of serviceability related files has been relocated.

System action: From now on the system will write serviceability related information to the new location.

Operator response: No action is required if the base output directory for serviceability related files is acceptable. Otherwise, if it is required to relocate the base output directory, modify the entry in log.properties which should be located at <Program_Files_dir>\ibm\tivoli\common\cfg (Windows) or /etc/ibm/tivoli/common/cfg/log.properties (UNIX). Changes to this file will take effect once the corresponding component is restarted.

EEZC0005W Unable to load IBMTCD10.dll.

Explanation: The system was not able to load "IBMTCD10.dll" which is required to load the value of the locale specific "Program Files" directory from the Windows Registry.

System action: Processing continues. The application will attempt to divert serviceability related output to another location.

Operator response: IBMTCD10.dll should be available to the application. Make sure the directory containing "IBMTCD10.dll" is available in the PATH environment variable.

EEZD0002E An exception was caught while trying to contact the JMS queue of the end-to-end automation manager.

Explanation: The automation engine was unable to establish contact with the end-to-end automation manager. This contact is required to forward EIF events from other automation domains.

System action: The automation engine is unable to contact the server. It has to be restarted when the problem has been resolved.

Operator response: Check the correct configuration for JMS User ID and JMS Password. Restart the automation engine.

EEZD0003E A 'JMSSecurityException' was caught while trying to contact the JMS queue of the end-to-end automation manager.

Explanation: The automation engine was unable to establish contact with the end-to-end automation manager. This contact is required to forward EIF events from other automation domains.

System action: The automation engine is unable to contact the server. It has to be restarted when the problem has been resolved.

Operator response: Check the correct configuration for JMS User ID and JMS Password. Restart the automation engine.

EEZD0004E Method was invoked with a null or empty parameter.

Explanation: A null or empty (list of) parameter was encountered.

System action: The method returns without processing the request any further.

Operator response: Provide the trace file and this message text to IBM support.

EEZD0005E A 'NoPermissionException' was caught while trying to contact the server with hostname: *server hostname* on port: *server port*.

Explanation: The automation engine was unable to create an initial contact with the server in order to start the system automation.

System action: The automation engine is unable to contact the server. It has to be restarted after the problem has been resolved.

Operator response: Check the setup of permissions in order to contact the server as an EJB client. Check the contents of the file specified by the variable 'CLIENTSAS' in the start script of the automation engine process.

EEZD0006E An 'InterruptedException' was caught in method *methodName* of class *className*. The received message was *message*.

Explanation: The processing was interrupted by this exception and cannot complete.

System action: The current task ends.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZD0007E A 'RemoteException' was caught in method *methodName* of class *className*.

Explanation: The processing was interrupted by this exception and cannot complete.

System action: The current task ends.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZD0008E An 'AlreadyBoundException' of service *handlerName* was caught during startup of the automation engine.

Explanation: This service has already bound before.

System action: The automation engine ends.

Operator response: Verify the automation engine is only started once on your system.

EEZD0009E A 'MalformedURLException' of service *handlerName* was caught during startup of the automation engine.

Explanation: This service name is not valid in your environment.

System action: The automation engine ends.

Operator response: Verify the URL. If you cannot resolve the problem provide the trace file and this message text to IBM support.

EEZD0010E The dispatcher thread *threadName* of the automation engine is unable to dispatch a message or request to the automation J2EE framework. The number of retries was: *Number of already attempted retries*. Original exception text was: *Original Exception text*

Explanation: The automation J2EE framework is not accessible at the moment. This message will occur periodically as long as the problem persists.

System action: The thread will keep on trying to contact the automation J2EE framework to dispatch the request or event.

Operator response: Ensure that the application server hosting the automation J2EE framework of the automation manager is up and running. Verify that the automation J2EE framework is running. Check the configuration of the automation engine.

EEZD0013E The required library *<Name of Library>* cannot be loaded by the automation engine.

Explanation: The library, which contains modules required by the automation engine, cannot be loaded.

System action: The core logic engine is unable to initialize. The automation engine is not able to receive any further commands and will terminate.

Operator response: Ensure that the library exists in the lib directory of the automation engine install root.

EEZD0014E A remote exception was caught while executing command line interface function *<Name of CLI method>*.

Explanation: The command line interface function caught a remote exception while trying to connect to the automation engine.

System action: The command line interface was unable to execute the requested function.

Operator response: Ensure that the automation engine is running on this system and that you are allowed to execute commands against it.

EEZD0015E An error occurred preparing a request to the core logic engine used by the automation engine. Internal error message is: *<Reason for encoder error>*.

Explanation: A request to the Core Logic Engine could not be created.

System action: The current action fails.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZD0016E An IOException was caught while trying to load policy *Name of policy* from the policy pool.

Explanation: The policy could not be read.

System action: The policy cannot be parsed. The automation engine is not able to activate this policy and will continue to run with the currently activated policy.

Operator response: Resolve the error indicated by the exception text.

EEZD0017E No or invalid policy pool directory specified. The value specified is: *<Policy Pool Directory>*

Explanation: The value specified for the Policy Pool is either empty or does not specify a directory.

System action: The automation engine was not able to find or load a policy document and will continue to run with the current policy.

Operator response: Ensure that the indicated environment variable points to the correct directory on the local machine where policy documents are placed. Ensure that the automation engine process has sufficient rights to open this directory.

EEZD0019E Unable to activate policy *<Name of policy>*

Explanation: The policy cannot be activated. This policy either does not comply with the required XML syntax or has not passed the consistency check done by the policy checker.

System action: The policy cannot be activated. The automation engine will continue operation with its currently activated policy

Operator response: See detailed error messages logged for this policy above. Resolve the error(s) in this policy before trying to activate it again.

EEZD0020E Domain name *<Name of domain in policy>* in policy *<Name of policy>* does not match automation engine domain *<Name of automation engine domain>*.

Explanation: The policy cannot be activated. It is valid only for the indicated domain.

System action: The policy cannot be activated. The automation engine will continue operation with its currently activated policy.

Operator response: Activate a policy for the currently active domain, or reconfigure and restart the automation engine to represent another domain.

EEZD0021E Expression *<unsupported filter value>* used during query is not supported.

Explanation: The specified filter is not supported.

System action: The enumeration of resources or relationships fails.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZD0022E Unknown event *<Event Type>* with reason *<Event Reason>* was received by the automation engine.

Explanation: This event cannot be handled by the automation engine.

System action: The automation engine continues processing.

Operator response: Contact IBM support and provide the trace file and this message text.

EEZD0023E An EEZConfigurationException was received while trying to set the user credentials for the list of domains specified in the policy.

Explanation: The policy cannot be activated.

System action: The automation engine is not able to activate this policy and will continue to run with the currently activated policy.

Operator response: Contact IBM support and provide the trace file and this message text.

EEZD0024E A Generic user ID or password is not specified in the *<file name>* file.

Explanation: A user and password for the GENERIC domain must always be specified in the properties file.

System action: The load of the properties file failed.

Operator response: Specify a user and password value for the GENERIC domain.

EEZD0025E Key: *<key>* occurs more than once in the properties file.

Explanation: To avoid confusion, only unique keys can be specified in the properties file.

System action: The load of the properties file failed.

Operator response: Delete the duplicate keys from the properties file.

EEZD0027E The required thread *thread name* of the automation engine process terminated. The exception with name *exception name* caused the thread to terminate.

Explanation: The thread terminated because of the indicated unhandled exception.

System action: The automation engine will try to restart this required thread. However, some tasks which were just processed by the thread during this termination may have not been completed. This can cause a loss of events or requests against other automation domains. The automation engine will try to recover from this by doing some synchronization.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0028E A `NotBoundException` of service `<handlerName >` was caught during shutdown processing of the automation engine.

Explanation: This service has not been bound before.

System action: The automation engine ends.

Operator response: None.

EEZD0029E No referenced resources hosted by a first-level automation domain were returned from the core logic engine.

Explanation: The currently activated policy contains no, or invalidly specified, resource references, or a problem occurred during execution of a task within the core logic engine.

System action: The automation engine cannot execute the current task.

Operator response: Check the currently activated policy and activate it again.

EEZD0030E Problem occurred parsing the core logic engine response for orders. An order was received, but an unknown action was found.

Explanation: The response to an internal workitem from the engine could not be parsed correctly.

System action: The automation engine cannot execute the current task.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0032E Method *Method name* cannot be executed, since the automation engine process is starting.

Explanation: A public method of the automation engine has been called, but the automation engine has not finished its startup processing.

System action: The automation engine cannot execute the current task.

Operator response: Wait until the automation engine startup completed message has been returned. If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0033E Method *Method name* cannot be executed, since automation engine process is stopping.

Explanation: A public method of the automation engine has been called, but the automation engine is currently shutting down on request.

System action: The automation engine cannot execute the current task.

Operator response: No action required.

EEZD0034E Unable to create a registry for the automation engine process on port: *Port Number* to bind its services.

Explanation: An exception was caught during startup of the automation engine while trying to create a registry object. This can be caused by another automation engine process running on this system, or another process using the port.

System action: The automation engine cannot complete the startup task. The automation engine process will terminate.

Operator response: Ensure that only one instance of the automation engine process is started on your system. If this is not the cause of this error, find out which process is blocking the port number indicated in the message text.

EEZD0035E Unable to create the Command Line Interface (CLI) handler of the automation engine process.

Explanation: A `RemoteException` was caught during startup of the CLI handler.

System action: The automation engine cannot complete the startup task. The automation engine process will terminate.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0036E Unable to create the main thread of the automation engine process.

Explanation: A `RemoteException` was caught during startup of the main thread.

System action: The automation engine cannot complete the startup task. The automation engine process will terminate.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0037E The following services are bound:
Service Name.

Explanation: The automation engine is unable to create a new registry. The listed services are already bound

System action: The automation engine cannot complete the startup task. The automation engine process will terminate.

Operator response: Ensure that only one instance of the automation engine process is started on your system.

EEZD0038E Unable to start the EIF receiver for the automation engine process on system:
Hostname port: Port Number to bind its services.

Explanation: An exception was caught during startup of the automation engine while trying to start the EIF receiver. This can be caused by another automation engine process running on this system, another process using the port, or a hostname that cannot be resolved.

System action: The automation engine cannot complete the startup task. The automation engine process will terminate.

Operator response: Ensure that only one instance of the automation engine process is started on your system. If this is not the cause of this error, find out which process is blocking the port number indicated in the message text. If necessary configure another Event port number in the automation engine configuration. Ensure that the hostname can be resolved on the local system and that the hostname is not the loopback address.

EEZD0040E Not all required properties are specified within the properties file with name: *file name*. Value for property *Property name* is not specified.

Explanation: The properties file cannot be loaded if not all of the required properties (key/value pairs) are specified in the file.

System action: The load of the properties file failed.

Operator response: Specify the required property correctly in the properties file.

EEZD0041E Property *Property name* exceeds the maximum length of *Maximum length of Property*

Explanation: The properties file contains invalid values for the specified property.

System action: The load of the properties file failed.

Operator response: Specify all required properties

correctly in the properties file.

EEZD0042E Value for property *Property name* is not a numeric value.

Explanation: The specified property must be a number.

System action: The load of the properties file failed.

Operator response: Specify all required properties correctly in the properties file.

EEZD0050E An IOException occurred during startup of the automation engine process.
Original Exception text is: *IO Exception text*

Explanation: The IOException cannot be handled by the automation engine process.

System action: The automation engine process terminates.

Operator response: Ensure that the automation engine has sufficient rights to access the indicated file.

EEZD0060E An EEZIllegalArgumentException was caught when trying to execute a list of requests. Following is a list of requests causing this exception: *List of failing requests*

Explanation: The requests cannot be executed on resources of the targeted automation domain.

System action: The automation engine will not try to invoke these requests again. They will be ignored. The automation engine continues working.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0061E An EEZExecutionFailedException was caught trying to execute a list of requests. Following is a list of requests causing this exception: *List of failing requests*. The following message was returned by the exception: *Exception Message*. Original message contained in the caught exception: *Original message*

Explanation: The requests cannot be executed on specified resources of the targeted automation domain.

System action: The automation engine will not try to invoke these requests again. They will be ignored. The automation engine continues working.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0062E An EEZExecutionTimeoutException was caught trying to execute a list of requests. Following is a list of requests causing this exception: *List of failing requests.*

Explanation: The requests could not be completed within a timeout period.

System action: The automation engine tries to invoke these requests again. The automation engine continues working.

Operator response: No action required.

EEZD0063E An EEZIllegalArgumentException was caught trying to subscribe to other automation domains. Following is a list of domains the automation engine tried to subscribe to: *List of domains to subscribe to*

Explanation: The subscription cannot be executed.

System action: The automation engine will not try to subscribe to these domains again. They will be ignored. The automation engine continues working.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0064E An EEZExecutionFailedException was caught trying to subscribe to other automation domains. Following is a list of domains causing this exception: *List of domains to subscribe to.* The following message was returned by the exception: *Exception Message*

Explanation: The subscription cannot be executed.

System action: The automation engine will not try to subscribe to these domains again. They will be ignored. The automation engine continues working.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0065E An EEZExecutionTimeoutException was caught trying to subscribe to other automation domains. Following is a list of domains causing this exception: *List of domains to subscribe to.*

Explanation: The requests could not be completed within a timeout period.

System action: The automation engine tries to invoke these requests again. The automation engine continues working.

Operator response: No action required.

EEZD0066E An EEZIllegalArgumentException was caught trying to subscribe to resources hosted on other automation domains. Following is a list of resources the automation engine tried to subscribe to: *List of resources to subscribe to*

Explanation: The subscription cannot be executed.

System action: The automation engine will not try to subscribe to these domains again. They will be ignored. The automation engine continues working.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0067E An EEZExecutionFailedException was caught trying to subscribe to resources hosted on other automation domains. Following is a list of resources the automation engine tried to subscribe to: *List of resources to subscribe to.* The following message was returned by the exception: *Exception Message*

Explanation: The subscription cannot be executed.

System action: The automation engine will not try to subscribe to these domains again. They will be ignored. The automation engine continues working.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0068E An EEZExecutionTimeoutException was caught trying to subscribe to resources hosted on other automation domains. Following is a list of resources the automation engine tried to subscribe to: *List of resources to subscribe to*

Explanation: The requests could not be completed within a timeout period.

System action: The automation engine tries to invoke these requests again. The automation engine continues working.

Operator response: No action required.

EEZD0069E A Security Exception was caught trying to subscribe to resources hosted on automation domain with name *Automation domain name.* Following is a list of resources the automation engine tried to subscribe to: *List of resources to subscribe to.*

Explanation: The subscription cannot be executed.

System action: The automation engine will not try to subscribe to these domains again. They will be ignored.

EEZD0072E • EEZD0084E

The automation engine continues working.

Operator response: If you cannot resolve the exception, contact IBM support and provide the trace file and this message text.

EEZD0072E An EEZUserSecurityException was caught trying to contact another automation domain. Original message text is: *Security Exception text*

Explanation: The automation engine was unable to contact the automation domain to execute commands.

System action: The automation engine will not try to invoke the command again.

Operator response: Resolve the error indicated by the exception text. Ensure that the User credentials in the configuration of the automation engine contains correct User IDs and valid passwords to access the indicated automation domain. After fixing the problem it is also required to refresh the automation engine by using the command line interface option -RECONFIG.

EEZD0073E An EEZSystemSecurityException was caught trying to contact another automation domain. Original message text is: *Security Exception text*

Explanation: The automation engine was unable to contact the automation domain to execute commands.

System action: The automation engine will not try to invoke the command again.

Operator response: Resolve the error indicated by the exception text. Ensure that the User credentials in the configuration of the automation engine contains correct User IDs and valid passwords to access the indicated automation domain. After fixing the problem it is also required to refresh the automation engine by using the command line interface option -RECONFIG.

EEZD0074E The properties file contains an invalid line.

Explanation: A line in the properties file is not marked as a comment and cannot be read as a key-value pair.

System action: The load of the properties file failed.

Operator response: Correct a line in the properties file that should be a comment or that is missing the separator character.

EEZD0080E An Exception was caught trying to install the license certificate <Certificate file name>. Original exception was: *Original exception text*

Explanation: The automation engine was unable to install the mentioned license certificate.

System action: The automation engine is unable to handle this license key. It will not perform any function until a license key was successfully installed.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZD0081E An Error occurred trying to install the license certificate <Certificate file name>. Original return code of the license usage management library was: *LUM return code*

Explanation: The automation engine was unable to install the mentioned license certificate.

System action: The automation engine is unable to handle this license key. It will not perform any function until a license key was successfully installed.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZD0082E An Error occurred trying to initialize the license usage management library. Original return code was: *LUM return code*

Explanation: The automation engine was unable to initialize the license use management.

System action: The automation engine is unable to handle license keys and will terminate.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZD0083E No valid installed license key found on your system. Message returned by the license usage management library was: *LUM return message* Original return code was: *LUM return code*.

Explanation: The automation engine was unable to detect an installed license on your system.

System action: Without installed license key the automation engine will not start.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZD0084E The automation engine does not have a valid license enrolled. If you purchased a license for this product, then you have a production license. This license will be automatically installed on your system during the installation process.

Explanation: The automation engine was unable to

detect an installed license on your system.

System action: Without installed license key the automation engine will not start.

Operator response: If you installed this product with help of the product installer the installation process must have encountered a problem. Provide the trace file and this message text to IBM support.

EEZD0100E An error occurred within the core logic engine used by the automation engine. Message code <InternalMessageCode> has been returned.

Explanation: The preprocessing of an internal workitem failed.

System action: The current action fails.

Operator response: Contact IBM Support, and provide the trace file and this message text.

EEZD0101E An error occurred during initialization of the core logic engine used by the automation engine. The logic deck <LogicDeckName> cannot be activated. Message code <InternalMessageCode> has been returned.

Explanation: The core logic engine required by the automation engine could not be created.

System action: The automation engine is not able to receive any further commands and will terminate.

Operator response: Contact IBM Support, and provide trace file and this message text.

EEZD0102E The internal encoder used by the automation engine did not return a result.

Explanation: The current workitem could not be encoded by the automation engine.

System action: The automation engine is not able to execute the current task.

Operator response: Contact IBM Support and provide trace file and this message text.

EEZD0103E The CORE LOGIC ENGINE used by the automation engine did not return any result.

Explanation: The current workitem could not be executed by the automation engine.

System action: The automation engine is not able to execute the current task.

Operator response: Contact IBM Support, and provide trace file and this message text.

EEZD0104E A problem occurred parsing the result returned by the CLE. The automation engine is unable to get the current preferred member of choice group: *Choice group name*

Explanation:

System action: The automation engine cannot complete this task.

Operator response: None

EEZD0105E A problem occurred parsing the result returned by the CLE. The automation engine is unable to get orders from the response given by the Core Logic Engine.

Explanation:

System action: The automation engine cannot complete this task.

Operator response: None

EEZD0106E An error occurred during initialization of the core logic engine used by the automation engine. The logic deck <LogicDeckName> cannot be found or opened.

Explanation: The core logic engine required by the automation engine could not be created.

System action: The automation engine is not able to receive any further commands and will terminate.

Operator response: Check that a correct file name is specified for the Logic definition file of the advanced domain configuration of the automation engine.

EEZD0107E An error occurred during initialization of the automation engine. The domain identification file <DomainIdentificationFile> cannot be found or opened.

Explanation: The automation engine needs the domain identification file in order to set up connections to other automation domains.

System action: The automation engine is not able to receive any further commands and will terminate.

Operator response: Check that a correct file name is specified for the domain identification file of the automation engine.

EEZD0200E A runtime exception was caught during execution of method *method name*.
Original exception text is: *exception name*.

Explanation:

System action: The automation engine cannot complete this task.

Operator response: Resolve the runtime exception, and ensure correct installation of the product.

EEZD1004W The event notifying the automation server about the terminating automation engine cannot be delivered.

Explanation: The automation engine assumes that the automation server has already been terminated.

System action: The automation engine continues to shut down.

Operator response: Check the shutdown sequence. The automation engine should be terminated before the automation server.

EEZD1005W Unable to contact the server on hostname: *server hostname* on port: *server port*.

Explanation: The automation engine assumes a running server in order to start the system automation.

System action: The automation engine will retry to contact this server periodically.

Operator response: Ensure that this server is running and reachable under this hostname. Check the configuration for Host name or IP address of the automation engine.

EEZD1010W One or more resource reference(s) in your currently active policy point to resources that cannot be found on the referenced automation domain with name *Name of automation domain*. Following is the list of resource references that cannot be resolved: *List of Resources which are invalid*

Explanation: The resource reference(s) may have been misspelled in your currently active policy, or they just have not been created on the referenced domain yet.

System action: The automation engine is not able to do any automation task or handle any requests against the specified resource reference(s).

Operator response: Check if the resource references are possibly misspelled. In this case, resolve this error within the policy and activate it again.

EEZD1011W No policy can be found in the policy pool: *Policy Pool Directory* which matches this domain with name: *Domain Name*

Explanation: The automation engine found at least one policy file with extension '.xml' or '.XML' but no policy file matches the current domain represented by the automation engine.

System action: The automation engine is not able to provide a list of available policies.

Operator response: Check that a correct file name is specified for the Policy pool of the advanced domain configuration of the automation engine. Ensure that policy documents inside this policy pool are defined for the domain name configured for the automation engine.

EEZD1012W No policy can be found in the policy pool: *Policy Pool Directory*

Explanation: The automation engine cannot find any file with extension '.xml' or '.XML'

System action: The automation engine is not able to provide a list of available policies.

Operator response: Check that a correct file name is specified for the Policy pool of the advanced domain configuration of the automation engine. Ensure that policy documents inside this policy pool are defined for the domain name configured for the automation engine. Ensure that the user id used to start the automation engine has sufficient rights to access this directory.

EEZD1013W Problems occurred while trying to load policy *Name of policy* from the policy pool.

Explanation: The policy contains errors.

System action: The policy cannot be loaded and parsed. The automation engine is not able to activate this policy.

Operator response: Check the named policy for errors.

EEZD1070W An EEZApplicationException was caught trying to contact another automation domain. Retry number is: *Retry number* Original message text is: *Application Exception text*

Explanation: The automation engine was unable to contact the automation domain to execute commands this time. It might have been caused by a temporary problem.

System action: The automation engine will try to invoke the command again.

Operator response: No action required

EEZD1071W A RemoteException was caught trying to contact another automation domain.
Retry number is: *Retry number* **Original message text is:** *Application Exception text*

Explanation: The automation engine was unable to contact the automation domain to execute commands this time. It might have been caused by a temporary problem.

System action: The automation engine will try to invoke the command again.

Operator response: No action required

EEZD2001I The automation engine has been started by a user with ID *User ID*. The automation engine is idling. No policy activated.

Explanation: No additional information is available for this message.

EEZD2002I The automation engine is shutting down on request by a user with ID *User ID*.

Explanation: No additional information is available for this message.

EEZD2003I The automation engine has been terminated.

Explanation: No additional information is available for this message.

EEZD2004I The automation engine was called to refresh its configuration.

Explanation: No additional information is available for this message.

EEZD2005I The automation engine has been started as a converter of EIF messages by a user with ID *User ID*.

Explanation: No additional information is available for this message.

EEZD2010I The currently active policy *Currently active policy name* will be deactivated before the new policy *New policy name* can be activated.

Explanation: Only one policy can be active at any time. Therefore the old policy will be deactivated first.

System action: The old policy will be deactivated. All defined resource references, groups and relationships will be deleted. All requests issued by any operator against any resource reference or group will be deleted as well.

Operator response: If desired, requests will have to be issued again by an operator for the resource references and groups defined by the new policy. The same is true if the currently active policy is re-activated.

EEZD2011I The policy with name *Active policy name* has been activated by a user with ID *User ID*.

Explanation: The policy has been activated successfully.

System action: The policy has been activated. If another policy was active before, all resource state information and all operator requests against the resource defined in the old policy have been deleted.

Operator response: No action required.

EEZD2012I The active policy has been deactivated by a user with ID *User ID*.

Explanation: The policy has been deactivated successfully.

System action: The policy has been deactivated.

Operator response: No action required.

EEZD2020I The automation engine was informed about a joined automation domain with name *Automation domain name*.

Explanation: The automation domain hosts resources that are referenced by at least one resource reference in the currently active policy. The automation engine will try to get the current state information of these resources.

System action: The automation engine will handle this domain join event.

Operator response: No action required.

EEZD2021I The automation engine was informed about a leaving automation domain with name *Automation domain name*.

Explanation: The automation domain hosts resources that are referenced by at least one resource reference in the currently active policy. The automation engine will update the state of these resource references accordingly.

System action: The automation engine will handle this domain leave event.

Operator response: No action required.

EEZD2022I The automation engine was informed about a communication state change of domain with name *Automation domain name*. The new communication state is: *New communication State name*. The old communication state was: *Old communication State name*.

Explanation: The automation domain hosts resources that are referenced by at least one resource reference in the currently active policy. The automation engine will update the state of these resource references accordingly.

System action: The automation engine will handle this communication state change event.

Operator response: No action required.

EEZD2023I The automation engine was called to synchronize itself with the automation domain with name *Automation domain name*.

Explanation: The automation domain hosts resources that are referenced by at least one resource reference in the currently active policy. The automation engine will verify and update the state of these resource references accordingly.

System action: The automation engine will handle this synchronization request.

Operator response: No action required.

EEZD2024I The automation engine was informed about a domain state change of domain with name *Automation domain name*. The new domain state is: *Communication State name*.

Explanation: The automation domain hosts resources which are referenced by at least one resource reference in the currently active policy. The automation engine will handle requests to these resource references according to the domain state.

System action: Requests will not be sent to a domain which is in the OFFLINE state.

Operator response: No action required.

EEZD2025I The automation engine was informed about a policy change on one automation domain with name *Automation domain name*.

Explanation: The indicated automation domain hosts resources that are referenced by at least one resource reference in the currently active policy. The automation engine will invalidate the existing observed states of these resource references and get the current data by subscribing to this domain again.

System action: The automation engine will handle this event.

Operator response: No action required.

EEZD2030I Request: *Request Name* was issued by User ID: *User Id* against Resource Class with name: *Resource Name*

Explanation:

System action: The automation engine will handle this request.

Operator response: No action required.

EEZD2031I Request was removed by User ID: *User ID* from Resource Class with name: *Resource Name*

Explanation:

System action: The automation engine will handle this remove request.

Operator response: No action required.

EEZD2032I Resource *New preferred member* is set to be the new preferred member of choice group *Choice group name*.

Explanation:

System action:

Operator response: No action required.

EEZD2040I The automation engine was called to install the license certificate from file: *Certificate file name*.

Explanation: A successful installation of the license certificate is required in order for the end-to-end automation engine to run.

System action: The automation engine will handle this request.

Operator response: No action required.

EEZD2041I The license certificate has been successfully installed.

Explanation: A successful installation of the license certificate is required in order for the end-to-end automation engine to run.

System action:

Operator response: No action required.

EEZD9999I DUMMY Message

Explanation: This message should never occur.

System action: Contact IBM support and provide this message text together with the trace.

Operator response:

EEZF0001E The WebSphere infrastructure has reported a severe error situation: *runtimeExceptionMessage*

Explanation: The application was interrupted by a RuntimeException and cannot complete its task.

System action: The current task ends. The transaction is rolled back.

Operator response: Check if the description of the error situation indicates that the server database or another subsystem is unavailable. If the problem persists, contact IBM support.

EEZF0003E A critical error has occurred in class: *className*, method: *methodName*. The logger object could not be initialized.

Explanation: This component could not initialize and access a logger object.

This indicates either a configuration or programming error.

System action: The process cannot be completed. All parts of this component are affected. The system is not operational.

Operator response: Please check that the path settings are correct and all required libraries exist. Collect trace data and this message text and contact IBM support if the problem persists.

EEZF0006E A exception has occurred in class: *className*, method: *methodName*. The connection to the automation engine could not be established.

Explanation: The end-to-end automation application could not connect to the automation engine.

This indicates a network or configuration problem.

System action: The current task ends.

Operator response: Check if the description of the error situation indicates that the automation engine backend is not up and running properly. Please start or restart it in this case. If the problem persists, contact IBM support.

EEZF0007E An error has occurred in class: *className*, method: *methodName*. The connection to the automation engine could not be established.

Explanation: The JNDI lookup of the object representing the automation engine failed and returned null.

This indicates a network or configuration problem.

System action: The current task will not be completed. The system and automation engine remain operational.

Operator response: Check if the description of the error situation indicates that the automation engine backend is not up and running properly. Please start or restart it in this case. If the problem persists, contact IBM support.

EEZF0011E An error has occurred in class: *className*, method: *methodName*. Illegal invocation of this method.

Explanation: This method is required by the J2C specification, but it is not implemented in this component and must not be called by any exploiter. This indicates a programming error or an illegal access to this component.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZF0012E An error has occurred in class: *className*, method: *methodName*. Illegal parameter at invocation of this method.

Explanation: The method has been invoked with a null parameter.

The method must be invoked with a parameter that is not null.

This indicates a programming error.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZF0014E An error has occurred in class: *className*, method: *methodName*. Illegal parameter combination at invocation of this method.

Explanation: This method is required by the J2C specification, but it is not implemented in this component and must not be called by any exploiter. This indicates a programming error or an illegal access to this component.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZF0015E An error has occurred in class: *className*. The variable 'locale' has not been set.

Explanation: This variable is mandatory for executing a request at the automation engine. This indicates a programming error on the EJB exploiter side.

System action: The current task will not be completed. The system and automation engine remain operational.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZF0016E An error has occurred in class: *className*, **method:** *methodName*. Illegal automation engine operation requested.

Explanation: The J2C architecture tried to execute an operation on the automation engine that is either unknown or not supported. This indicates a programming error.

System action: The current task will not be completed. The system and automation engine remain operational.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZF0017E A *exception* has occurred in class: *className*, **method:** *methodName*.

Explanation: The method encountered an expected error during invocation of an operation on the automation engine backend. This indicates a connection problem.

System action: The current task ends.

Operator response: Check if the automation engine backend is up and running properly. Please start or restart it if not. If the problem persists, contact IBM support.

EEZF0019E An error has occurred in class: *className*, **method:** *methodName*. Illegal parameter at invocation of this method.

Explanation: Method `createManagedConnection()` must not be called with a null `ConnectionRequestInfo` object. This is an indication of an internal JCA error.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZF0031E A *exception* has occurred in class: *className*, **method:** *methodName*. A connection object could not be allocated.

Explanation: The call to `getConnection()` returned with an exception.

This indicates an internal J2C programming or architectural error.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZF0032E An error has occurred in class: *className*, **method:** *methodName*. A connection factory object could not be allocated.

Explanation: The `ManagedConnectionFactory` of this J2C encountered an internal error.

The `ConnectionManager` instance was null. This indicates an internal J2C programming or architectural error.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZF0501W An exception was encountered and ignored in order to continue operation. Exception string: *exceptionString*

Explanation: The invoked method is designed to ignore exceptions and continue operation. It logs the exception for problem determination purposes.

System action: Ignores the exception.

Operator response: Evaluate the exception details.

EEZF0502W The connection to the automation engine backend failed in class: *className*, **method:** *methodName*. Exception: *exception*.

Explanation: The application was interrupted by a `NameNotFoundException` and cannot complete its task.

System action: The system might repeat this operation until the automation engine is found.

Operator response: Ensure that the automation engine is up and running.

If the problem persists, please collect logs and traces and provide them to IBM support.

EEZI0001E The WebSphere infrastructure has reported a severe error situation: *runtimeExceptionMessage*.

Explanation: The application was interrupted by a `RuntimeException` and cannot complete its task.

System action: The current task ends. The transaction is rolled back.

Operator response: Check the description of the error situation if it indicates that the server database or another subsystem is unavailable. If the problem persists, contact IBM support.

EEZI0003E A critical error has occurred in class: *className*, method: *methodName*. Unable to initialize Logger.

Explanation: No Logger object could be initialized and accessed.

System action: The process cannot be completed. All parts of this component are affected

Operator response: Contact IBM support and provide this message text.

EEZI0005E Failing Logger initialization in: *variable text*, in class: *className*. Information: *someInfo*

Explanation: Critical error. No logger object could be obtained. The entire application might be affected.

System action: Method terminates with a ConfigurationFailedException.

Operator response: Ensure the correct classpath configuration. Contact IBM support if the problem persists.

EEZI0012E Internal error. Null parameter passed in method: *methodName*, in class: *className*.

Explanation: Method getConnection() must not be called with null parameters. This is an indication of a programming error on the EJB exploiter side.

System action: Method terminates with an IllegalArgumentException.

Operator response: Invoke getConnection() with a fully initialized EEZFLAConnectionSpec object as a valid parameter. Contact IBM support if the problem persists.

EEZI0013E Internal error. Illegal parameter passed in method: *methodName*, in class: *className*.

Explanation: The EEZFLAConnectionSpec parameter contained an uninitialized EEZFLAConfigData member object.

System action: Method terminates with an IllegalArgumentException.

Operator response: Invoke getConnection() with a fully initialized EEZFLAConnectionSpec object as a valid parameter. Contact IBM support if the problem persists.

EEZI0014E Illegal invocation of method: *methodName*, in class: *className*.

Explanation: Method invoke() must not be called with this parameter combination. It is not supported.

System action: Method terminates with an IOException.

Operator response: Invoke invoke() with the signature(InteractionSpec, Record) as a valid parameter combination. Contact IBM support if the problem persists.

EEZI0015E Critical error in class: *className*, method: *methodName*. A connection to the Adapter could not be established.

Explanation: The call to EEZAdapterConnection.open(..) returned value 0.

System action: The method terminates with a ConnectionFailedException.

Operator response: See the WebSphere and Adapter logs and traces. Contact IBM support if the problem persists.

EEZI0016E Critical error in class: *className*, method: *methodName*. Unknown AdapterException return code in *variable text*.

Explanation: The operation has terminated with an AdapterException, but the internal return code is unknown.

System action: The method terminates with a ExecutionFailedException.

Operator response: See the WebSphere and Adapter logs and traces. Contact IBM support if the problem persists.

EEZI0017E Critical error in class: *className*, method: *methodName*. The operation could not be performed because of *exception*.

Explanation: An exception other than a subtype of EEZApplicationException occurred during interaction with the backend.

System action: The method terminates with a ExecutionFailedException.

Operator response: See the WebSphere and Adapter logs and traces. Contact IBM support if the problem persists.

EEZI0018E Internal error. Illegal parameter passed in method: *methodName*, in class: *className*.

Explanation: The EEZFLAConnectionRequestInfo parameter contained an uninitialized EEZFLAConfigData member object.

System action: Method terminates with an IllegalArgumentException.

Operator response: Ensure that the createManagedConnection() method is called with a fully initialized EEZFLAConnectionRequestInfo object as a valid parameter. Contact IBM support if the problem persists.

EEZI0019E Internal error. Illegal invocation of method: *methodName*, in class: *className*.

Explanation: Method createConnection() must not be called without parameters. This is an indication of an internal JCA error.

System action: Method terminates with an IllegalOperationException.

Operator response: Invoke createConnection() with a fully initialized ConnectionManager object as a valid parameter. Contact IBM support if the problem persists.

EEZI0021E Security violation detected in class: *className*, method: *methodName*. SSL required but not enabled.

Explanation: The ConnectionFactory of this JCA requires SSL-secure connections, but the current request does not comply.

System action: The method terminates with a SystemSecurityException.

Operator response: Check the security properties of the FLA ConnectionFactory and ensure compliant requests. Contact IBM support if the problem persists.

EEZI0022E Security violation detected in class: *className*, method: *methodName*. The SSL configuration file could not be found.

Explanation: The ConnectionFactory of this JCA requires SSL-secure connections, but the file containing the necessary properties could not be found.

System action: The method terminates with a ConfigurationException.

Operator response: Check the properties of the FLA ConnectionFactory and ensure that the file is in the correct location. Contact IBM support if the problem persists.

EEZI0023E Security violation detected in class: *className*, method: *methodName*. The SSL configuration file could not be opened.

Explanation: The ConnectionFactory of this JCA requires SSL-secure connections, but the file containing the necessary properties could not be opened and read.

System action: The method terminates with a ConfigurationException.

Operator response: Ensure the properties file is not corrupt and has the appropriate read access rights. Contact IBM support if the problem persists.

EEZI0031E Connector exception detected in class: *className*, method: *methodName*. The content is: *exceptionDetails*. A Connection object could not be allocated.

Explanation: The call to getConnection() returned with an exception that is not attributable to an internal application exception.

System action: The method terminates with a ResourceException.

Operator response: See the WebSphere logs and traces for further details. Contact IBM support if the problem persists.

EEZI0032E Connector exception detected in class: *className*, method: *methodName*. A ConnectionFactory object could not be allocated.

Explanation: The ManagedConnectionFactory of this JCA encountered an internal error. The ConnectionManager instance was null.

System action: The method terminates with a ConfigurationException.

Operator response: Ensure the properties file is not corrupt and has the appropriate read access rights. Contact IBM support if the problem persists.

EEZI0041E Internal error. Illegal parameter passed in method: *methodName*, in class: *className*.

Explanation: The parameter passed to this object was not initialized.

System action: Method terminates with an IllegalArgumentException.

Operator response: Invoke this method with a fully initialized object as a valid parameter. Contact IBM support if the problem persists.

EEZI0042E Internal error. Illegal call to method
methodName, in class *className*.

Explanation: This method is specified and required by the J2C specification, but must not be called this way.

System action: Method terminates with an `IllegalOperationException`.

Operator response: Make sure this method is not called. Contact IBM support if the problem persists.

EEZI0044E Critical error in *methodName*, in class *className*. SSL problem. Property *property* is null.

Explanation: The SSL properties file could not be read correctly. One or more properties do not exist or are incorrect.

System action: The J2C Connector will fail to load and not be operational.

Operator response: Make sure all settings in the SSL properties file are correct and restart the server. Contact IBM support if the problem persists.

EEZI0046E Critical error in *methodName*, in class *className*. SSL problem.

Explanation: An SSL connection could not be established. One reason might be corrupt or incorrect SSL files.

System action: The current task ends.

Operator response: Make sure all settings in the SSL properties file are correct and that all SSL files are in the correct location and not corrupted. Contact IBM support if the problem persists.

EEZI0501W An exception was encountered and ignored in order to continue operation.
Exception string: *exceptionString*

Explanation: The invoked method is designed to ignore exceptions and continue operation. It logs the exception for problem determination purposes.

System action: Ignores the exception.

Operator response: Evaluate the exception details.

EEZI0545W Possible error in *methodName*, in class *className*. SSL problem. Property *property* equals null.

Explanation: The SSL properties file could not be read correctly. One or more properties do not exist or are incorrect.

System action: The J2C Connector will start, but will only be operational for non-SSL operations.

Operator response: Make sure all settings in the SSL

properties file are correct, and restart the server if SSL operations are desired.

EEZJ0001E The WebSphere infrastructure has reported a severe error situation:
RuntimeException message

Explanation: The application was interrupted by a `RuntimeException` and cannot complete its task.

System action: The current task ends. The transaction is rolled back.

Operator response: Check the description of the error situation if it indicates that the server database or another subsystem is unavailable. If the problem persists, contact IBM support.

EEZJ0002E The WebSphere infrastructure has reported an error situation: *Exception message*

Explanation: The application was interrupted by an unexpected exception or error that is not a `RuntimeException`.

System action: The current task ends, but the database operations that have been performed already remain valid (no transaction rollback).

Operator response: Contact IBM support.

EEZJ0003E Operation *operationName* encountered a `FinderException` because automation domain *domainName* is unknown in the scope of the management server. The operation continues processing of the other automation domains.

Explanation: Possible causes of the problem are:
1) The automation domain name was incorrect.
2) The automation domain has been deleted in the meantime.

System action: The operation task ends as far as the indicated automation domain is concerned. The operation continues processing of the other automation domains.

Operator response: Refresh the list of existing automation domains and verify that the domain name is contained in the list of existing domains. If not, and if the domain still exists and participates in automation, then restart the end-to-end automation adapter for this domain.

EEZJ0004E Expected a nonempty list of input data but received none in class: *className*, method: *methodName*, parameter: *parameterName*

Explanation: A null or empty list parameter was

encountered. This is an indication of a programming error on the EJB client side.

System action: The server method ends without processing the request.

Operator response: Contact IBM support.

EEZJ0005E Expected nonempty input but received no input in class: *className*, method: *methodName*, parameter: *parameterName*

Explanation: A parameter with a null value was encountered. This is an indication of a programming error on the EJB client side.

System action: The server method ends without processing the request.

Operator response: Contact IBM support.

EEZJ0006E Domain type *domainType* of automation domain *domainName* is unknown.

Explanation: The domain type of an automation domain is unknown.

System action: The server method ends without processing the request.

Operator response: Collect trace data and contact IBM support.

EEZJ0007E Within the list of resource requests, a request was encountered that contains a null or empty automation domain name.

Explanation: One of the requests within the parameter list contains a null or empty automation domain name.

System action: All requests in the list are ignored.

Operator response: Retry the operation. If the problem persists, trace the application to generate more detailed data about the erroneous operation, and contact IBM support.

EEZJ0008E The application is unable to connect to JMS topic *topicName*. Topic connection factory is *topicConnectionFactoryName*.

Explanation: JMS method `TopicConnectionFactory.createTopicConnection()` failed.

System action: The session bean cannot publish any messages to the topic. This may result in a loss of event data.

Operator response: Retry the operation. If the problem persists, trace the application to generate more detailed data about the erroneous operation, and contact IBM support.

EEZJ0009E Within the list of resource requests for automation domain *firstDomainName*, a request was encountered for automation domain *differentDomainName*

Explanation: Request lists must contain requests against a single automation domain only. The request list that causes the problem contains requests against multiple automation domains.

System action: All requests in the list are ignored.

Operator response: Retry the operation. If the problem persists, trace the application to generate more detailed data about the erroneous operation, and contact IBM support.

EEZJ0010E The `EEZDomainNameList` parameter received in class: *className*, method: *methodName* contains an element that is not a string.

Explanation: An incorrect parameter value was detected. This is an indication of a programming error on the EJB client side.

System action: The method ends but the session continues to exist.

Operator response: Contact IBM support.

EEZJ0011E The subscription method *methodName* in class *className* was called before the subscriber id was set in the session.

Explanation: Before a subscribe or unsubscribe method can be called, the subscriber id must be set within the session. This is an indication of a programming error on the EJB client side.

System action: The method ends but the session continues to exist.

Operator response: Restart the application that failed and retry the operation. If the problem persists, contact IBM support.

EEZJ0012E An attempt to activate automation policy *policyName* within automation domain *domainName* has been detected. This domain is not an end-to-end automation domain and does not support policy activation via System Automation interfaces.

Explanation: A caller tried to activate an automation policy on a domain instance that does not support policy activation.

System action: The policy activation request is ignored.

Operator response: Verify that the selected domain supports policy activation via the System Automation

interfaces. If it does not, use the automation product's native policy activation mechanism instead.

EEZJ0013E Subscriber *subscriberId* was unable to unsubscribe from some resources in domain *domainName* because the automation domain is not accessible at this time.

Explanation: The automation domain is currently not accessible, so the unsubscribe request could not be forwarded to the domain. However, the subscription cleanup within the management server was successful. Appropriate cleanup mechanisms in the domain (at domain adapter startup, for example) will take care of the orphaned subscription at the domain level.

System action: The unsubscribe operation continues to unsubscribe from resources that reside within other automation domains.

Operator response: Determine why the automation domain is not accessible at this time. If necessary, restart the end-to-end automation adapter for that domain in order to trigger resynchronization. If the domain has left, no further action is required.

EEZJ0014E Subscriber *subscriberId* was unable to unsubscribe from all resources in automation domain *domainName* because the domain is not accessible at this time.

Explanation: The automation domain is currently not accessible, so the unsubscribe request could not be forwarded to the domain. However, the subscription cleanup within the management server was successful. Appropriate cleanup mechanisms in the domain (at domain adapter startup, for example) will take care of the orphaned subscription at the domain level.

System action: The unsubscribe operation continues to unsubscribe from all resources that the subscriber has subscribed to previously and that reside within domains other than the failing one.

Operator response: Determine why the automation domain is not accessible at this time. If necessary, restart the end-to-end automation adapter for that domain in order to trigger resynchronization. If the domain has left, no further action is required.

EEZJ0015E An attempt to invoke operation *methodName* within automation domain *domainName* has been detected. The type of this domain does not support the requested operation.

Explanation: A caller tried to invoke an operation that is not supported.

System action: The operation request is ignored.

Operator response: Restart the application that failed

and retry the operation. If the problem persists, contact IBM support.

EEZJ0016E Unable to create an initial context.

Explanation: The JNDI naming directory is not accessible, and the attempt to create an initial context failed.

System action: The current task ends.

Operator response: Restart the application that logged this message. If this does not solve the problem, restart the WebSphere Application Server that provides the runtime environment for the automation manager. If the problem persists, contact IBM support.

EEZJ0017E Looking up object *jndiLookupName* in JNDI failed.

Explanation: Possible causes of the problem are:
1) The JNDI naming directory is not accessible.
2) The object was not bound to the JNDI correctly.

System action: The current task ends.

Operator response: Restart the WebSphere Application Server that provides the runtime environment for the automation manager. If the problem persists, contact IBM support.

EEZJ0018E Automation domain *domainName* does not exist.

Explanation: Possible causes of the problem are:
1) An invalid automation domain name was supplied.
2) The automation domain has been deleted in the meantime.

System action: The current task ends.

Operator response: Retry the operation. If the problem persists, trace the application to generate more detailed data about the erroneous operation, and contact IBM support.

EEZJ0019E Automation domain *domainName* is not accessible at this time.

Explanation: The automation domain exists, but it is currently not possible to communicate with it.

System action: The current task ends.

Operator response: Make sure that the automation domain is running. If it is a first-level automation domain, verify that the automation adapter is running. Retry the operation after the timeout period defined by the environment variable `com.ibm.eez.aab.watchdog-interval-seconds`. If the problem persists, restart the automation adapter (in case of a first-level automation domain) or the end-to-end automation engine (in case of an end-to-end automation domain).

EEZJ0020E Automation domain *domainName* seems to be not accessible at this time. Invocation of method *methodName* failed with a RemoteException.

Explanation: The automation domain exists, but it is currently not possible to communicate with it.

System action: The current task ends.

Operator response: Make sure that the automation domain is running. If it is a first-level automation domain, verify that the automation adapter is running. Retry the operation after the timeout period defined by the environment variable `com.ibm.eez.aab.watchdog-interval-seconds`. If the problem persists, restart the automation adapter (in case of a first-level automation domain) or the end-to-end automation engine (in case of an end-to-end automation domain).

EEZJ0021E Automation domain *domainName* cannot be accessed because of a problem within the J2EE framework.

Explanation: An attempt to create a session failed within the J2EE framework.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0022E An unrecoverable error occurred during startup of application *productName*. The application stops. Details about the error: *exceptionDetails*.

Explanation: An exception was encountered.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0023E An attempt to activate policy *policyName* in automation domain *domainName* resulted in an error which indicates that the policy is invalid.

Explanation: The automation domain indicates that an error was detected while processing the specified end-to-end automation policy.

System action: The current task ends.

Operator response: Verify the correctness of the end-to-end automation policy, and activate it from the operations console user interface.

EEZJ0024E An attempt to activate policy *policyName* in automation domain *domainName* resulted in an error which indicates that the policy cannot be found.

Explanation: The automation domain indicates that

the specified end-to-end automation policy cannot be found in the file system.

System action: The current task ends.

Operator response: Verify that the end-to-end automation policy file exists and contains a valid policy. Then activate it from the operations console user interface.

EEZJ0025E The operation `setPreferredMember` has ended since the automation domain name specified by the choice group key: *choiceGroupDomainName* did not match the domain name specified by the preferred member key: *preferredMemberDomainName*

Explanation: The resource keys that were provided do not point to the same automation domain. It is necessary, however, that the choice group and its members reside within the same domain.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0026E Operation *operation name* is not supported by class *class name*.

Explanation: A caller tried to invoke an operation that is not supported.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0027E An attempt to deactivate the currently active automation policy within automation domain *domainName* has been detected. This domain is not an end-to-end automation domain and does not support policy deactivation via System Automation interfaces.

Explanation: A caller tried to deactivate an end-to-end automation policy on an automation domain instance that does not support policy management via System Automation interfaces.

System action: The policy deactivation request is ignored.

Operator response: Verify that the selected automation domain supports automation policy deactivation via the System Automation interfaces. If it does not, use the automation product's native policy deactivation mechanism instead.

EEZJ0028E An attempt to retrieve an emitter object from the CEI EmitterFactory with JNDI name *emitterFactoryJndiName* failed with exception *exceptionString*.

Explanation: Although it was possible to retrieve an EmitterFactory from JNDI, it is not possible to create an emitter from that EmitterFactory.

System action: The current task ends.

Operator response: Verify that the CEI application is started. If the problem persists, restart the CEI application.

EEZJ0029E An attempt to publish an event was stopped since there is an active transaction. Event automation domain name is *domainName* and event reason is *eventReason*.

Explanation: The application does not support sending of JMS messages within a transactional boundary.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0030E The end-to-end automation manager is not fully initialized and refuses to accept requests.

Explanation: The EEZEAR application is either starting or stopping. During these periods, no method requests are accepted.

System action: The current task ends.

Operator response: If the EEZEAR application is starting, retry the request. If the EEZEAR application is stopping, restart the application and retry the request.

EEZJ0031E Refused to invoke operation *methodName* on end-to-end automation domain *domainName* because the caller is not in the EEZEndToEndAccess role.

Explanation: The target of this operation is an end-to-end automation domain. This operation may be invoked against end-to-end automation domains only by operators that are in the EEZEndToEndAccess role.

System action: The operation request is ignored.

Operator response: If the operator is not allowed to invoke operations against end-to-end resources, no action is required. If the operator should be allowed to invoke operations against end-to-end resources, the operator's userid or a user group that contains the operator's userid has to be added to role EEZEndToEndAccess.

EEZJ0032E Within the list of resource keys for automation domain *firstDomainName*, a resource key was encountered for automation domain *differentDomainName*

Explanation: In the context of this operation, each element of the list of resource keys must point to the same automation domain. This condition is not satisfied.

System action: The current task ends.

Operator response: Retry the operation. If the problem persists, trace the application to generate more detailed data about the erroneous operation, and contact IBM support.

EEZJ0033E Automation domain *domainName* requires user authentication.

Explanation: The automation domain requires that authentication information be supplied for each task. The authentication information consists of a userid and a password. The failing task did not supply that information.

System action: The current task ends.

Operator response: Case 1: If user authentication checking is enabled in the automation domain, ensure that user credential information for the automation domain is supplied. If the failing task was invoked from the TSA operations console, the user credential information must be supplied using the user preferences managed by the TSA operations console. If the failing task was invoked from the end-to-end automation engine, the user credential information must be supplied in the automation domain identification file used by the end-to-end automation engine. Case 2: If user authentication checking has been disabled in the automation domain, restart the adapter for that automation domain.

EEZJ0034E You are not authorized to perform the operation.

Explanation: The authorization failed while accessing the end-to-end automation manager.

System action: The requested operation is cancelled.

Operator response: Ensure that the permissions and user roles defined in the WebSphere Application Server are set up correctly. If the problem persists, contact your system administrator.

EEZJ0035E You are not authorized to perform the operation.
error details.

Explanation: The authorization failed while accessing the end-to-end automation manager.

System action: The requested operation is cancelled.

Operator response: Ensure that the permissions and user roles defined in the WebSphere Application Server are set up correctly. If the problem persists, contact your system administrator.

EEZJ0036E A WebSphere user transaction with an unexpected status was encountered while operation *operationName* was processed. The expected status is *expectedStatus* but the actual status is *actualStatus*.

Explanation: In the process of using a WebSphere user transaction, an unexpected transaction state was encountered.

System action: The current task ends.

Operator response: Retry the operation. If the problem persists, restart the WebSphere Application Server.

EEZJ0100E The processing of an event resulted in an exception: *exceptionDetails*

Explanation: The EventHandlerBean received an exception when processing an event.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0101E Cannot create or use an FLA connection to domain *domainName*. Details about the exception: *exceptionDetails*.

Explanation: The EventHandlerBean received an exception when processing an AdapterJoin event. It was not able to create or use a connection to a first-level automation domain.

System action: The processing of the AdapterJoin event ends.

Operator response: Contact IBM support.

EEZJ0102E Not able to add a subdomain to the domain *domainName*. Details about the exception: *exception*.

Explanation: The EventHandlerBean tried to locate this automation domain, but it received an exception. Therefore it is not able to add a subdomain to this automation domain.

System action: The current task ends but event processing continues.

Operator response: Contact IBM support.

EEZJ0103E Encountered a FinderException for the domain *domainName*.

Explanation: The EventHandlerBean tried to locate this automation domain, but it received a FinderException, because the automation domain is unknown in the scope of the end-to-end automation management server.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0104E Received an exception related to a transaction when processing an event of domain *domainName*. Details about the exception: *exception*.

Explanation: The transaction that was started when processing an event resulted in an exception.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0105E Not able to communicate with automation domain *domainName*. Details about the exception: *exception*.

Explanation: The EventHandlerBean received a domain join event of an automation domain, but it was not able to communicate with this automation domain. An exception was thrown instead.

System action: The processing of the domain join event ends.

Operator response: Contact IBM support.

EEZJ0106E Received a CreateException trying to create a domain for the domain name *domainName*.

Explanation: The EventHandlerBean received a CreateException while trying to create an automation domain object.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0107E Forwarding an event to the end-to-end automation domain *domainName* failed. Details about the exception: *exception*.

Explanation: The EventHandlerBean tried to forward an event to the automation engine. This operation failed.

System action: The current task ends. But the event processing continues.

Operator response: Contact IBM support.

EEZJ0108E **Activating policy *policyName* failed.**
Details about the exception: *exception*

Explanation: The EventHandlerBean tried to activate an end-to-end automation policy on an automation engine. This operation failed.

System action: The current task ends. But the event processing continues.

Operator response: Try to activate the policy using the operations console. If the problem persists, contact IBM support.

EEZJ0109E **Resynchronizing the end-to-end automation domain *domainName* failed.**
Details about the exception: *exception*.

Explanation: The EventHandlerBean tried to resynchronize the automation engine. This operation failed.

System action: The current task ends. But the event processing continues.

Operator response: Contact IBM support.

EEZJ0110E **FinderException received while trying to find subscriptions for entity *entityName*.**

Explanation: The EventHandlerBean tried to find subscriptions for this entity, but it received a FinderException.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0111E **CreateException received while trying to create a connection to the end-to-end automation domain *domainName*.**

Explanation: The EventHandlerBean received a CreateException while trying to create a connection to the automation engine.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0112E **RemoteException received when communicating with the end-to-end automation domain *domainName*.**

Explanation: The EventHandlerBean received a RemoteException when it called a function of the automation engine.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0113E **Calling checkHealth returned a null object for domain *domainName*.**

Explanation: The EventHandlerBean received a null object when calling checkHealth for an automation domain that just sent a domain join event. The domain join processing failed for this automation domain.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0114E **The domain object returned by checkHealth has a different domain name than the according domain join event. The event domain name is *domainName*.**

Explanation: The EventHandlerBean received an incorrect object from checkHealth. The domain join processing failed for this automation domain.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0115E **Exception received while trying to publish an event. Details about the exception: *exception details*.**

Explanation: The EventHandlerBean received an exception when it tried to publish an event.

System action: Processing continues.

Operator response: Contact IBM support.

EEZJ0501W **An exception was encountered and ignored in order to continue operation. Details about the exception: *exceptionString***

Explanation: The invoked method is designed to ignore exceptions and continue operation. It logs the exception for problem determination purposes.

System action: Processing continues.

Operator response: Evaluate the exception details.

EEZJ0502W **Lookup of CEI EmitterFactory object *jndiLookupName* in JNDI failed. This indicates that the CEI service is not installed or that it is not active.**

Explanation: Possible causes of the problem are:
1) JNDI is not accessible.
2) The object was not bound to the JNDI correctly.
In case of CEI this can occur when the CEI service has been disabled or CEI is not installed at all.

System action: The current task ends.

Operator response: Make sure the CEI application is in the desired state. If events have to be forwarded to

TEC, the CEI application must be running.

EEZJ0503W The conversion to CBE event format failed.

Explanation: The resulting CBE event is not valid.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0504W The emitter is unable to send the event to CEI.

Explanation: The emitter cannot deliver the event to the CEI event server.

System action: The current task ends.

Operator response: Ensure that the CEI application is installed within this server and that it is running.

EEZJ0505W Unable to filter event: *event*.

Explanation: The filter is unable to filter the specified event.

Possible causes of the problem are:

- 1) The event is not a valid XML document.
- 2) The filter configuration does not match this event.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0506W Cannot load the filter configuration file *filter filename*.

Explanation: The loading of the filter configuration XML or XSLT from the EAR failed.

Possible causes of the problem are:

- 1) The file is not a valid XML file.
- 2) The file cannot be found.
- 3) Parsing of the filter file failed.

System action: The current task ends.

Operator response: Contact IBM support.

EEZJ0507W Unable to close the CEI emitter.

Explanation: The emitter fails to release held resources.

System action: The current task ends.

Operator response: If the problem persists, contact IBM support.

EEZJ0508W The CEI application is not running.

Explanation: Since the CEI application is not running, no events can be sent to TEC.

System action: The current task ends. An event is not sent to TEC.

Operator response: Make sure the CEI application is in the desired state. If events have to be forwarded to TEC, the CEI application must be running.

EEZJ0600W A RemoveException was received while trying to remove an entity from the database when processing an event received from automation domain *domainName*.

Explanation: The EventHandlerBean received a RemoveException while trying to remove an entity after processing an event.

System action: Processing continues.

Operator response: Evaluate the exception details.

EEZJ0601W The policy name stored in the J2EE framework and the policy name supplied by a policy changed event are not equal. The policy name stored in the J2EE framework is *aab policyName*. The policy name supplied by the event is *event policyName*.

Explanation: The J2EE framework received a policy changed event that contains a policy name that does not match the policy name that was stored previously in the J2EE framework.

System action: Processing continues.

Operator response: Verify that the policy names are set correctly. If necessary, activate the policy again.

EEZJ0602W Not able to communicate with automation domain *domainName*.

Explanation: The EventHandlerBean tried to communicate with an automation domain, but it received an exception.

System action: Processing continues.

Operator response: Evaluate the exception details.

EEZJ0603W Automation domain *oldDomainName* has left and automation domain *newDomainName* has joined. These domains have the same access data. Apparently the domain has been renamed.

Explanation: The EventHandlerBean received a domain join event. The access data of this event, such as the hostname and port, is the same as that of an existing automation domain with a different name. The EventHandlerBean created a new object for the automation domain that joined and will soon remove the object for the automation domain that left.

System action: Processing continues.

Operator response: Verify that the automation domain has not been renamed by mistake.

EEZJ1000I Application *productName* has started working.

Explanation: The application starts its asynchronous work.

System action: No system action required.

Operator response: No operator action required.

EEZJ1001I Application *productName* was shut down by the J2EE container and has stopped working.

Explanation: The application stops its asynchronous work.

System action: No system action required.

Operator response: If required, restart the application.

EEZJ1002I Domain *domainName* has been inactive for a long period of time and has been removed from the automation scope.

Explanation: The timeout defined by the environment variable `com.ibm.eez.domain-removal-hours` has been reached for this automation domain.

System action: No system action required.

Operator response: No operator action required. When the automation domain that has been removed from the automation scope joins the automation scope again, it is recreated.

EEZJ1003I The communication state of automation domain *domainName* has changed from *previousCommState* to *newCommState*.

Explanation: The communication health state has changed.

System action: The system publishes a related event.

Operator response: Depending on the current state values and the desired communication state of the automation domain, it might be necessary to restart the automation adapter.

EEZJ1004I Config: The timeout for backend automation calls is *timeoutValue* seconds.

Explanation: Controls how many seconds each call to the backend may take at most. Default: 60, minimum: 30, maximum: 3600.

System action: No system action required.

Operator response: If needed, set or modify the environment variable `com.ibm.eez.aab.invocation-timeout-seconds`.

EEZJ1005I Config: The timeout to determine domain communication health state is *timeoutValue* seconds.

Explanation: Controls the number of seconds of inactivity after which the health of the communication to the automation domain is checked automatically. Default: 300, minimum: 60, maximum: 86400.

System action: No system action required.

Operator response: If needed, set or modify the environment variable `com.ibm.eez.aab.watchdog-interval-seconds`.

EEZJ1006I Config: The timeout before removing domains that have left is *timeoutValue* hours.

Explanation: Controls the number of hours of inactivity after which the automation domain's representation in the management server is removed automatically. Default: 48, minimum: 1, maximum: 1000.

System action: No system action required.

Operator response: If needed, set or modify the environment variable `com.ibm.eez.aab.domain-removal-hours`.

EEZJ1007I The CEI application is not installed. It is therefore not possible to send any events to the IBM Tivoli Enterprise Console (TEC). The original error message is: *errorMessage*

Explanation: CEI is not, or not properly, installed within the scope of this server.

System action: The application will not try again to send events to TEC.

Operator response: If the application is not supposed to send events to TEC, there is no action required. If the application is supposed to send events to TEC, ensure that the CEI application is installed within this WebSphere Application Server and that it is running. When the problem with CEI is resolved, you need to restart the WebSphere Application Server before any System Automation events can be sent to TEC.

EEZJ1008I The domain state of domain *domainName* has changed from *previousDomainState* to *newDomainState*

Explanation: The state of the automation domain has changed.

System action: The system publishes a related event.

Operator response: Depending on the current state values and the desired state of the automation domain, it might be necessary to restart the domain.

EEZJ1009I The CEI application is not installed. It is therefore not possible to send any events to the IBM Tivoli Enterprise Console (TEC).

Explanation: CEI is not, or not properly, installed within the scope of this server.

System action: The application will not try again to send events to TEC.

Operator response: If the application is not supposed to send events to TEC, there is no action required. If the application is supposed to send events to TEC, ensure that the CEI application is installed within this WebSphere Application Server and that it is running. When the problem with CEI is resolved, you need to restart the WebSphere Application Server before any System Automation events can be sent to TEC.

EEZJ1100I Some data of domain *domainName* has changed.

Explanation: The domain join event of the automation domain contains different data than the domain object. The domain object will be updated with the data of the event by the EventHandlerBean.

System action: Processing continues with the updated domain object.

Operator response: None.

EEZK0003E String *someString* is too long: the maximum length of *nameOfTheString* Strings is *maxLength*.

Explanation: Setting the string to the specified value did not succeed due to string length.

System action: The current task ends.

Operator response: Verify the input parameters. If the problem persists, provide the trace file and this message text to IBM support.

EEZK0004E String named *someStringName* must not be null and must not exceed the maximum length of *maxLength*.

Explanation: Setting the String to null is not allowed.

System action: The current task ends.

Operator response: Verify the input parameters. If the problem persists, provide the trace file and this message text to IBM support.

EEZK0005E An exception that is not an instance of **EEZApplicationException** has been passed to the **EEZApplicationTransientException**. The type of the message is *exceptionType*.

The exception message is:
exceptionMessage.

Explanation: This is an unexpected behavior.

System action: The current task will continue. The exception will be processed.

Operator response: If any other error occurs, please provide the logs and traces as an aid to analysis.

EEZL0001E The WebSphere infrastructure has reported a severe error situation:
runtimeExceptionMessage

Explanation: The application was interrupted by a RuntimeException and cannot complete its task.

System action: The current task ends. The transaction is rolled back.

Operator response: Check the description of the error situation if it indicates that the server database or another subsystem is unavailable. If the problem persists, contact IBM support.

EEZL0002E The WebSphere infrastructure has reported an error situation:
exceptionMessage

Explanation: The application was interrupted by an unexpected exception or error that is not a RuntimeException.

System action: The current task ends, but the database operations that have been performed already remain valid (no transaction rollback).

Operator response: Contact IBM support.

EEZL0003E A critical error has occurred in class:
className, method: *methodName*. The logger object could not be initialized.

Explanation: This component could not initialize and access a logger object. This indicates either a configuration or programming error.

System action: The process cannot be completed. All parts of this component are affected. The system is not operational.

Operator response: Please check that the path settings are correct and all required libraries exist. Collect trace data and this message text and contact IBM support if the problem persists.

EEZL0004E An error has occurred in class:
className, method: *methodName*,
parameter: *parameterName*

Explanation: The method has been invoked with an empty or null parameter list. The method must be invoked with a parameter list that

is not null and filled.

This indicates a programming error.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0005E An error has occurred in class:
className, **method:** *methodName*,
parameter: *parameterName*

Explanation: The method has been invoked with a null parameter.

The method must be invoked with a parameter that is not null.

This indicates a programming error.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0015E An error has occurred in class:
className.

Explanation: Configuration data object is null.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0016E An error has occurred in class:
className.

Explanation: First-level automation name has not been set.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0017E An error has occurred in class:
className.

Explanation: Host address has not been set.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0018E An error has occurred in class:
className.

Explanation: Adapter plugin class has not been set.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0019E An error has occurred in class:
className.

Explanation: Port has not been set.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0020E An error has occurred in class:
className.

Explanation: Timeout value has not been set.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0021E An error has occurred in class:
className.

Explanation: User Credentials object is null.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0022E An error has occurred in class:
className.

Explanation: Username has not been set.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0023E An error has occurred in class:
className.

Explanation: Password has not been set.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0024E An error has occurred in class:
className, **method:** *methodName*. **Illegal return object.**

Explanation: The JCA has returned an illegal argument to the EJB, which has caused a `ClassCastException`.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0025E An error has occurred in class: *className*, method: *methodName*. Illegal parameter at invocation of this method.

Explanation: The method has been invoked with a null parameter.
The method must be invoked with a parameter that is not null.
This indicates a programming error.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0030E A exception has occurred in class: *className*, method: *methodName*. The nested exception is null.

Explanation: No exception object was linked to the ResourceException that has been caught.
This is an unexpected behavior and indicates a programming error on the J2C side.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0031E An error has occurred in class: *className*, method: *methodName*. Invalid nested exception: *nestedException*.

Explanation: An invalid exception object was linked to the ResourceException that has been caught.
This is an unexpected behavior and indicates a programming error on the J2C side.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0032E An error has occurred in class: *className*, method: *methodName*. No Connection object could be obtained.

Explanation: The call to EEZConnectionFactory.getConnection(..) returned null.
This is an unexpected behavior and indicates a programming error at J2C side.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0033E An error has occurred in class: *className*, method: *methodName*. No Interaction object could be obtained.

Explanation: The call to EEZConnection.createInteraction() returned null.
This is an unexpected behavior and indicates a

programming error at J2C side.

System action: The current task ends.

Operator response: Please collect logs and traces and provide them to IBM support.

EEZL0034E An error has occurred in class: *className*, method: *methodName*. JNDI name: *jndiName* did not return a ConnectionFactory object.

Explanation: The JNDI lookup of this J2C has encountered an internal error. The ConnectionFactory object could not be retrieved.

This indicates a JNDI configuration error.

System action: The current task ends. No connection to the first-level automation will be possible until this problem is fixed.

Operator response: Please ensure the JNDI settings for the J2C connection factories are correct and restart the server. Contact IBM support if the problem persists.

EEZL0040E Error occurred during XML (de)serialization process.
Exception: *exception* detected in *className* method *methodName*.

Explanation: The XML decoder has received an XML string that contained unsupported encoding.

System action: The method terminates with an ExecutionFailedException.

Operator response: Please provide logs and traces and contact IBM support if the problem persists.

EEZL0501W An exception was encountered and ignored in order to continue operation.
Exception string: *exceptionString*

Explanation: The invoked method is designed to ignore exceptions and continue operation. It logs the exception for problem determination purposes.

System action: Ignores the exception.

Operator response: Evaluate the exception details.

EEZL0510W An exception was encountered at XML serialization in class *className* method *methodName*. **Exception string:** *exceptionDetails*

Explanation: This might be subject to back-level toleration and can be ignored.

System action: The exception is ignored. The process will be continued.

Operator response: Evaluate the exception details. If the problem persists or causes further errors, please contact IBM support.

EEZN0010E The connection to the LDAP server on port *port* on host *host* cannot be established.

Explanation: The application was interrupted by a RuntimeException and cannot complete its task.

System action: The connection process ends abnormally.

Operator response: Check whether the LDAP server is running and configured to the correct port and retry if applicable.

EEZN0011E Host *host* could not be found.

Explanation: The connection to the LDAP server cannot be established. Possible reasons are that the server is off-line, the name is incorrect, or a network failure.

System action: The connection process ends abnormally.

Operator response: Ensure that the host name or IP address is specified correctly. Verify that the network and the LDAP server are operational and retry if applicable.

EEZN0012E Either the user name *username* or the password was not accepted by the LDAP server.

Explanation: Either the user name or the password or both were incorrect.

System action: Access denied.

Operator response: Retry with a correct user name and password.

EEZN0013E An error occurred during logon to the LDAP server: *error* or the password was not accepted by the LDAP server. The contacted host was *host* with port address *port*. You tried to log in as user *username*.

Explanation: An attempt to log on to the LDAP server resulted in an unknown error. One possible error is an incorrect password.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and retry.

EEZN0201E An exception was encountered. Details about the exception: *exceptionString*. Attempted to create CEI Response File *ceirsp*.

Explanation: An attempt to create the CEI Response File resulted in a "FileNotFoundException" error. One possible

error is incorrect user rights to create files on the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try it again.

EEZN0202E An exception was encountered. Details about the exception: *exceptionString*. Attempted to find or open the CEI Response File *ceirsp*.

Explanation: An attempt to find or open the CEI Response File to store all the CEI properties resulted in an "Input/Output" error. One possible error is an incorrect file path or incorrect user rights to open this file on the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try again.

EEZN0301E ISCPreqProcess has returned the value: *iscPrereqRC*.

Explanation: The ISC Prerequisite Check Process returned an error code.

System action: No system action required.

Operator response: No operator action required.

EEZN0301I ISC Portal Server is already started!!!!

Explanation: No explanation.

System action: No system action required.

Operator response: No operator action required.

EEZN0302E An exception was encountered. Details about the exception: *errorString*. The ISC command executed by the following system command *javaCmd* was not successful.

Explanation: No explanation.

System action: No system action required.

Operator response: No operator action required.

EEZN0302I Output of the following *cmd* is: *line*.

Explanation: No explanation.

System action: No system action required.

Operator response: No operator action required.

EEZN0303E An exception was encountered. When trying to get the "BufferedReader-InputStream" for the following command: " *cmd* " the result was "null".

Explanation: One possible error is an incorrect file path or incorrect user rights to write to this file on the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try again.

EEZN0304E An exception was encountered. When trying to get the "BufferedReader-ErrorStream" for the following command: " *cmd* " the result was "null".

Explanation: One possible error is an incorrect file path or incorrect user rights to write to this file on the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try again.

EEZN0401E An exception was encountered. Details about the exception: *exceptionString*.

Explanation: The application is not able to get the InstallShield database from the services of ISMP.

System action: No system action required.

Operator response: No operator action required.

EEZN0402E An exception was encountered. Details about the exception: *exceptionString*. The application is not able to set the following InstallShield variable *isVariable* in the InstallShield database.

Explanation: No explanation

System action: No system action required.

Operator response: No operator action required.

EEZN0403E An exception was encountered. Details about the exception: *exceptionString*.

Explanation: Error: cannot get registry service to read vpd.properties.

System action: No system action required.

Operator response: No operator action required.

EEZN0404E An exception was encountered. Details about the exception: *exceptionString*.

Explanation: The operation performed by ISMP FileService fails.

System action: No system action required.

Operator response: No operator action required.

EEZN0500W An exception was encountered and ignored in order to continue operation. Details about the exception: *exceptionString*

Explanation: The invoked method is designed to ignore exceptions and continue operation. It logs the exception for problem determination purposes.

System action: Processing continues.

Operator response: Evaluate the exception details.

EEZN0501E An exception was encountered. Details about the exception: *exceptionString*. An attempt to execute a specified system command via Java resulted in an "Input/Output" error. The Java error was: *exceptionString*. Attempted to execute the specified system command *javaCmd*.

Explanation: One possible error is incorrect user rights to execute system commands on the operating system. Another possible reason is that the attributes of the command are incorrect.

System action: No system action required.

Operator response: Use the information given in the explanation text to correct the error and try again.

EEZN0501I Found text: *findText* in file: *fileName*.

Explanation: No explanation.

System action: No system action required.

Operator response: No operator action required.

EEZN0502E An exception was encountered. Details about the exception: *exceptionString*. The current thread initiated by the following system command *javaCmd* was interrupted by another thread while it was waiting, then the wait terminated and an "InterruptedException" was thrown.

Explanation: No explanation

System action: No system action required.

Operator response: No operator action required.

EEZN0503E **Unsupported platform. Within the beanID = *beanId*, platform *platform* is not supported.**

Explanation: No explanation

System action: No system action required.

Operator response: No operator action required.

EEZN0504E **An exception was encountered. Details about the exception: *exceptionString*. An "Input/Output" error occurred trying to create the following log file or, if it exists, trying to find its directory path: *EEZLogFile*.**

Explanation: One possible error is an incorrect file path or incorrect user rights to open this file on the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try again.

EEZN0505E **An exception was encountered. Details about the exception: *exceptionString*. An "Input/Output" error occurred trying to write to the following log file: *EEZLogFile*.**

Explanation: One possible error is an incorrect file path or incorrect user rights to write to this file on the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try again.

EEZN0506E **An exception was encountered. Details about the exception: *exceptionString*. An attempt to read a line from an "Inputstream" into a from a "BufferedReader" resulted in a "Input/Output" error.**

Explanation: One possible error is an incorrect "Inputstream" from the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try it again.

EEZN0507E **An exception was encountered. Cannot find the following DepcheckWizardAction bean *depcheckWzrdActID*.**

Explanation: No explanation.

System action: No system action required.

Operator response: No operator action required.

EEZN0508E **An exception was encountered. The product " *prodName* " does not exist.**

Explanation: No explanation.

System action: No system action required.

Operator response: No operator action required.

EEZN0509E **An exception was encountered. Details about the exception: *exceptionString*.**

Explanation: An "Input/Output" error occurred trying to store properties.

System action: No system action required.

Operator response: No operator action required.

EEZN0510E **An exception was encountered. Details about the exception: *exceptionString*. An "Input/Output" error occurred trying to get the URL to the " *rsrcName* " resource within the application archive.**

Explanation: No explanation.

System action: No system action required.

Operator response: No operator action required.

EEZN0511E **An exception was encountered. Details about the exception: *exceptionString*.**

Explanation: An "Input/Output" error occurred trying to load properties.

System action: No system action required.

Operator response: No operator action required.

EEZN0512E **An exception was encountered. Details about the exception: *exceptionString*. An attempt to create a "FileInputStream" from the properties file *prop* resulted in a "FileNotFoundException" error.**

Explanation: One possible error is an incorrect file path or incorrect user rights to get this "FileInputStream" from the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try it again.

EEZN0513E **An error was found. The file " *fileName* " does not exist.**

Explanation: One possible error is an incorrect file path or incorrect user rights to open this file on the operating system.

System action: Access denied.

EEZN0514E • EEZP0002E

Operator response: Use the information given in the explanation text to correct the error and try it again.

EEZN0514E An exception was encountered. The log.properties file ("logPropFileName") does not exist.

Explanation: One possible error is an incorrect file path or incorrect user rights to open this file on the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try it again.

EEZN0515E An exception was encountered. The log.properties file ("logPropFileName") does not contain the property prop.

Explanation: No explanation.

System action: No system action required.

Operator response: Use the information given in the explanation text to correct the error and try it again.

EEZN0516E An exception was encountered. The log.properties file ("logPropFileName") does not exist.

Explanation: One possible error is an incorrect file path or incorrect user rights to open this file on the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try it again.

EEZN0517E An exception was encountered. Details about the exception: *exceptionString*. Unable to resolve hostname ("hostName").

Explanation: You need to edit the config files manually.

System action: No system action required.

Operator response: You need to edit the config files manually.

EEZN0518E An exception was encountered. Details about the exception: *exceptionString*. An attempt to create a "FileReader" from the file "*fileName*" resulted in a "FileNotFoundException" error.

Explanation: One possible error is an incorrect file path or incorrect user rights to get this "FileReader" from the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try it again.

EEZN0519E An exception was encountered. Details about the exception: *exceptionString*. An "Input/Output" error occurred trying to read from the following file: *fileName*.

Explanation: One possible error is an incorrect file path or incorrect user rights to write to this file on the operating system.

System action: Access denied.

Operator response: Use the information given in the explanation text to correct the error and try again.

EEZN1000I Application *productName* has started working.

Explanation: The application starts its asynchronous work.

System action: No system action required.

Operator response: No operator action required.

EEZP0001E The specified <Source> "*source*" in the <Relationship> "*source*" "*relationshipType*" "*target*" does not exist as a <ResourceReference>, <ResourceGroup> or <ChoiceGroup>.

Explanation: The <Source> and <Target> of a <Relationship> must exist as exactly one <ResourceReference>, <ResourceGroup> or <ChoiceGroup>.

System action: This policy cannot be activated.

Operator response: Verify this <Relationship> in the policy.

EEZP0002E The specified <Target> "*target*" in the <Relationship> "*source*" "*relationshipType*" "*target*" does not exist as a <ResourceReference>, <ResourceGroup> or <ChoiceGroup>.

Explanation: The <Source> and <Target> of a <Relationship> must exist as exactly one <ResourceReference>, <ResourceGroup> or <ChoiceGroup>.

System action: This policy cannot be activated.

Operator response: Verify this <Relationship> in this policy.

EEZP0003E The specified *<policyElement>* name "*nameOfElement*" was found more than once as the name of a *<ResourceReference>*, *<ResourceGroup>* or *<ChoiceGroup>*.

Explanation: The value of the name attributes of *<ResourceReference>*, *<ResourceGroup>* and *<ChoiceGroup>* must be unique.

System action: This policy cannot be activated.

Operator response: Verify this name attribute in this policy.

EEZP0004E The specified member "*groupMember*" of the *<groupElement>* name "*groupName*" does not exist as a *<ResourceReference>*, *<ResourceGroup>* or *<ChoiceGroup>*.

Explanation: The member in a group must exist as exactly one *<ResourceReference>*, *<ResourceGroup>* or *<ChoiceGroup>*.

System action: This policy cannot be activated.

Operator response: Verify this member name in this policy.

EEZP0005E Syntax error in line *lineNumber* column *columnNumber*. Original parser exception: *errorMessage*

Explanation: A syntax error occurred while parsing this policy.

System action: This policy cannot be activated.

Operator response: Correct the syntax error in this policy.

EEZP0006E The specified policy file "*policyFile*" cannot be found.

Explanation: The policy cannot be loaded from this location.

System action: This policy cannot be activated.

Operator response: Verify the policy XML file name and its path.

EEZP0007E Original Parser Exception: *exceptionMessage*

Explanation: An internal problem occurred while parsing this policy.

System action: This policy cannot be activated.

Operator response: Verify that the product is correctly installed.

EEZP0008E An unsupported character *character* was found in the string "*completeString*". This string was found in the element *<elementName>* of the parent element *<parentElement>*.

Explanation: The character found in the string is not supported.

System action: This policy cannot be activated.

Operator response: Remove the unsupported character from this string in this policy.

EEZP0009E The specified name "*nameOfElements*" was found in the elements *<policyElement>* and *<otherPolicyElement>*.

Explanation: The value of the name attribute must be unique.

System action: This policy cannot be activated.

Operator response: Verify this name attribute in this policy.

EEZP0010E The specified *<ResourceReference>* "*referenceName*" was found as a member of multiple *<ChoiceGroup>* elements.

Explanation: A *<ResourceReference>* can only be a member of one *<ChoiceGroup>*.

System action: This policy cannot be activated.

Operator response: Check that the *<ResourceReference>* is a member of at most one *<ChoiceGroup>* element in this policy.

EEZP0011E The specified *<groupForm>* "*groupName*" was found as a member of multiple other groups.

Explanation: A group can only be a member of one group.

System action: This policy cannot be activated.

Operator response: Check that the group is a member of at most one group element in this policy.

EEZP0012E The two *<ResourceReference>* elements "*reference*" and "*otherReference*" point to the same *<ReferencedResource>* "*resource*".

Explanation: Only one *<ResourceReference>* can point to the same *<ReferencedResource>* child element values.

System action: This policy cannot be activated.

Operator response: Check that only one *<ResourceReference>* has the same *<ReferencedResource>* child element values for

<AutomationDomain>, <Name>, <Class> and <Node> in this policy.

EEZP0013E The specified member "*memberName*" was found multiple times in the same <groupForm> "*groupName*".

Explanation: All <Members> child elements must be unique in one group.

System action: This policy cannot be activated.

Operator response: Check that the group has no duplicate <Members> child elements in this policy.

EEZP0014E The specified <ResourceReference> "*reference*" was found as a member of the <ResourceGroup> "*resourceGroupName*" and the <ChoiceGroup> "*choiceGroupName*".

Explanation: A <ResourceReference> can only be a member of multiple <ResourceGroup> elements or one <ChoiceGroup> element.

System action: This policy cannot be activated.

Operator response: Check that the <ResourceReference> is not a member of a <ResourceGroup> and a <ChoiceGroup> at the same time in this policy.

EEZP0015E The specified <Relationship> <Type> "*relationType*" with <Source> "*source*" and <Target> "*target*" was found in a loop.

Explanation: <Relationship> elements of the same <Type> where one <Relationship> element <Target> is the next <Relationship> element <Source> must not form a loop.

System action: This policy cannot be activated.

Operator response: Check that the <Relationship> elements are not defined as a loop in this policy.

EEZP0016E The specified element <childElement> was found more than once as a child element of <parentElement> name "*parentName*".

Explanation: At most one element of this type is allowed in this group.

System action: This policy cannot be activated.

Operator response: Check that at most one element of this type is specified in this group in this policy.

EEZP0017E The specified element <parentElement> name "*parentName*" was found without <Members> child elements.

Explanation: At least one <Members> child element must be specified in this group.

System action: This policy cannot be activated.

Operator response: Check that at least one <Members> child element is specified in this group in this policy.

EEZP0018E The policy document does not contain a <ResourceReference> element.

Explanation: At least one <ResourceReference> element must be specified in this policy.

System action: This policy cannot be activated.

Operator response: Check that at least one <ResourceReference> element is specified in this policy.

EEZP0019E The specified element <ChoiceGroup> name "*groupName*" was found with more than one <Members> child element with the "preferred" attribute equal to "true".

Explanation: One <ChoiceGroup> member must have the "preferred" attribute equal to "true".

System action: This policy cannot be activated.

Operator response: Check that exactly one <ChoiceGroup> member has the "preferred" attribute equal to "true".

EEZP0020E The specified <Relationship> with the <Type> "*relationType*", the <Source> "*source*" and the <Target> "*target*" was found multiple times in the policy document.

Explanation: All <Relationship> elements must be unique.

System action: This policy cannot be activated.

Operator response: Check that at most one <Relationship> of this type is specified in this policy.

EEZP0021E A 'UTFDataFormatException' was caught in method *methodName* of class *className*. The received message was *message*.

Explanation: The processing was interrupted by this exception and cannot complete.

System action: The policy cannot be loaded.

Operator response: Ensure the correct data format of the policy document by only using editors which create

UTF-8-compliant documents.

EEZP0022E The specified *<groupType>* name "*groupName*" was found in a loop.

Explanation: Group elements cannot form a loop with their members.

System action: This policy cannot be activated.

Operator response: Check that the group *<Members>* child elements are not defined as a loop in this policy.

EEZP0023E The specified element *<ChoiceGroup>* name "*groupName*" has no *<Members>* child element with the "preferred" attribute equal to "true".

Explanation: One *<ChoiceGroup>* member must have the "preferred" attribute equal to "true".

System action: This policy cannot be activated.

Operator response: Check that exactly one *<ChoiceGroup>* member has the "preferred" attribute equal to "true".

EEZP0500W The specified member "*memberName*" of the *<ChoiceGroup>* "*choiceGroupName*" was also found as a *<Source>* or *<Target>* of a *<Relationship>*.

Explanation: The member of a *<ChoiceGroup>* should not be the *<Source>* or *<Target>* of a *<Relationship>* at the same time.

System action: Application continues.

Operator response: To avoid complexity, delete the *<Relationship>* or delete this *<ChoiceGroup>* member in this policy.

EEZP0502W The two *<Relationship>* elements with *<Type>* "StartAfter" and *<Type>* "StopAfter" were found with the same *<Source>* "*source*" and *<Target>* "*target*".

Explanation: The two *<Relationship>* elements with *<Type>* "StartAfter" and *<Type>* "StopAfter" should not have the same *<Source>* and *<Target>*. With this configuration the *<Target>* is started before the *<Source>* and the *<Target>* is stopped before the *<Source>*.

System action: Application continues.

Operator response: Verify this behavior. The common usage of "StartAfter" together with "StopAfter" is the following: 1. The *<Source>* of the "StartAfter" is the *<Target>* of the "StopAfter". 2. The *<Target>* of the "StartAfter" is the *<Source>* of the "StopAfter".

EEZP0503W The *<ResourceReference>* with name "*Reference Name*" has a *<DesiredState>* of Online, but is a member of the group with name "*Group Name*" which defines the desired state of all of its members.

Explanation: The *<DesiredState>* of this *<ResourceReference>* will be set to Offline, which is the default for a *<ResourceReference>* that is a member of a group.

System action: Application continues.

Operator response: To avoid this warning do not specify an Online desired state for a *<ResourceReference>* that is a member of a group.

EEZP0504W The group with name "*member group name*" has a *<DesiredState>* of Online, but is a member of another group with name "*hosting group name*" which defines the desired state of all of its members.

Explanation: The *<DesiredState>* of this group will be set to Offline, which is the default for a group that is a member of another group.

System action: Application continues.

Operator response: To avoid this warning do not specify an Online desired state for a group that is a member of another group.

EEZP2000I Usage:

Explanation: No additional information is available for this message.

EEZP2001I eezpolicychecker *<Policy XML file name with path>*

Explanation: No additional information is available for this message.

EEZP2002I (C) COPYRIGHT International Business Machines Corp. 2005
All Rights Reserved.

Explanation: No additional information is available for this message.

EEZP2003I Policy has been verified.

Explanation: No additional information is available for this message.

EEZP2004I Policy could not be verified.

Explanation: No additional information is available for this message.

EEZP2005I The following policy errors were found:

Explanation: No additional information is available for this message.

EEZP2006I The following policy warnings were found:

Explanation: No additional information is available for this message.

EEZP2007I ERROR:

Explanation: No additional information is available for this message.

EEZP2008I WARNING:

Explanation: No additional information is available for this message.

EEZP2009I POLICY CHECKER

Explanation: No additional information is available for this message.

EEZP2010I This program verifies a policy document (*.xml) for use with IBM Tivoli System Automation for Multiplatforms.

Explanation: No additional information is available for this message.

EEZP2011I EXPLANATION:

Explanation: No additional information is available for this message.

EEZP2012I USER ACTION:

Explanation: No additional information is available for this message.

EEZT0001W SA detected a domain status change in domain " *domainName* ". Domain status is *domainStatus*.

Explanation: The status of this automation domain has changed.

System action: None.

Operator response: Verify that the new state is intended for this automation domain.

EEZT0002I SA received a domain join event from domain " *domainName* " and processed that event successfully.

Explanation: This automation domain has successfully joined to System Automation.

System action: None.

Operator response: No action required.

EEZT0003W SA received a domain leave event from domain " *domainName* " and processed that event successfully.

Explanation: This automation domain has successfully left System Automation.

System action: None.

Operator response: Verify that the new state is intended for this automation domain.

EEZT0004E SA cannot communicate with the automation adapter on IP address " *ipAddress* ". For example, a firewall could be preventing that kind of communication.

Explanation: An automation adapter was started and it tried to join to System Automation. The join attempt failed because System Automation cannot communicate with the automation adapter. Possible cause of the problem is: a firewall could be preventing communication with this domain, or SSL communication is required to be setup for communication with the automation adapter.

System action: None.

Operator response: Determine the cause of the communication problem.

EEZT0005E SA encountered a security problem while trying to access domain " *domainName* ".

Explanation: System Automation cannot access this automation domain. This is caused by a security problem. Either the EEZAutomationManager does not have the required authorization, or some credential definition in the domain identification file that is used by the end-to-end automation engine is invalid. For example, the password may have expired.

System action: None.

Operator response: Browse the domain identification file used by the end-to-end automation engine. Check if the password that is associated with the domain mentioned in the message text is valid. If the password is valid, then check the message log files of the end-to-end automation engine and of the J2EE framework for additional hints about the problem.

EEZT0006I SA domain "*domainName*" has activated a new policy. The filename of the policy is "*policyFileName*".

Explanation: A new policy has been activated for this automation domain.

System action: None.

Operator response: Verify that the action is intended for this automation domain.

EEZT0007E SA domain "*domainName*" has encountered a severe problem. Check the domain log file for details.

Explanation: See message.

System action: None.

Operator response: Check the automation domain log file for details.

EEZT0008I SA domain "*domainName*" no longer requires someone to check its log file.

Explanation: An operator has inspected the logfile.

System action: None.

Operator response: None.

EEZT0009W SA detected a problem in domain "*domainName*". The resource reference "*resourceName*" is unknown.

Explanation: The resource reference definition in the end-to-end automation policy does not match any resource in the specified automation domain. Possible causes of this problem are: 1) the end-to-end automation policy is incorrect, possibly caused by a spelling error. 2) the resource does not or does not yet exist in the specified automation domain.

System action: None.

Operator response: Verify the correctness of the resource reference definition in the policy. Verify that the referenced resource exists in the automation domain.

EEZT0010I SA changed the preferred member of resource group "*groupName*" in domain "*domainName*" as requested.

Explanation: See message.

System action: None.

Operator response: None.

EEZT0011W SA detected a domain communication status change in domain "*domainName*". Domain communication status is *communicationState*. Domain status is *domainState*.

Explanation: The communication state for this automation domain has changed.

System action: None.

Operator response: Verify that the new state is intended for this automation domain.

EEZT0012W SA domain "*domainName*" no longer has an active policy.

Explanation: The automation policy has been deactivated for this automation domain.

System action: None.

Operator response: Verify that the action is intended for this automation domain.

EEZT0013E SA detected a problem in domain "*domainName*". The resource reference "*resourceName*" is unable to contact the referenced resource. An operator needs to reset the resource reference manually.

Explanation: The operational state of the resource reference indicates that an operator needs to fix some problem and then reset the resource reference manually.

System action: None.

Operator response: Inspect the message log of the end-to-end automation domain in order to find the reason the resource reference is broken. Once the root cause of the problem is fixed, reset the resource reference.

EEZT0750W SA event with unknown reason in domain "*domainName*".

Explanation: See message.

System action: None.

Operator response: Analyze the event details and contact IBM support.

EEZT0753I SA automation manager started successfully.

Explanation: See message.

System action: None.

Operator response: None.

EEZT0754W SA automation manager stopped successfully.

Explanation: See message.

System action: None.

Operator response: Verify that the automation manager was stopped intentionally.

EEZT0756I SA detected a new configuration in domain " *domainName* ". New request for resource " *resourceName* " with action *action* and source *source* created.

Explanation: See message.

System action: None.

Operator response: None.

EEZT0758I SA detected a deleted configuration in domain " *domainName* ". Request for resource " *resourceName* " with action *action* and source *source* deleted.

Explanation: See message.

System action: None.

Operator response: None.

EEZT0759W SA detected a status change of resource " *resourceName* " in domain " *domainName* ". **Common observed state:** *CommonObservedState*, **Common desired state:** *CommonDesiredState*, **Common operational state:** *CommonOperationalState*, **Common compound state:** *CommonCompoundState*

Explanation: The state of this automation resource in the specified automation domain has changed.

System action: None.

Operator response: Verify that the new state is intended for this automation resource.

EEZT0770I SA started the automation of resource " *resourceName* " in domain " *domainName* ". **Common observed state:** *CommonObservedState*, **Common desired state:** *CommonDesiredState*, **Common operational state:** *CommonOperationalState*, **Common compound state:** *CommonCompoundState*

Explanation: This resource is now under control of System Automation.

System action: None.

Operator response: Verify that the new state is intended for this automation resource.

EEZT0771I SA completed the automation of resource " *resourceName* " in domain " *domainName* ". **Common observed state:** *CommonObservedState*, **Common desired state:** *CommonDesiredState*, **Common operational state:** *CommonOperationalState*, **Common compound state:** *CommonCompoundState*

Explanation: This resource is no longer under control of System Automation.

System action: None.

Operator response: Verify that the new state is intended for this automation resource.

EEZT0778I SA set the resource " *resourceName* " online as requested in domain " *domainName* ". **Common observed state:** *CommonObservedState*, **Common desired state:** *CommonDesiredState*, **Common operational state:** *CommonOperationalState*, **Common compound state:** *CommonCompoundState*

Explanation: System Automation has set the resource online successfully.

System action: None.

Operator response: None.

EEZT0779I SA set the resource " *resourceName* " offline as requested in domain " *domainName* ". **Common observed state:** *CommonObservedState*, **Common desired state:** *CommonDesiredState*, **Common operational state:** *CommonOperationalState*, **Common compound state:** *CommonCompoundState*

Explanation: System Automation has set the resource offline successfully.

System action: None.

Operator response: Verify that the new state is intended for this automation resource.

EEZT0786W SA detected a warning in domain " *domainName* ". The resource " *resourceName* " has problems caused by a dependency to another resource. **Common observed state:** *CommonObservedState*, **Common desired state:** *CommonDesiredState*, **Common operational state:** *CommonOperationalState*, **Common compound state:** *CommonCompoundState*

Explanation: An automation operation on the specified resource cannot be performed because it

depends on another resource which prevents the operation.

System action: None.

Operator response: Inspect the status of the resources on which the specified resource depends.

EEZT0789E SA detected a problem in domain " *domainName* ". The resource " *resourceName* " cannot be set to the requested state. Common observed state: *CommonObservedState*, Common desired state: *CommonDesiredState*, Common operational state: *CommonOperationalState*, Common compound state: *CommonCompoundState*

Explanation: An automation operation on the specified resource cannot be performed.

System action: None.

Operator response: Inspect the individual status fields of this resource.

EEZT0790E SA detected a problem in domain " *domainName* ". The resource " *resourceName* " needs a manual reset by the operator. Common observed state: *CommonObservedState*, Common desired state: *CommonDesiredState*, Common operational state: *CommonOperationalState*, Common compound state: *CommonCompoundState*

Explanation: An automation operation on the specified resource cannot be performed.

System action: None.

Operator response: Reset the automation resource manually. Invoke the operation "reset from broken" in the operation console.

EEZU0001E The following RuntimeException occurred:
Exception text

Explanation: The processing was interrupted by a RuntimeException and cannot complete correctly.

System action: The current task ends.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZU0002E The following error occurred while writing file *filename*:
Exception text

Explanation: The processing was interrupted by an error and cannot complete correctly.

System action: The current task ends.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZU0003E The following error occurred while reading file *filename*:
Exception text

Explanation: The processing was interrupted by an error and cannot complete correctly.

System action: The current task ends.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZU0004E An error has occurred while accessing the end-to-end automation manager:
Exception text

Explanation: An error has occurred while accessing the end-to-end automation manager component running on the management server. The requested action could not be processed.

Possible reasons are:

- 1) The management server is down or not working correctly.
- 2) There are some inconsistencies regarding the level of the operations console and the end-to-end automation manager.

System action: The requested action is cancelled.

Operator response: Check your connection to the management server and ensure that you have the right permissions. Also check that the management server is up and running. Verify that the levels of the operations console and the end-to-end automation manager are appropriate. If the problem persists, contact your system administrator.

EEZU0005E The credential vault service was not found or could not be loaded:
Exception text

Explanation: The credential vault cannot be accessed because the corresponding service was not found or could not be loaded due to an initialization error.

System action: The current task ends.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZU0006E The page with the ID *Page UUID* could not be found:
Exception text

Explanation: The application tried to load the page with the specified ID to display the log data. However,

the page with this ID could not be found.

System action: The application continues, but the log data cannot be displayed.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZU0007E The credential vault cannot be accessed:
Exception text

Explanation: Possible causes of the problem are:
1) The credential vault is not accessible for technical reasons.
2) The credential vault is not accessible for security reasons.

System action: The current task ends.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZU0008E The credential secret for automation domain Automation domain name is not set:
Exception text

Explanation: A user credential for a certain automation domain was requested but is not set for the user.

System action: The current task ends.

Operator response: If the problem persists after a new log-on, provide the trace file and this message text to IBM support.

EEZU0010E Unable to receive events from the end-to-end automation manager. The following error occurred while trying to read an event:
Exception text

Explanation: An error has occurred while trying to access the event path to the management server. The operations console is not able to receive any events and is therefore not able to update the status information for resources if the status changes.

Possible reasons are:

- 1) The management server is down.
- 2) The JMS service of the management server is not working properly.

System action: Processing continues, but no events can be received.

Operator response: Check your connection to the management server and ensure that you have the right permissions. Also check that the management server is up and running. If the problem persists, contact your system administrator.

EEZU0011E Unable to set up the event path between the operations console and the management server:
Exception text

Explanation: The connection to the right JMS service on the management server could not be established. This connection is used to receive events about status changes from the end-to-end automation manager. Possible reasons are:

- 1) The management server is down.
- 2) The JMS service of the management server is not working properly.

System action: Processing ends.

Operator response: Check your connection to the management server and ensure that you have the right permissions. Also check that the management server is up and running. If the problem persists, contact your system administrator.

EEZU0012E An error occurred trying to look up the JMS service on the management server to establish the event path:
Exception text

Explanation: An error has occurred while trying to access the management server. Possible reasons are:

- 1) The management server is down.
- 2) The JMS service of the management server is not working properly.
- 3) The JMS topic for end-to-end automation management has not been configured correctly on the management server.

System action: Processing ends.

Operator response: Check your connection to the management server and ensure that you have the right permissions. Also check that the management server is up and running and that JMS is configured correctly. If the problem persists, contact your system administrator.

EEZU0013E An error has occurred while trying to establish the connection to the end-to-end automation manager:
Exception text

Explanation: An error has occurred while connecting to the end-to-end automation manager component running on the management server.

Possible reasons are:

- 1) The management server is down or not working correctly.
- 2) There are some inconsistencies regarding the level of the operations console and the end-to-end automation manager.

System action: Processing ends.

Operator response: Check your connection to the management server and ensure that you have the right permissions. Also check that the management server is

up and running. Verify that the levels of the operations console and the end-to-end automation manager are appropriate. If the problem persists, contact your system administrator.

EEZU0014E An error has occurred while trying to establish the connection to the end-to-end automation manager:
Exception text

Explanation: An error has occurred while connecting to the end-to-end automation manager component running on the management server.

Possible reasons are:

- 1) The management server is down or not working correctly.
- 2) The management server is not accessible.

System action: Processing ends.

Operator response: Verify your entries for server name and port are correct. Ensure that you have the right permissions. Also check that the management server is up and running.

Refer to the 'Related errors' section for more details about the problem.

If the problem persists, contact your system administrator.

EEZU0015E The log data cannot be displayed because the service to launch a new page was not found or could not be loaded.

Explanation: The log data is normally displayed on a new page within the Integrated Solutions Console, but the service to launch a new page was not found or could not be loaded due to an initialization error.

System action: The application continues, but the log data cannot be displayed.

Operator response: If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZU0017E There is no log data available for automation domain *Automation domain*.

Explanation: No log file exists for the automation domain. The log file is normally located on the node where the automation domain's automation adapter is running, or if it is the end-to-end automation domain, where the end-to-end automation engine is running.

System action: The application continues without displaying log data.

Operator response: Ensure that logging is set up correctly for this automation domain; for example, check the `eezjlog.properties` file. If the problem persists, contact your system administrator.

EEZU0018E Creating EIF event receiver failed, error message is: *Exception text*.

Explanation: The operations console accesses first-level automation domains directly (direct access mode). To be able to receive events from first-level automation domains an Event Integration Facility (EIF) event receiver must be created. Creating the event receiver failed.

System action: The operations console will not receive events.

Operator response: Examine the error message to find the cause of failure. If you cannot resolve the problem, provide the trace file and this message text to IBM support.

EEZU0019E The operations console was notified of new domain *new domain* that has the same name as the known domain *known domain*.

Explanation: The operations console accesses first-level automation domains directly (direct access mode). It was notified about a new domain that has the same name as a domain that is already known by the operations console. However, the connection information of the of the form '`domainname@ip-address:port`' suggest that the new domain automates a different cluster than the known domain. Every domain operated from an operations console must have a unique name.

System action: The domain is not allowed to join and therefore, will not show up in the topology view.

Operator response: Try to determine from the information of new domain where the domain is located. If the new domain automates a different cluster than the known domain, have the name of the new domain changed, and its automation adapter restarted to notify the operations console.

EEZU0020E The operations console was notified of domain *domain* from adapter *adapter* with version *adapter version* that is lower than the required minimum version *minimum version*.

Explanation: The operations console accesses first-level automation domains directly (direct access mode). It was notified about a domain from an adapter with a version that is too low for reliable operation.

System action: The domain is not allowed to join and therefore, will not show up in the topology view.

Operator response: Try to locate the adapter that tried to join the domain and have it upgraded to a version that is equal or higher than the required minimum version. Then have the automation adapter restarted to notify the operations console.

EEZU0021E The operations console contacted a domain *domain* with adapter *adapter* at version *adapter version* that is lower than the required minimum version *minimum version*.

Explanation: The operations console accesses first-level automation domains directly (direct access mode). It contacted a domain from an adapter with a version that is too low for reliable operation.

System action: The operations console must not communicate with the domain which has a too low version and therefore, the domain will remain disabled in the topology view.

Operator response: Try to locate the adapter of the domain and have it upgraded to a version that is equal or higher than the required minimum version. Then have the automation adapter restarted to notify the operations console.

EEZU0500W The selected automation domain *domain name* no longer exists.

Explanation: You selected an automation domain that no longer exists. Possible reasons are that the automation domain has been deleted in the meantime.

System action: The current task continues.

Operator response: If the automation domain is still displayed, use menu item 'Refresh all' to obtain the currently available domains.

EEZU0501W The selected resource *resource name* no longer exists.

Explanation: You selected a resource that no longer exists. Possible reasons are that the resource has been deleted in the meantime or the automation policy has been changed or deactivated.

System action: The current task continues.

Operator response: If the resource is still displayed, use menu item 'Refresh all' to obtain the currently available resources.

EEZU0502W The selected node *node name* no longer exists.

Explanation: You selected a node that no longer exists. Possible reasons are that the node has been deleted in the meantime.

System action: The current task continues.

Operator response: If the node is still displayed, use menu item 'Refresh all' to obtain the currently available nodes.

EEZU0503W The request has been submitted but has not been processed yet.

Explanation: A request has been submitted but was not processed by the corresponding automation manager. Reasons for this can be a slow network or an automation manager that is not responding.

System action: The application continues.

Operator response: If the request is not processed soon, send the request again. If the problem persists, check the connections to the automation manager and inspect the log files of the automation manager for problems.

EEZU0504W The order to cancel the operator request has been submitted, but the request is still not cancelled yet.

Explanation: A cancel request has been submitted but was not processed by the corresponding automation manager. Reasons for this can be a slow network or an automation manager that is not responding.

System action: The application continues.

Operator response: If the request is not processed soon, cancel the request again. If the problem persists, check the connections to the automation manager and inspect the log files of the automation manager for problems.

EEZU0505W The order to change the automation policy has been submitted, but the policy change has not been completely processed yet.

Explanation: The order to change the automation policy has been submitted to the corresponding automation manager, but the processing of this change has not finished yet. Reasons for this can be a slow network or an automation manager that is not responding.

System action: The application continues. When the processing of the policy change has been completed the screen will automatically refresh to reflect the change.

Operator response: If the problem persists, check the connections to the automation manager and inspect the log files of the automation manager for problems.

EEZU0506W Domain *Domain name* became unavailable.

Explanation: The operations console accesses first-level automation domains directly (direct access mode). A domain that had been contacted successfully before, became unavailable when the operations console tried to perform a request on a first-level automation domain. The automation adapter or the node of the domain may have shut down without being able to notify the operations console.

System action: The request and any further request will not be performed on the domain until it becomes available.

Smart refresh messages

This chapter lists the messages that are generated by the Smart refresh function.

AZPCTP001E The refresh interval {0} is not a valid number of seconds.

Explanation: The refresh interval represents the number of seconds that will pass before the page is refreshed. The value must be a number between 5 and 999.

Operator response: Enter a valid refresh interval and run the command again.

AZPCTP009E The refresh mode has been changed to the timed refresh mode because this browser type does not support smart refresh mode.

Explanation: This browser type does not support smart refresh mode, so the refresh mode was automatically changed to timed refresh.

Operator response: Change the browser to one of the browsers supported by the Tivoli Management Portal.

AZPCTP012E Your browser does not support XMLHTTP Objects.

Explanation: Without support for XMLHTTP objects, the smart refresh mode cannot be used.

Operator response: Switch to a supported browser or do not use the smart refresh mode. Use the timed refresh mode instead.

AZPCTP013E There was an error connecting to the server while checking the smart refresh.

Explanation: When checking to see if a smart refresh was needed, the HTTP connection to the server could not be established. This might be a temporary connectivity error, or this browser might not support XMLHTTP objects.

Operator response: If this error is always seen when using this browser, then this browser probably does not support XMLHTTP objects. These objects are required for smart refresh mode. To avoid this error, switch to a supported browser or do not use the smart refresh mode. If this error is not always seen when using this browser, this might be a temporary connectivity problem with the server.

AZPCTP014E The action {0} is not a valid action for this servlet.

Explanation: The servlet has received a POST request with an invalid action. The only valid action is

Operator response: Select menu item 'Refresh all' to obtain the new state of the domains.

GET_NEWCONTENTFLAG.

Operator response: This might be an erroneous request or an attack on this servlet

Programmer response: Check that the action being sent is GET_NEWCONTENTFLAG.

AZPCTP015E No action was received in the POST request for this servlet.

Explanation: The servlet has received a POST request with no specified action. The only valid action is GET_NEWCONTENTFLAG.

Operator response: This might be an erroneous request or an attack on this servlet

Programmer response: Check that the action being sent is GET_NEWCONTENTFLAG.

AZPCTP016E No sessionId was received in the POST request for this servlet.

Explanation: The servlet has received a POST request with no specified sessionId.

Operator response: Contact IBM software support. There is a communication problem between the browser and the Tivoli Management Portal that cannot be resolved by the user.

AZPCTP017E The refresh interval {0} is not a valid number of seconds.

Explanation: The refresh interval represents the number of seconds that will pass before the page is refreshed. Only the numerical values between 5 and 999 are valid.

Operator response: Enter a valid refresh interval and repeat the command.

AZPUTL001E Unable to retrieve a connection to the database.

Explanation: Cannot open a connection to the cloudscape database. The necessary JAR files containing the driver (db2j.jar) and license (license.jar) for the cloudscape database might have been deleted or moved out of <WAS_HOME>/AppServer/lib/app.

Operator response: Check to see if the necessary JAR files are present.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie New York 12601-5400
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS"

WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

- IBM, AIX, DB2, NetView, Tivoli, Tivoli Enterprise, Tivoli Enterprise Console, WebSphere, and z/OS are trademarks of International Business Machines Corporation in the United States, other countries, or both.
- Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Other company, product, and service names may be trademarks or service marks of others.

Glossary

This glossary defines technical terms and abbreviations used in end-to-end automation management documentation.

If you do not find the term you are looking for, refer to the *IBM Dictionary of Computing Terms*, located at <http://www.ibm.com/networking/nsg/>

In the glossary, glossary items are shown in *italics* where they are used in definitions for other items.

A

administrator. A person who configures and administers applications and components and creates and maintains policies.

automated resource. A resource that is managed by an automation manager.

automation engine. The subcomponent of end-to-end automation management that makes the automation decisions based on the policy.

C

choice group. An end-to-end automation resource group whose members are alternatives. Only one of the members can be active at a time. If the desired state of the choice group is Online, the end-to-end automation manager tries to keep the active resource online but will only start the resource in place if it fails. An operator can start a different member of a choice group from the operations console.

configuration dialog. The graphical user interface that is used to configure the end-to-end automation management component.

D

desired state. The automation goal of a resource that is pursued by the automation manager.

direct access mode. The operations console mode in which only resources that are automated by the Base component of IBM Tivoli System Automation for Multiplatforms can be managed and monitored from the console.

domain health indicator. A resource whose state is used to indicate whether or not a domain is healthy. If the observed state of such a resource differs from its

desired state, an error or a warning appears on the operations console for the domain by which it is hosted.

E

end-to-end automation management. (1) The component of IBM Tivoli System Automation for Multiplatforms that allows the automation of resources that are hosted by different first-level automation domains. (2) The automation tasks that are performed by the subcomponents of the end-to-end automation management component.

end-to-end automation mode. The operations console mode in which end-to-end automation management is installed and active. In this mode, resources that are hosted by the end-to-end automation domain and by first-level automation domains can be monitored and managed from the operations console.

Extensible Markup Language. A meta-language for defining specialized markup languages. An XML (Extensible Markup Language) document has a hierarchical structure of tags and attributes according to a specific syntax. XML documents are plain text documents and thus platform and language independent.

F

first-level automation mode. The operations console mode in which only resources that are hosted by a first-level automation domain can be monitored and managed from the console. To run the operations console in this mode, the end-to-end automation management component must be installed but end-to-end automation management must not be active.

first-level automation resource. A resource that is automated by a first-level automation product, such as the Base component of IBM Tivoli System Automation for Multiplatforms or IBM Tivoli System Automation for z/OS.

I

installation phase. The tasks that are performed after the **Install** button is clicked in the installation wizard (contrast with *pre-installation phase*).

ISC. See *Integrated Solutions Console*.

Integrated Solutions Console. An application running in WebSphere Application Server in which multiple

IBM applications can be installed which can be accessed from a single Web interface. The operations console is installed in Integrated Solutions Console.

L

LDAP. See *Lightweight Directory Access Protocol*.

Lightweight Directory Access Protocol. A directory service protocol that runs over TCP/IP.

Lightweight Third Party Authentication. An authentication framework that allows single sign-on across a set of servers that fall within a domain.

LTPA. See *Lightweight Third Party Authentication*

M

member. A constituent of a group.

monitor resource. A first-level automation resource whose current state can be monitored from the operations console but whose desired state cannot be changed through start and stop requests that are submitted from the console.

N

nested groups. Groups that are members of other groups.

node. A physical system on which first-level automation resources are located.

O

operations console. The graphical user interface from which automated resources can be monitored and managed.

operator. A person who monitors and manages automated resources from the operations console.

P

policy. The list of resources that are managed by end-to-end automation management, their relationships, and the set of rules that determines the behavior of the end-to-end automation manager.

policy checking tool. The tool that ensures that each policy is syntactically correct and complies with the end-to-end automation-specific rules. In end-to-end automation management, each policy is checked automatically before it is activated.

preferred member. The member of a choice group which is kept online by the automation manager if the

desired state of the choice group is online. Operators can change the preferred member from the operations console.

pre-installation phase. The tasks that are performed before the **Install** button is clicked in the installation wizard (contrast with *installation phase*).

R

referenced resource. A first-level automation resource for which a resource reference has been defined in the active end-to-end automation policy.

relationship. A start or stop dependency between two resources that results in a specific automation behavior. Relationships are defined in end-to-end automation policies.

request. A method for changing the automation goal for a resource. In end-to-end automation management, Online and Offline request are used to start or stop resources.

resource. (1) In IBM Tivoli System Automation, any component, job, device, or target system capable of being monitored or automated through IBM Tivoli System Automation. (2) Any element that can be automated, monitored and managed as one unit, for example, individual resources, resource groups, and choice groups.

resource group. A collection of resources that have the same desired state and are monitored and managed as one unit. In end-to-end automation management, resource groups can comprise resource references, resource groups, and choice groups.

resource reference. A virtual resource that is managed by the end-to-end automation manager. Resource references are defined in a policy and reference to actual resources that are hosted by first-level automation domains (see *referenced resource*).

S

Secure Sockets Layer. A security protocol that ensures the data that is transferred between a client and a server remains private.

SSL. See *Secure Sockets Layer*

T

top-level resource. In end-to-end automation, a resource that is displayed at the topmost level in the resource tree for a domain. Typically, these are resources that are either not members of a group, or groups that are not nested within other groups. By default, such resources are used as *domain health indicators*.

U

user credentials. The information used for user authentication, for example, the signon information for a user (user ID and password).

V

vote. An internal request to change the *desired state* of a resource. Votes are generated by an automation manager against resources that have to be started or stopped because an Online or Offline request was issued against an automation resource that has start or stop dependencies.

X

XML. See *Extensible markup language*.

XML policy. An XML file that contains the policy definitions for an end-to-end automation policy.

Index

A

- access roles
 - introduction 143
- Additional info page
 - state-related information 287
- AIX systems
 - multiple CPUs 300
- automation adapters
 - starting 189
- automation engine
 - command-line options 183
 - eezdmn command 183
 - JMS authentication 158
 - log and trace files 293
 - stopping 183
 - user ID 148
 - user ID in sas.client.props 156
 - XML log file, viewing 294
- automation J2EE framework
 - environment variables 301
 - log and trace files 294
 - starting 190
 - stopping 190
- automation management server
 - JMS authentication 159
- automation manager
 - maximum number of
 - connections 107
- AutomationDomain
 - XML element 267
- AutomationDomainName
 - XML element 267
- AutomationPolicy
 - XML element 268

C

- choice group
 - request Offline 252
- choice groups
 - changing the preferred member 257
 - definition 27
 - Online request against a member 251
 - overview 256
 - starting the preferred member 256
- ChoiceGroup
 - XML element 268
- command-line options
 - automation engine 183
- communication flows
 - first-level automation mode 24
 - policy activation 19
 - request submission 23
- communication state 230
- compound state 225
 - icons 226
 - resources 232
 - values 225
- configuration dialog 193

- configuration properties
 - configuration dialog 193
 - properties files 263
- contact information
 - displaying 238
- conversion-only mode 15, 24, 185, 187
- CORBA.NO_RESPONSE
 - errors 302
- credential vault
 - Integrated Solutions Console 247

D

- DB2
 - automation manager database
 - cataloging 64
 - creating 61
 - client installation 62
 - installation parameters 72
 - local setup 51
 - operations console database
 - cataloging 63
 - creating 61
 - remote setup 51
 - server installation 59
 - setup options 51
 - software prerequisites
 - for remote setup 54
 - local setup 54
 - TCP node
 - cataloging 63
 - user registry 51
 - WebSphere Application Server
 - requests 99
- DB2 access
 - user ID of the automation management server 157
- default directories 261
- Description
 - XML element 270
- desired state
 - resources 232
- DesiredState
 - XML element 270
- direct access mode
 - overview 16
- directories
 - default paths 261
- disk space requirements
 - AIX 56
 - Linux 56
 - Windows 55
- domain health indicators
 - defining 244
- domain identification file 263
- domain state 229
- domains
 - communication state 230
 - displaying, troubleshooting 296
 - domain state 229
 - hiding 243

- domains (*continued*)
 - operational state 227

E

- Eclipse Help System server
 - starting 190
 - stopping 190
- EEZBus
 - resolving problems 306
- eezdmn command
 - options
 - ? 188
 - co 187
 - monitor 186
 - reconfig 187
 - shutdown 185
 - start 184
 - xd 188
 - quick reference 184
 - using 183
- EEZEAR
 - starting 190
 - stopping 190
- electronic distribution 48
- end-to-end automation domain
 - name 80
- end-to-end automation manager
 - configuring 193
- end-to-end automation mode
 - overview 15
- end-to-end automation policies
 - XML reference 267
- environment variables
 - automation J2EE framework 301
- Errors and warnings
 - View button 238
- event path error
 - resolving 306
- external shutdown
 - SA z/OS resources 167
- external startup
 - SA z/OS resources 167

F

- first-level automation mode
 - communication flow 24
 - overview 15
- fix packs
 - installing 113
- ForcedDownBy relationships 33
 - defining 178

G

- goal-driven automation
 - overview 27

H

- home page
 - IBM Tivoli System Automation 113
- HTTP session timeout
 - modifying 106

I

- IBM Tivoli Enterprise Console
 - installation parameters 79
- IBM.Equivalency
 - resource class 251
- IllegalMonitorStateException
 - troubleshooting 300
- info link 238
- information area
 - overview 222
- information pages
 - setting up 181
- initial resource events 38
- installation
 - planning 47
 - post-installation tasks 103
 - prerequisites 81
 - product CD 47
 - troubleshooting 119
 - verifying 99
- installation directory 71
- installation log files
 - locating 119
- installation wizard
 - launching 83
- installing
 - service fix packs 113
- Integrated Solutions Console
 - changing passwords 154
 - credential vault 247
 - deleting user groups 155
 - deleting user IDs 155
 - event path error 306
 - logging on 209
 - server port number 296
 - starting 190
 - stopping 190

J

- J2EE framework
 - starting 190
 - stopping 190
- JMS authentication 157
 - automation engine 158
 - automation management server 159
 - operations console 158

L

- LDAP
 - directory tree structure 67
 - installation parameters 74
 - pre-installation tasks 68
 - required user groups 68
 - sample configuration 69
 - setting up 67
 - software prerequisites 54

- LDAP (*continued*)
 - user registry 51
- log files
 - automation engine 293
 - automation J2EE framework 294
 - locations 293
 - viewing 237
- log viewer 294
- logging on
 - Integrated Solutions Console 209
 - operations console 209
- LTPA settings
 - LTPA password 105
 - LTPA timeout 105

M

- main menu 223
- manager automation flag
 - SA z/OS resources 167
- Members
 - XML element 272
- messages
 - generated by the Smart refresh function 373
 - interpreting 323
- middleware software CDs
 - contents 59
- monitor resources 39

N

- Name
 - XML element 273
- name filters
 - administering 241
 - applying 241
 - defining 239
- Node
 - XML element 271
- nodes
 - excluding from automation 255
 - including in automation 255
 - observed state 231
- NOSTART option
 - SA z/OS resources 167

O

- observed state
 - nodes 231
 - resources 232
- operating systems 52
- operational state 225
 - domains 227
 - resources 232
- operations console
 - accessing 209, 210
 - direct access mode 16
 - end-to-end automation mode 15
 - first-level automation mode 15
 - icons 309
 - information area
 - overview 222
 - installation parameters 76
 - installation problems 119

- operations console (*continued*)

- JMS authentication 158
 - layout 212
 - main menu 223
 - refreshing 245
 - smart refresh 223
 - time zone settings 304
 - topology tree 214
 - topology tree icons 216
 - trace files 295
 - using views 238
 - operator instructions
 - displaying 238
 - operator request icons 322
 - ORB request timeout 302
 - ORB service 302
 - Owner
 - XML element 273
- ## P
- packaging
 - product CD 47
 - passive application groups
 - SA z/OS resources 167
 - passwords
 - changing, on Integrated Solutions Console 154
 - policies
 - checking
 - from a command line 179
 - from the operations console 249
 - deactivating 250
 - defining choice groups 176
 - defining groups 174
 - defining relationships 177
 - defining resource groups 174
 - defining resources 171
 - ForcedDownBy relationships 178
 - modifying 251
 - rules 281
 - SA z/OS resources 167
 - sample policy 127, 282
 - schema 168
 - StartAfter relationships 177
 - StopAfter relationships 178
 - UTF-8 format 168
 - worksheet 265
 - XML declaration 169
 - XML elements
 - AutomationDomain 173, 267
 - AutomationDomainName 170, 267
 - AutomationPolicy 169, 268
 - ChoiceGroup 268
 - Class 173, 269
 - Description 270
 - DesiredState 270
 - InfoLink 271
 - Members 272
 - Name 173, 273
 - Node 173, 271
 - Owner 273
 - PolicyAuthor 171, 273
 - PolicyDescription 171, 274
 - PolicyInformation 170, 274
 - PolicyName 170, 275

- policies (*continued*)
 - XML elements (*continued*)
 - PolicyToken 171, 275
 - ReferencedResource 172, 276
 - Relationship 276
 - ResourceGroup 277
 - ResourceReference 172, 277
 - Source 279
 - Target 279
 - Type 280
 - XML reference 265, 267
 - XML template 168
- policy checking tool
 - starting 179
- policy pool directory 179
- PolicyAuthor
 - XML element 273
- PolicyDescription
 - XML element 274
- PolicyInformation
 - XML element 274
- PolicyName
 - XML element 275
- PolicyToken
 - XML element 275
- post-installation tasks
 - for administrators
 - access roles 143
 - assigning access permissions 146
 - assigning access roles 149
 - automation engine user ID 148
 - creating user groups 145
 - overview 143
 - modifying the LTPA settings 105
 - overview 103
 - setting up SSL 103
- product CD
 - contents 80
 - directories 80
 - installation wizard files 47
- properties files
 - file list 263
 - sas.client.properties 156
 - vpd.properties 56, 119
- PropFilePasswordEncoder utility 156

R

- ReferencedResource
 - XML element 276
- Refresh all 245
- Relationship
 - XML element 276
- relationships
 - displaying 237
 - ForcedDownBy 33, 178
 - StartAfter 30, 177
 - StopAfter 32, 178
- release notes 113
- request lists
 - displaying 253
 - viewing 253
- requests
 - canceling 253
 - displaying information about requests 252
 - Online 251

- requests (*continued*)
 - overview 251
 - stop 252
- resource groups
 - definition 27
- resource references
 - definition 27
 - SA for Multiplatforms resources
 - restrictions 167
 - SA z/OS resources
 - restrictions 167
- resource table
 - limiting the scope 238
- ResourceGroup
 - XML element 277
- ResourceReference
 - XML element 277
- resources
 - compound state 232
 - desired state 232, 235
 - icons 310
 - monitoring 225
 - observed state 232, 234
 - operational state 232
 - resetting, from unrecoverable errors 254
 - stopping 252
- resources section
 - overview 217
- return codes
 - eezdmn -co 187
 - eezdmn -reconfig 187
 - eezdmn -shutdown 185
 - eezdmn -start 185
 - eezdmn -xd 188

S

- SA for Multiplatforms
 - restrictions 167
- SA z/OS resources
 - restrictions
 - external shutdown 167
 - external startup 167
 - manager automation flag 167
 - NOSTART option 167
 - passive application groups 167
- sa_status_automation 290
- sa_status_compound 288
- sa_status_control 290
- sa_status_desired 289
- sa_status_observed 289
- sample policy 127, 282
- sas.client.props
 - automation engine user ID 156
 - encryption 156
- security concepts
 - overview 57
 - SSL 57
- server port number
 - Integrated Solutions Console 296
- service
 - installing 113
- smart refresh 223, 245
- messages 373
- Source
 - XML element 279

- SSL
 - security concepts 57
 - setting up 103
- start requests 251
- StartAfter relationships 30
 - defining 177
- starting
 - automation adapters 189
 - automation engine 183
 - Eclipse Help System server 190
 - EEZEAR 190
 - Integrated Solutions Console 190
 - J2EE framework 190
 - resources 251
 - WebSphere Application Server on AIX and Linux 190
 - WebSphere Application Server on Windows 189
- state change event 21
- stop requests 252
- StopAfter relationships 32
 - defining 178
- stopping
 - automation engine 183
 - Eclipse Help System server 190
 - EEZEAR 190
 - Integrated Solutions Console 190
 - J2EE framework 190
 - resources 252
 - WebSphere Application Server on AIX and Linux 190
 - WebSphere Application Server on Windows 189
- subscription 20
- support site
 - IBM Tivoli System Automation 113
- supported languages 81

T

- Target
 - XML element 279
- TCP/IP connectivity
 - hardware requirements 55
- testing
 - connection between WebSphere Application Server and DB2 99
- time zone settings 304
- timeouts
 - HTTP session timeout 106
 - LTPA timeout 105
 - resolving problems 301
- Tivoli Common Directory 71, 293
- top-level resources 217, 244
- topology tree
 - hiding domains 216, 243
 - icons 216, 309
 - Located here column 217
 - overview 215
 - Status column 217
- trace files
 - automation engine 293
 - automation J2EE framework 294
 - locations 293
 - operations console 295
- trademarks 376

- troubleshooting
 - DB2
 - connection problem 305
 - installation problems 119
 - WebSphere Application Server 300
 - connection problem 305
- Type
 - XML element 280

- WebSphere Application Server (*continued*)
 - stopping (*continued*)
 - on Windows 189
 - troubleshooting 300
 - Upgrade CD 48
- worksheet
 - for policy definition 265

U

- uninstallation 109
- unrecoverable errors
 - resetting resources 254
 - resolving problems 304
- user credentials
 - managing 247
- user groups
 - deleting, on Integrated Solutions Console 155
- user management
 - access roles
 - introduction 143
- user registry
 - DB2 51
 - software prerequisites 54
 - LDAP 51
 - software prerequisites 54
 - options 51
- users
 - assigning users to groups 154
 - authorizing 153
 - changing passwords 154
 - creating 153
 - deleting, on Integrated Solutions Console 155

X

- XML policy files
 - schema 168
 - template 168
- UTF-8 format 168

W

- Web browsers
 - configuring 209
 - JavaScript 209
 - multiple browser windows 209, 296
 - requirements 54
 - security level 209
 - security settings 209
- WebSphere Application Server
 - connection problem
 - troubleshooting 305
 - connection to DB2
 - test 99
 - installation parameters 74
 - installing
 - interim fixes 64, 65
 - post-installation tasks 66
 - Refresh Pack 2 64, 65
 - Version 6.0.0.0 64, 65
 - interim fixes 113
 - ISC_Portal 50
 - server1 50
 - start-up problems 305
 - starting
 - on AIX and Linux 190
 - on Windows 189
 - stopping
 - on AIX and Linux 190

Readers' Comments — We'd Like to Hear from You

IBM Tivoli System Automation for Multiplatforms
End-to-End Automation Management
User's Guide and Reference
Version 2.1.1

Publication No. SC33-8211-02

Overall, how satisfied are you with the information in this book?

| | Very Satisfied | Satisfied | Neutral | Dissatisfied | Very Dissatisfied |
|----------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| Overall satisfaction | <input type="checkbox"/> |

How satisfied are you that the information in this book is:

| | Very Satisfied | Satisfied | Neutral | Dissatisfied | Very Dissatisfied |
|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| Accurate | <input type="checkbox"/> |
| Complete | <input type="checkbox"/> |
| Easy to find | <input type="checkbox"/> |
| Easy to understand | <input type="checkbox"/> |
| Well organized | <input type="checkbox"/> |
| Applicable to your tasks | <input type="checkbox"/> |

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape

PLACE
POSTAGE
STAMP
HERE

IBM Deutschland Entwicklung GmbH
Department 3248
Schoenaicher Strasse 220
D-71032 Boeblingen
Germany

Fold and Tape

Please do not staple

Fold and Tape