Tivoli® System Automation for Multiplatforms

IBM

Version 2.2

End-to-End Automation Management Component Administrator's and User's Guide

SC33-8275-00

Tivoli® System Automation for Multiplatforms

IBM

Version 2.2

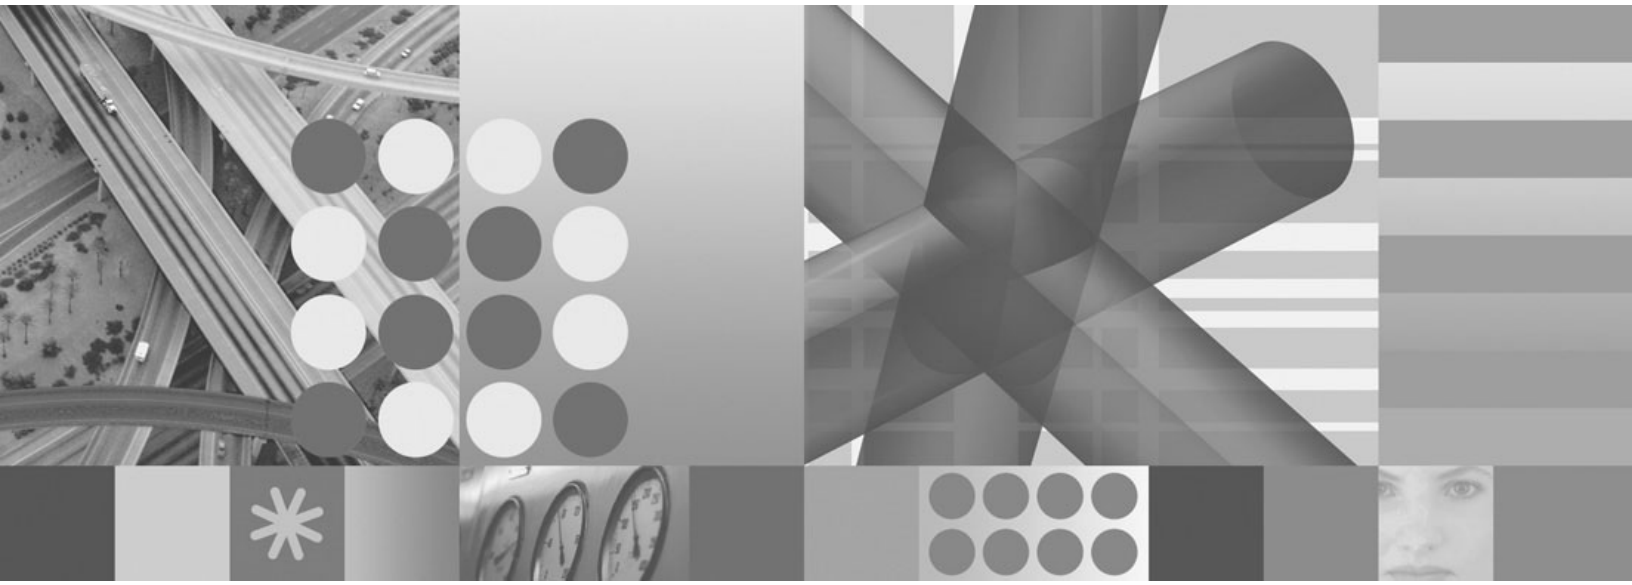

End-to-End Automation Management Component Administrator's and User's Guide

SC33-8275-00

Note!
Before using this information and the product it supports, read the information in Appendix C, "Notices," on page 223.

First Edition (October 2006)

This edition of the *End-to-End Automation Management Administrator's and User's Guide* applies to Version 2, Release 2 of IBM Tivoli System Automation for Multiplatforms, program number 5724–M00, and to all subsequent releases and modifications of this product until otherwise indicated in new editions.

IBM welcomes your comments. A form for readers' comments may be provided at the back of this publication, or you may address your comments to the following address:

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

FAX (Germany): 07031+16-3456
FAX (Other Countries): (+49)+7031-16-3456

Internet e-mail: eservdoc@de.ibm.com

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:
- Title and order number of this book
- Page number or topic related to your comment

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.

Contents

Figures

Tables

About this guide

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

Who should read this guide

This guide is for administrators who 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, "First steps," on page 45 describes in how you can use sample environment that is configured during the installation to learn about end-to-end automation management.
- Part 3, "Administering the end-to-end automation management component," on page 61 describes how to create policies, manage users, and start and stop the components of end-to-end automation management.
- Part 4, "Monitoring and managing automated resources," on page 121 describes how to exploit the functionality of end-to-end automation management.
- Part 5, "Working with the HACMP and MSCS adapters," on page 181 describes how to work with the adapters and with the objects of HACMP and MSCS domains.
- In the Appendixes you find reference information you may need for using and operating the end-to-end automation management component.

Where to find more information

In addition to this manual, the IBM Tivoli System Automation for Multiplatforms library contains the following books:

- *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*, SC33-8273
- *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*, SC33-8276
- *IBM Tivoli System Automation for Multiplatforms Base Component Administrator's and User's Guide*, SC33-8272
- *IBM Tivoli System Automation for Multiplatforms Base Component Reference*, SC33-8274

You can download the documentation at

`http://publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms2.2.html`

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
- 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.
domain health indicators	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. 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. 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.
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	A first-level automation resource that has the following characteristics: • its current state can be monitored from the operations console • its desired state cannot be changed through start and stop requests
resource	Any application, process, or service that is monitored and managed by a first-level or end-to-end automation manager. 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.
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.

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

Term	Description
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 *domain health indicators*.

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

WebSphere Application Server publications:

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/`

IBM DB2 publications:

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

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

The link to the PDF manuals is available in the **Other resources** section on the Web page.

What's new in release 2.2

For release 2.2, the IBM Tivoli System Automation for Multiplatforms library was restructured:

- This *End-to-End Automation Management Component Administrator's and User's Guide* provides the information you need for administering the end-to-end automation management component and for monitoring and managing automated resources.
- The installation and configuration tasks, which were formerly described in the *End-to-End Automation Management Component User's Guide and Reference*, are now described in the *Installation and Configuration Guide*.
- The reference information, which was formerly provided in the *End-to-End Automation Management Component User's Guide and Reference*, is now available in the *End-to-End Automation Management Component Reference*

In release 2.2, the following new features are introduced for the end-to-end automation management component of IBM Tivoli System Automation for Multiplatforms:

End-to-end automation manager command shell

The command shell introduces a command line interface that can be used to monitor and control resources, groups, and relationships which are automated by the end-to-end automation domain. This new interface can be used in addition to the existing graphical user interface (operations console) and does not interfere with it. This allows you, for example, to monitor and control resources from a computer that has no Web browser installed (which is a prerequisite for using the operations console), or to write shell scripts or Windows batch files to drive end-to-end-automation actions, like activating an automation policy under the control of a scheduler product.

Automation adapter for High Availability Cluster Multi-Processing (HACMP) clusters

Using this adapter, first-level automation domain clusters that are managed by HACMP can be integrated into the end-to-end automation environment of IBM Tivoli System Automation for Multiplatforms, and resources that are made highly available by HACMP can be incorporated into end-to-end automation policies. The adapter is delivered as a separately installable entity together with the end-to-end automation management component.

Automation adapter for Microsoft Server Clustering (MSCS) clusters

Using this adapter, first-level automation domain clusters that are managed by MSCS can be integrated into the end-to-end automation environment of IBM Tivoli System Automation for Multiplatforms, and resources that are made highly available by MSCS can be incorporated into end-to-end automation policies. The adapter is delivered as a separately installable entity together with the end-to-end automation management component.

Part 1. Introducing end-to-end automation management

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.

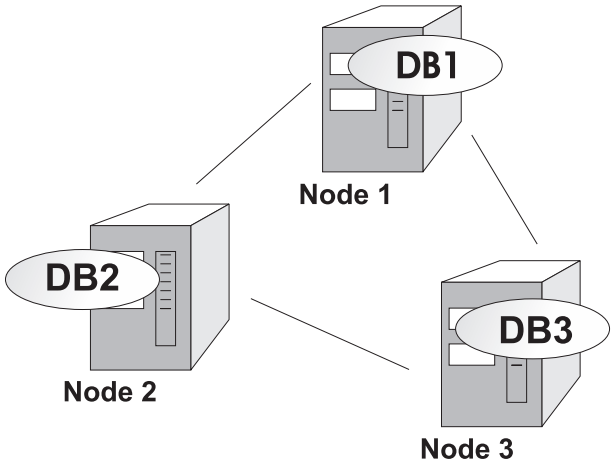

When you use SA for Multiplatforms, you no longer need to specify event correlation rules in sophisticated scripts. Such scripts would describe the desired behavior in complex lists such as

```
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.

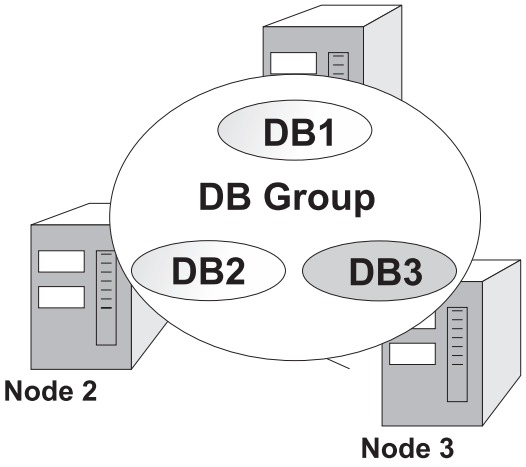

You can also define relationships between resources in the policy. Relationships can define:

- 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, and by High Availability Cluster Multi-Processing (HACMP), and Microsoft Server Clustering (MSCS).

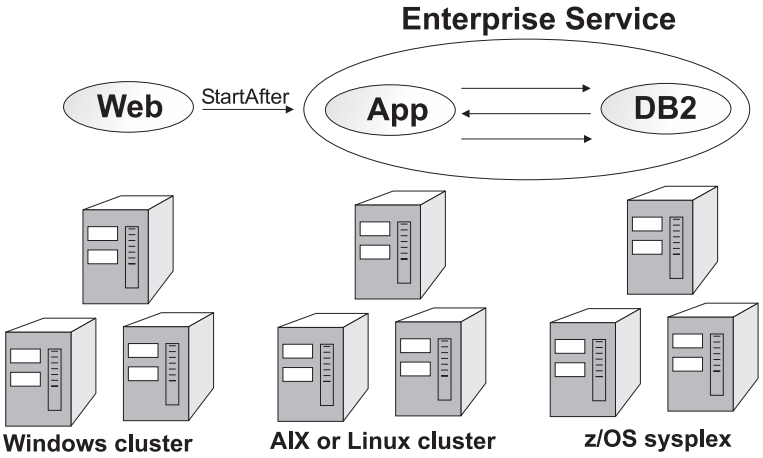

In this manual, an automation solution on a homogeneous cluster is called a first-level automation domain. End-to-end automation management does not replace these first-level automation domains but rather builds upon and integrates them.

In the example shown in the figure above the resource "Web", which is defined on a Windows cluster, has a startAfter relationship to the group "Enterprise Service", which consists of resources that are running on an AIX or Linux 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.
- The 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 who is responsible for ensuring 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.

To accomplish these tasks, operators can use either the operations console of SA for Multiplatforms, which provides a user interface that is designed to support an operator in performing the tasks, or the end-to-end automation manager command shell.

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.

Role of an application owner

In IBM Tivoli System Automation, an application owner is responsible for an application that is automated and, therefore, controlled as a resource at least by a first-level automation product and may even be referenced by a resource reference that is controlled by end-to-end automation.

In either case, application owners can no longer use the standard mechanisms to start and stop these applications. Instead, they must use the proper methods of the first-level automation manager to start and stop such applications (such as the command-line interface of the base component of SA for Multiplatforms). When the application resource is integrated into end-to-end automation, the application owner must use either the end-to-end automation operations console or the end-to-end command shell in order to issue requests to start and stop the application.

A feasible way of doing this is to integrate the automation manager commands (command-line interface commands of the SA for Multiplatforms base component or end-to-end automation manager command shell commands in line mode, respectively) into the startup and shutdown scripts of the application for which the application owner is responsible. This allows application owners to use application-typical scripts for starting and stopping and prevents them from having to remember SA for Multiplatforms specific commands.

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

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

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

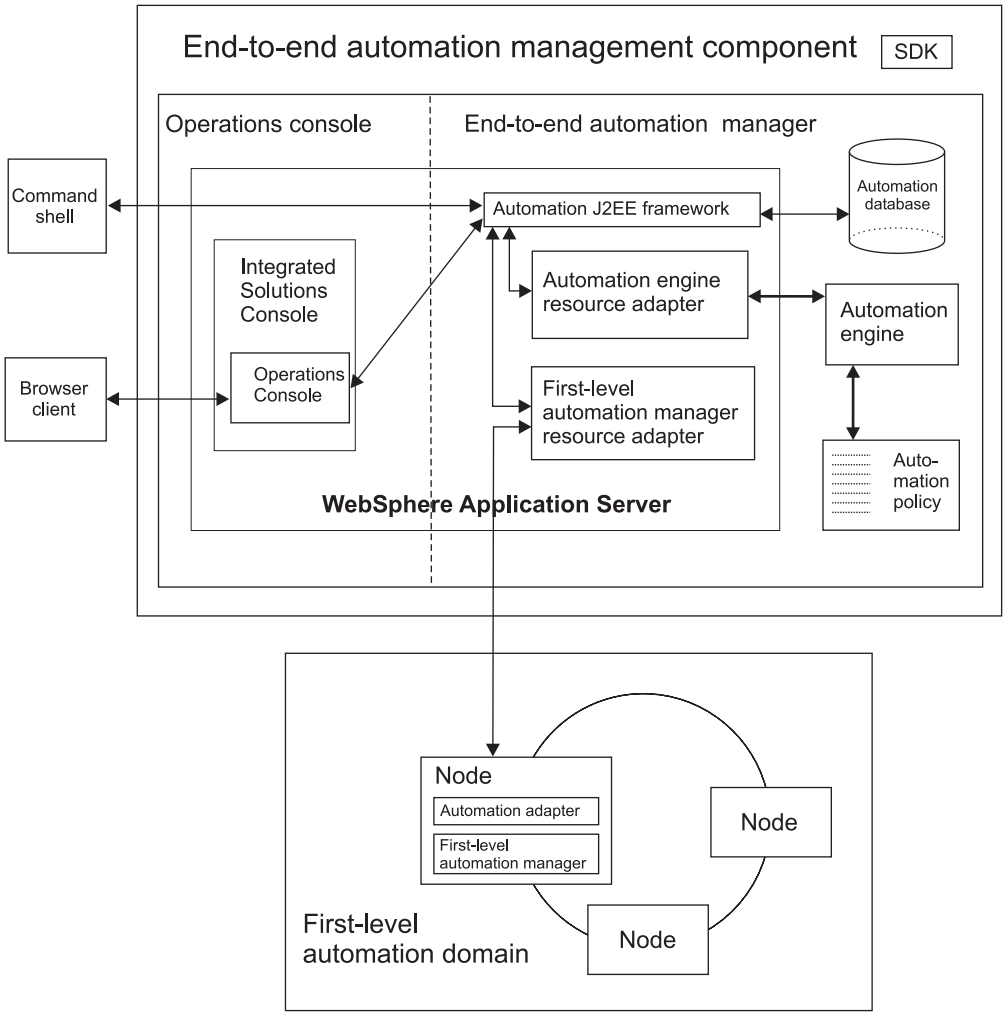

Figure 1. Components of end-to-end automation management

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 manager command shell

The end-to-end automation manager command shell allows you to perform the following tasks by issuing commands to the end-to-end automation manager:

- List resources and resource groups and their states
- List resource group members
- List relationships
- Display, activate, and deactivate policies
- Change the preferred member of a choice group
- Issue online and offline requests against resources and cancel requests
- Reset a resource from an unrecoverable error

You can use the command shell in addition to or instead of the operations console. Using the command shell has the following benefits:

- You can work with end-to-end automation domains from systems where no Web browser is available for displaying the operations console.
- You can use the commands in system scripts or Windows batch files, for example, to monitor or issue requests against resources or to activate a different policy. You can have the scripts launched automatically, for example, by a workload scheduler, such as Tivoli Workload Scheduler, or the cron daemon on UNIX platforms.
- Users who are not working with the operations console on a daily basis may find it easier to use than the operations console.

For information about the end-to-end automation manager command shell, refer to Chapter 24, "Using the end-to-end automation manager command shell," on page 179.

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 chapters 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 operation 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 and the end-to-end automation engine are 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 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 4, “Monitoring and managing automated resources,” on page 121.

First-level automation mode

In this mode, end-to-end automation management is installed but not active. You can use the operations console for monitoring and managing resources of domains that are automated by first-level automation products (Base component of SA for Multiplatforms, SA z/OS, HACMP, Microsoft® Server Clustering (MSCS)).

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 16, "Using the command-line interface of the automation engine," on page 103.

Direct access mode

In this mode, you can use the operations console for monitoring and managing resources that are managed by the following first-level automation products:

- Base component of SA for Multiplatforms
- HACMP
- Microsoft Server Clustering (MSCS)

On the system on which the operations console 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 Administrator's and User's Guide*.

Prerequisites for using the operations console in direct access mode:

- The operations console is installed.
- The automation adapters for the first-level automation domains are configured to send events to the 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.

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 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:

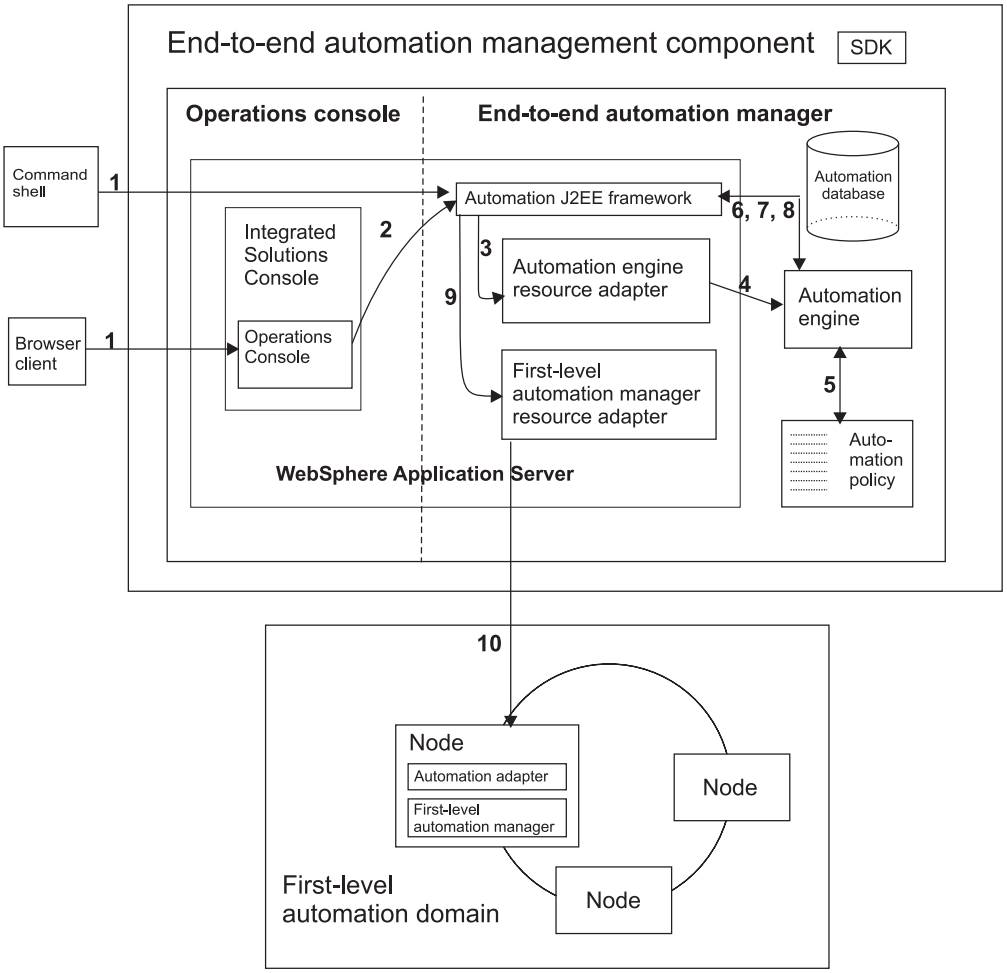

This is a description of the scenario shown in the figure above:

1. An operator requests the activation of an end-to-end automation policy using either the operations console or the end-to-end automation manager command shell.

   _____

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.

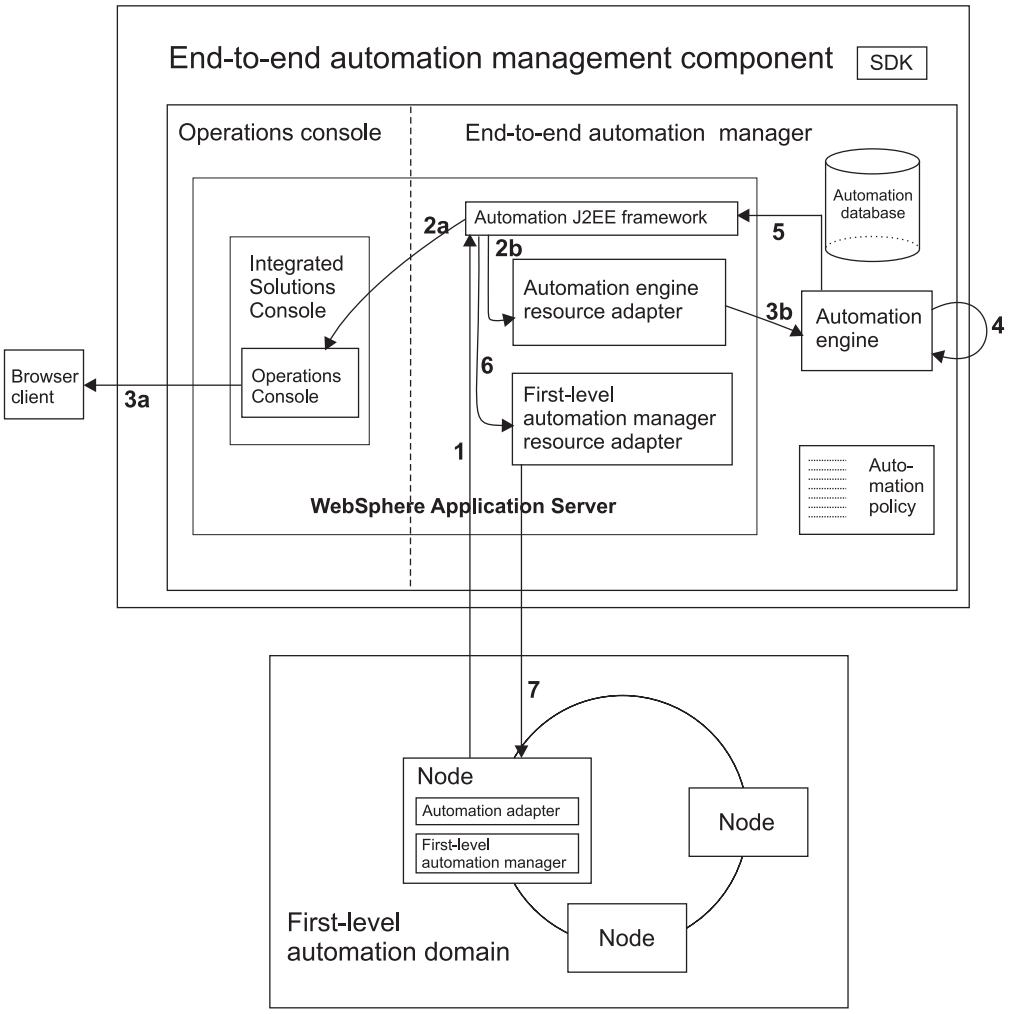

This is a description of the scenario shown in the figure above:

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:

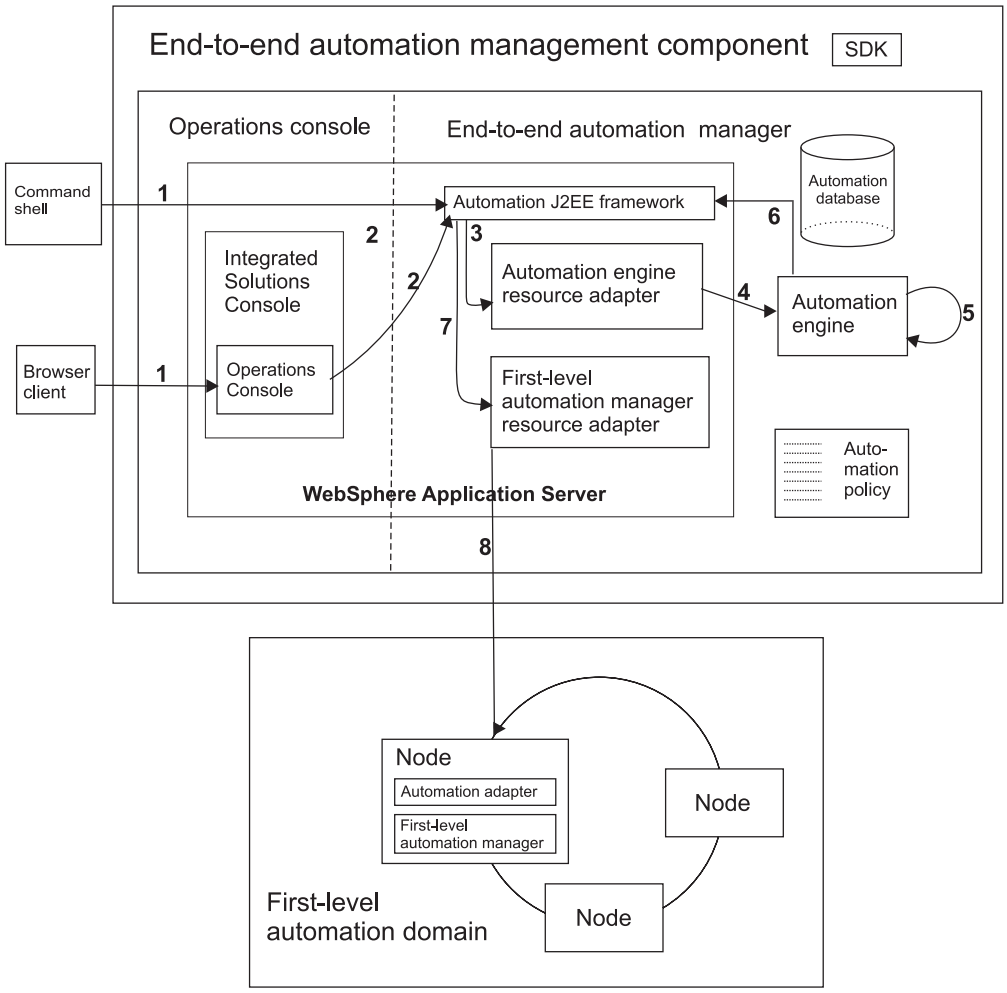

This is a description of the scenario shown in the figure above.

1. An operator submits a request against a resource reference using either the operations console or the end-to-end automation manager command shell.

   _____

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. If automation for the resource reference is not currently suspended, the automation engine calculates all resulting requests (for request-driven first-level automation domains) or commands (for command-driven first-level automation

domains) which must be issued against referenced first-level automation resources. These calculations take into account all relationships defined in the active end-to-end automation policy.

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

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

8. The requests or commands are passed to the first-level automation domains. The first-level automation managers will handle the requests or commands 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.

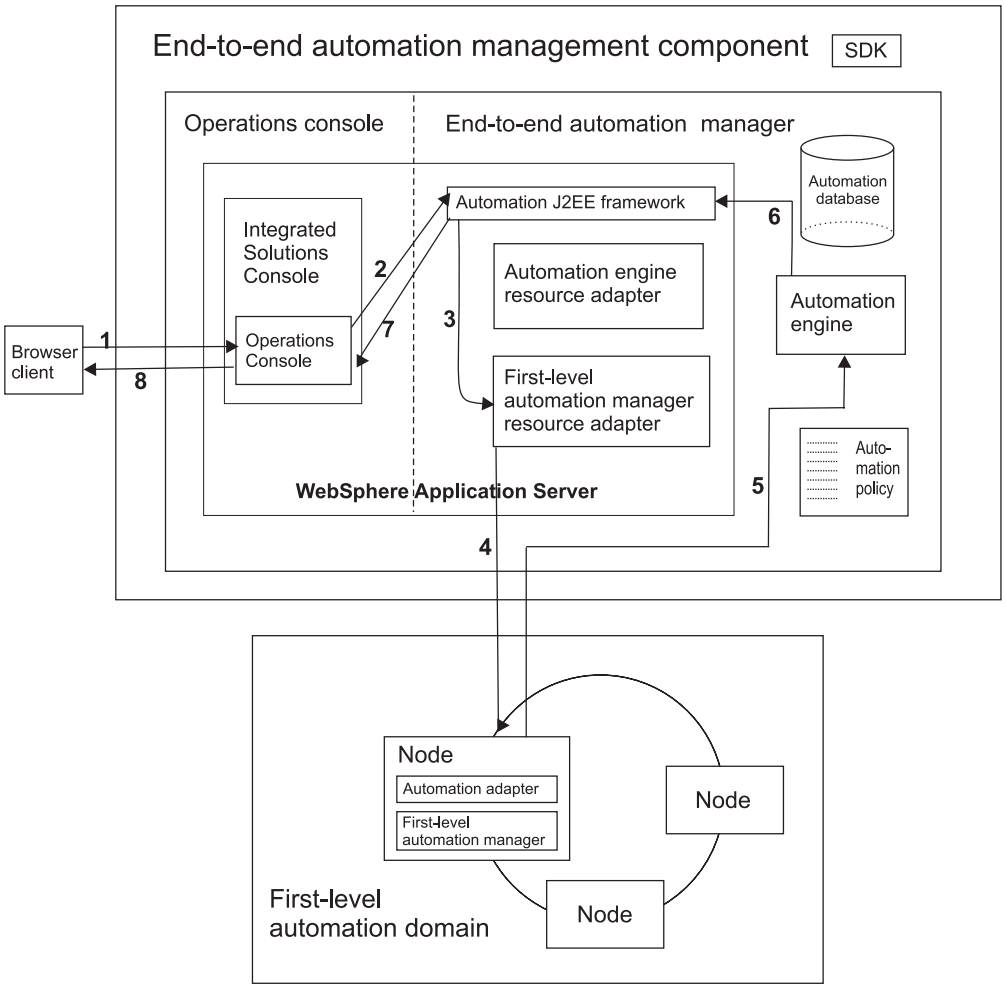

Figure 2. Operations console is used for managing first-level automation domains only

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 143.

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.

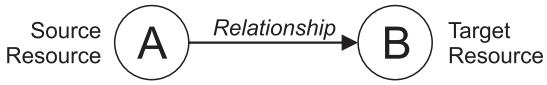

As the arrow in the figure above indicates, relationships always have a direction: In a StartAfter relationship, for example, target resource B would be started before source resource A.

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

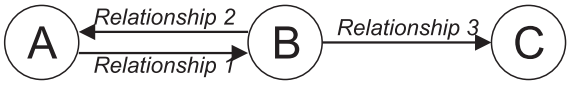

The arrows between the resources in the figure could, for example, represent the following three relationship definitions in the policy:

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:

  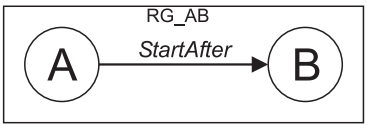

  When the desired state of their resource group is set to online, for example by a start request, both members A and B are started. Due to the StartAfter relationship from A to B, resource B is started first. Once the observed state of resource B is online, resource A is started.
- In the example shown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

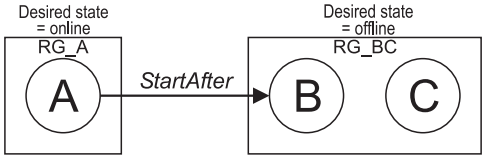

Due to the start sequence defined by the StartAfter relationship, resource B has to be started first. However, because RG_BC's desired state is set to offline, the following conflict exists:
RG_BC wants resource B to be offline whereas the StartAfter relationship forces B to be started. The end-to-end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

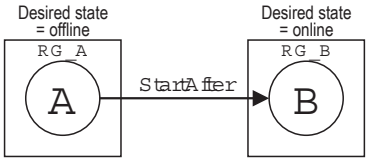

When RG_A's desired state is set to online, resource A can be started right away since resource B is already online.

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

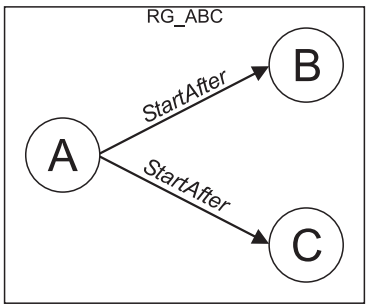

In this case, starting A requires that both resources B and C are online before the end-to-end automation manager can start resource A. If A, B, and C are members of the resource group RG_ABC, setting the desired state of RG_ABC to online causes that resources B and C are started in parallel first. When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

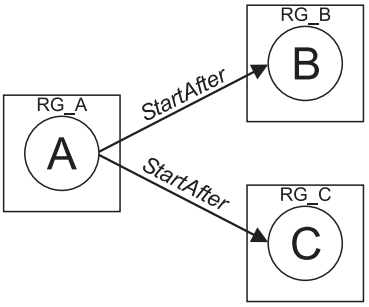

A has a StartAfter relationship to both B and C. Setting RG_A's desired state to online causes that due to the StartAfter relationship resource C and resource B are started. After both resources B and C are online, A is started.

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.

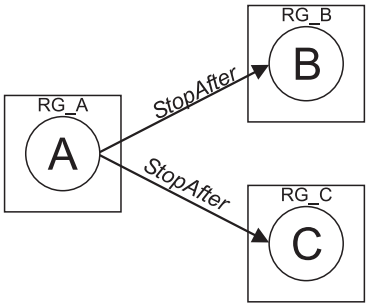

If the desired state of RG_A is online and you want to stop it, RG_A cannot be stopped as long as the desired state of both RG_B and RG_C is online. Only when both RG_B and RG_C have a desired state of offline, resource A can be stopped.

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:

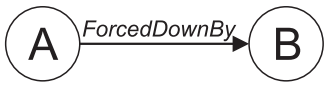

Resource A is forced offline when either the target resource goes offline. The stop of resources A and B can happen in parallel. The force down of resource A will be triggered when resource B enters any of the regular down states (Offline) after having previously been in an Online state or when resource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

2. 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.

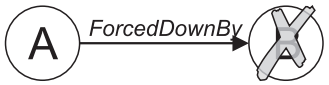

  Both resources are online. In case resource B goes offline, resource A will be forced down.
- In this example, resource A is member of resource group RG_A, and resource B is member of resource group RG_B, and A has a ForcedDownBy relationship with resource B. The force down behavior of the ForcedDownBy 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 an end-to-end automation 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 rank 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 rank of a request is determined by the value of its priority attribute (Prio), its source, and its type (online or offline):

Possible priority values:

Force Overrides requests with any other priority value. The value can only be set using the **resreq** command of the end-to-end automation manager command shell.

High Overrides low priority requests. The value is used as fixed value for requests that are issued from the operations console.

Low The value can only be set using the **resreq** command of the end-to-end automation manager command shell.

Possible sources of a request:

Operator
Default value that is set for requests that are submitted from the operations console or through a command that is issued in the end-to-end automation manager command shell or used in a system script or Windows batch file.

Automation
Default value that is set for requests that are generated by the end-to-end automation manager.

ExtSched
This value can be set for end-to-end automation manager command shell commands that are used in shell scripts or Windows batch files. These scripts are typically launched automatically, for example, by an external scheduler, such as Tivoli Workload Scheduler or the cron daemon on UNIX systems.

To determine the priority ranking of requests that were submitted against a resource, the end-to-end automation manager first evaluates the value of the priority attribute. If multiple requests have the same priority value, the value of the source attribute is evaluated: Operator requests have the highest priority, followed by Automation requests, and finally by ExtSched requests. If the requests could still not be prioritized, start requests take precedence over stop requests.

Table 3 illustrates the priority ranking of requests. The asterisks (*) indicate the default priority of requests that are issued from the operations console or end-to-end automation manager command shell if no priority is specified.

Table 3. Priority ranking of requests

Priority	Source	Request type
Force	Operator | <Other> Operator | <Other> Automation Automation ExtSched ExtSched	Online Offline Online Offline Online Offline
High	Operator | <Other> Operator | <Other> Automation Automation ExtSched ExtSched	Online* Offline* Online Offline Online Offline
Low	Operator | <Other> Operator | <Other> Automation Automation ExtSched ExtSched	Online Offline Online Offline Online Offline

Additional prioritization rules:

- 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.
- 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.

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. If automation for the resource reference is not suspended, the end-to-end automation manager evaluates the request and forwards it 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:

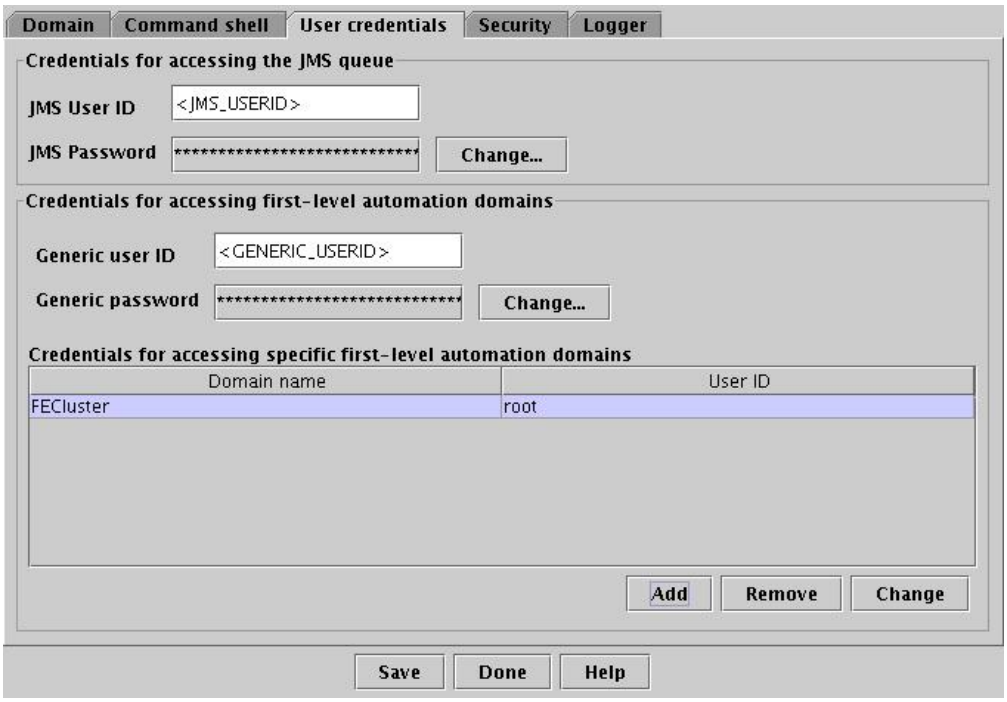

On the User credentials page shown above, specific credentials are only defined for the first-level automation domain FECluster. When the end-to-end automation manager issues requests against referenced resources that are hosted by FECluster, it uses the user ID `root` and the corresponding password.

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:

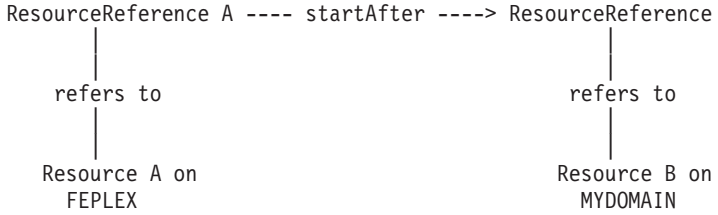

A policy is activated

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 engines uses the user credentials that were specified in the end-to-end automation manager configuration dialog. For information about the configuration dialog, see the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*, section "Configuring the end-to-end automation manager".

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 desire 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
```

Automation is suspended for the resource

When automation for end-to-end automation resources is suspended, the automation manager will not to react on observed state changes by issuing requests against the resource. A state change of a suspended resource can still act as a trigger for state changes of other resources that have a relationship to the suspended resource. This includes that resources having relationships to the suspended resource may still be started or stopped by automation.

If operator requests are submitted against suspended resources, they will be added to the resource's request list but the automation manager will not generate requests against the referenced resources.

For more information about the automation behavior that occurs when automation is suspended, see “Suspending and resuming automation for resources” on page 174.

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 SA for Multiplatforms base component resource group

To list all requests against resource groups in a base component 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 a SA for Multiplatforms base component resource

To list all requests that were issued directly against resources in a base component domain, 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. First steps

Chapter 6. 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 7, "Starting the sample end-to-end automation domain," on page 49)
- How to activate a policy (see Chapter 8, "Activating the sample end-to-end automation policy," on page 51)
- How to create and activate a new policy (see Chapter 9, "Creating and activating a new sample policy," on page 53)
- How to display the first-level automation domains and the resources that are hosted by the domains on the operations console (see Chapter 10, "Displaying a first-level automation domain on the operations console," on page 57)
- Which steps are required to adapt a policy and to activate the modified policy (see Chapter 11, "Creating a policy that references actual first-level resources," on page 59)

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 54.

Chapter 7. 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**:

     Start the automation engine with the following command:

     ```
     eezdmn
     ```

   ---

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 203 for a description of how to determine 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 21, "What you must know about the operations console," on page 127 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 8, "Activating the sample end-to-end automation policy," on page 51.

Chapter 8. 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 21, "What you must know about the operations console," on page 127. The complete list of icons that are displayed on the console is available in the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.
- 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 the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.
- To learn how to create and activate a new policy, perform the tasks described in Chapter 9, "Creating and activating a new sample policy," on page 53.

Chapter 9. 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 14, "Creating and modifying policies," on page 81.

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**:

     Stop the automation engine with the following command:

     ```
     eezdmn -shutdown
     ```

3. Start the end-to-end automation manager configuration dialog and open the Domain page. For information on how to start the configuration dialog, refer to the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*, section "Configuring the end-to-end automation manager".

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 7, "Starting the sample end-to-end automation domain," on page 49.

   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 167.

_____

9. Select the domain “My Domain” in the topology tree to display the new resource “My Reference” in the resource table.

_____

Chapter 10. Displaying a first-level automation domain on the operations console

To work with resources that are hosted by a first-level automation domain from the operation console, you perform the following steps:

1. Check that the user credentials for the first-level automation domain are specified on the User credentials page of the end-to-end automation manager configuration dialog. The end-to-end automation manager needs these credentials to authenticate itself to the first-level automation domain.

   The configuration dialog is described in the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*. For detailed information about the User credentials page, refer to the online help of the configuration dialog.

   ---

2. Check that the automation adapter is configured such that it contacts the end-to-end automation manager. For information about configuring the end-to-end automation adapter of the base component of IBM Tivoli System Automation for Multiplatforms, and the HACMP and MSCS adapters, see the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*. For the SA z/OS adapter, check that the value for **eif-send-to-hostname** is set correctly.

   ---

3. Check that the adapter is running or start it.

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 the troubleshooting section for information on how to resolve the problem.

Chapter 11. 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 53, 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 3. Administering the end-to-end automation management component

Chapter 12. Post-installation tasks for administrators

When the end-to-end automation management component is installed, only the user ID `iscadmin` is authorized for the operations console, where it has unlimited authority. To create and authorize additional users, you must perform the following tasks regardless of whether the operations console will be used in end-to-end 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 65.

   ---

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 66.

   ---

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 68. 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 69.

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 4 on page 64 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 4. WebSphere Application Server access roles for end-to-end automation management

Role	Permissions	Recommended group name
EEZMonitor	Grants minimum access rights. Users who have this role can perform query-type operations. They cannot activate a policy or perform actions that modify the state of a resource, for example, submit start requests.	EEZMonitorGroup
EEZOperator	Extends the EEZMonitor role. Users who have this role can issue requests against first-level and end-to-end automation resources but cannot activate or deactivate policies. Operators managing both first-level and end-to-end automation resources must also have the role EEZEndToEndAccess.	EEZOperatorGroup
EEZConfigurator	Extends the EEZMonitor role. Users given this role cannot submit requests against resources. 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.	EEZConfiguratorGroup
EEZAdministrator	Extends the EEZOperator and EEZConfigurator roles. Users who have this role can perform all operations provided on the operations console for first-level and end-to-end automation domains. Administrators who manage both first-level and end-to-end automation domains must also have the role EEZEndToEndAccess. **Note:** The user ID of the automation engine must have the role EEZAdministrator. You achieve this by adding the user ID to the EEZAdministratorGroup.	EEZAdministratorGroup
EEZEndToEndAccess	User who do not have this role, can view and monitor the end-to-end automations domain and the resources hosted by the domain. This role is only required if a user needs to start or stop end-to-end automation resources or activate and deactivate policies. 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. **Note:** The user ID of the automation engine must have the role EEZEndToEndAccess. You achieve this by adding the user ID to the EEZEndToEndAccessGroup.	EEZEndToEndAccessGroup

Table 4. 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 67

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:

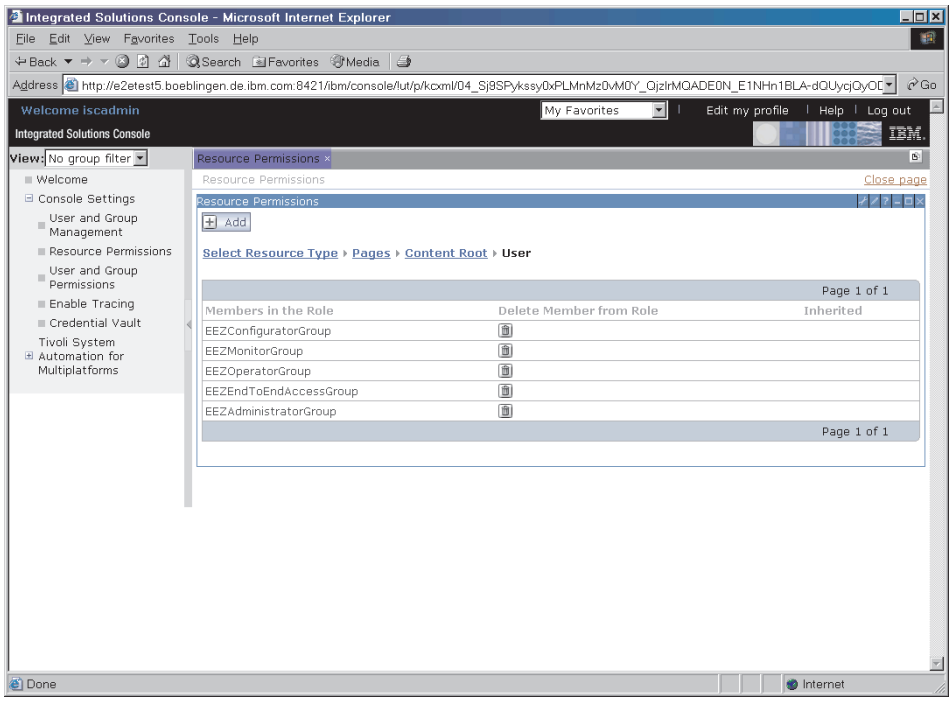

11. Restart Integrated Solutions Console.

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 5. 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.

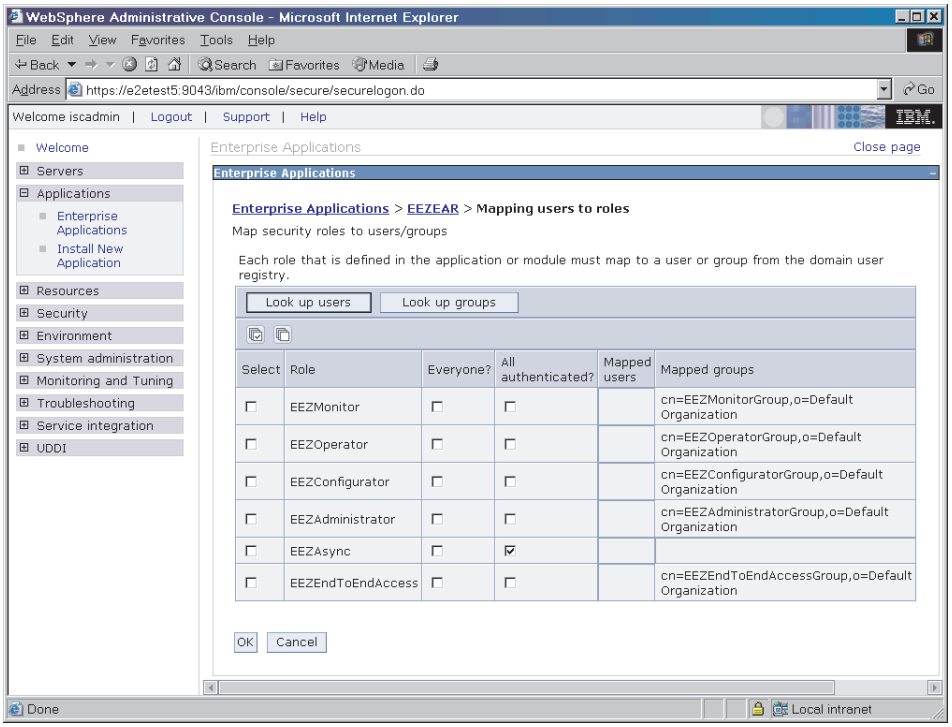

Attention: After you have completed this step, only mapped groups are authorized to work with the operations console. In particular, this means that the group `iscadmins`, which is managed by Integration Solution Console, is not longer authorized. To authorize the user ID `iscadmin` for the operations console, you must add it to an authorized group.

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 13. 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 12, "Post-installation tasks for administrators," on page 63.
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 164 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 63.
- 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 73 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 user authentication for command shell users

The end-to-end automation manager requires authentication when a user invokes the end-to-end automation manager command shell. The end-to-end automation

manager supports three authentication modes. On the Command shell page of the end-to-end automation manager configuration dialog you can select the desired authentication mode and, if you are using a shared user ID for authentication, change the password for the user ID.

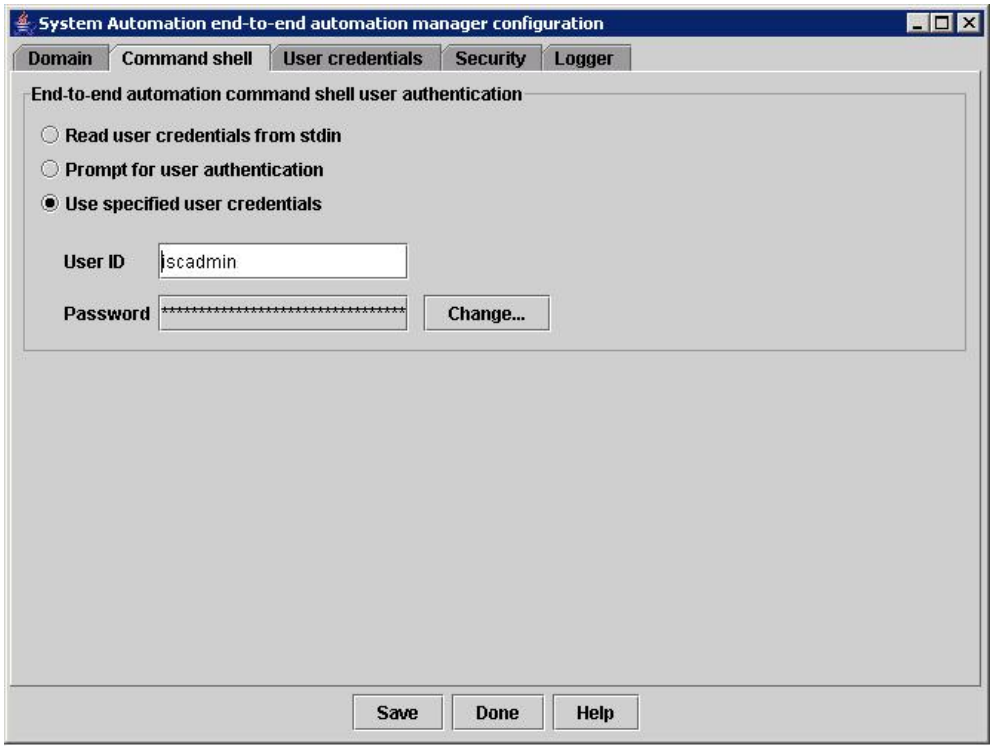

Figure 3. Command shell page of the end-to-end automation manager configuration dialog

The configuration dialog is described in the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*, chapter "Configuring the end-to-end automation management component").

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 end-to-end automation manager configuration dialog, see the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*. 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 end-to-end automation manager configuration dialog, see the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*. 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 14. 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 82
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 82

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. 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.	"The scope of end-to-end automation policies" and Appendix A, "Policy definition worksheet," on page 197
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 88

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.

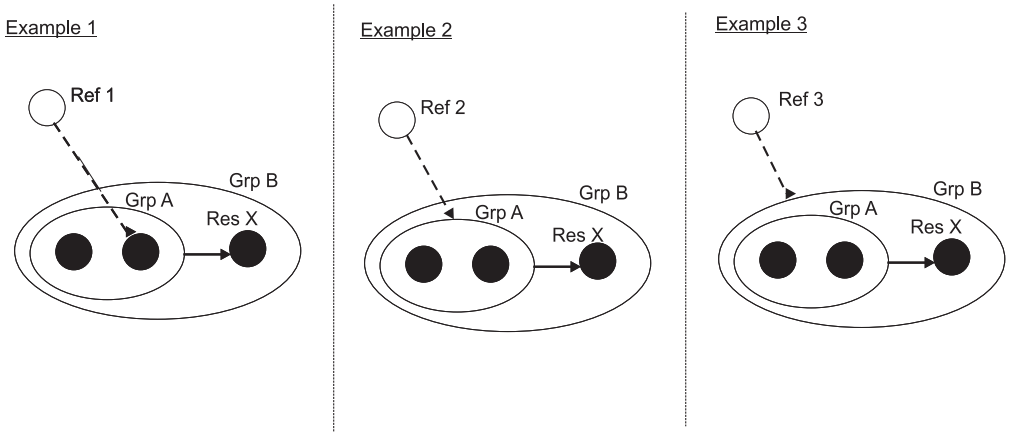

The examples in the figure above show three resource references that were created for resources or resource groups that are hosted by a first-level-automation domain. These examples are described in the following sections.

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 send 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.

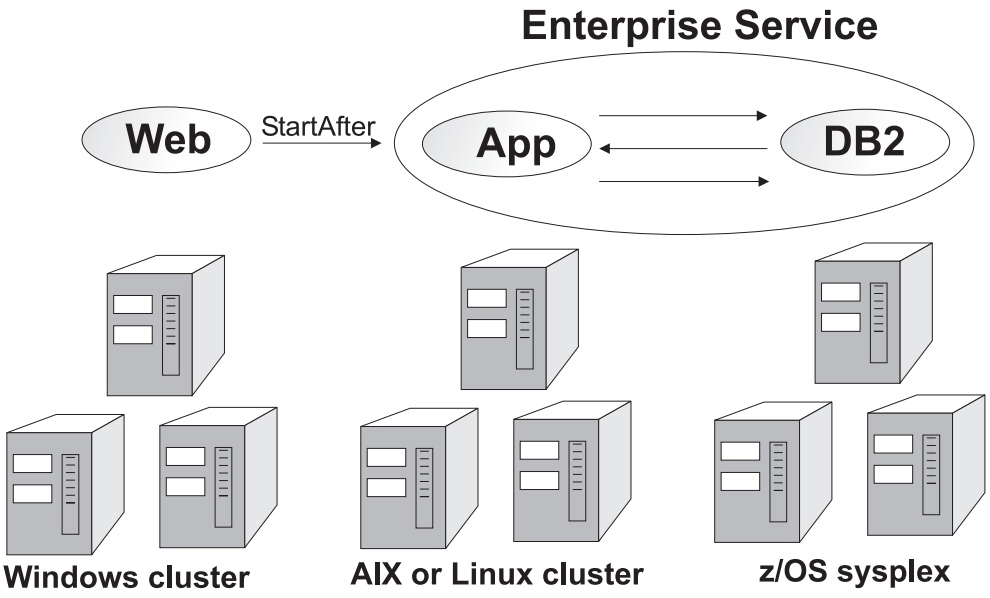

Relationships

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.

When you define a resource reference in an end-to-end automation policy, you must provide information about the first-level resource in the `<ReferencedResource>` subelement. You can easily obtain all the required

information on the operations console by displaying the General page for the first-level resource.

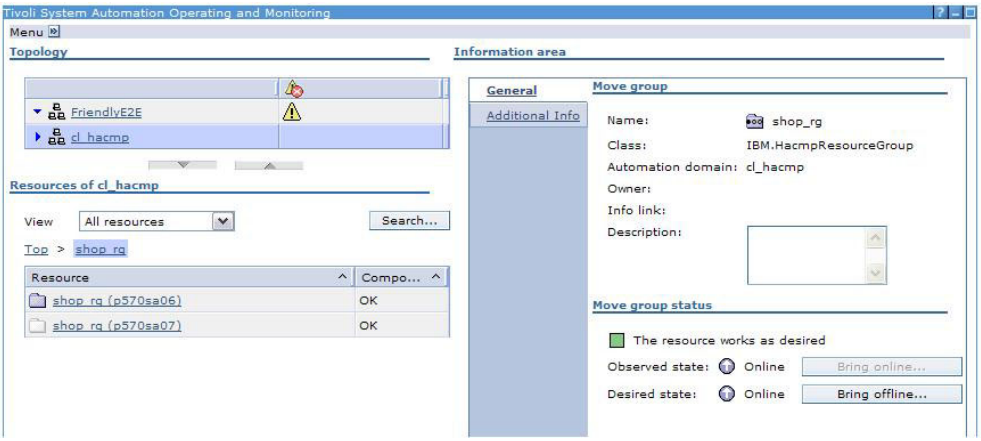

Figure 4. General page for a first-level resource

This is how you display the information for a first-level resource on the operations console:

1. Make sure that the adapter for the first-level domain whose resources you want to reference in the policy is correctly configured and running.
2. Open the operations console and select the first-level automation domain.
3. Select the first-level resource you want to reference in the automation policy.
4. Open the General page for the resource.
5. In the end-to-end automation policy, specify the information in the policy exactly as it appears on the page. Especially, if no node information is provided on the General page, do not specify the `<Node>` element in the end-to-end automation policy.

A worksheet for gathering the data you need for defining a policy is available in Appendix A, "Policy definition worksheet," on page 197.

Considerations for referencing first-level automation resources

The sections below list the considerations that apply when you create resource references for resources that are managed by first-level automation products of the IBM Tivoli System Automation product family.

For considerations that apply when you reference resources that are managed by other first-level automation products, refer to Part 5, "Working with the HACMP and MSCS adapters," on page 181.

Considerations for referencing SA for Multiplatforms base component resources

When you create resource references for SA for Multiplatforms base component 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 82.

Table 6. 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 7. 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" on page 89
2	Define the resources of the end-to-end automation domain	"Defining the resources of the end-to-end automation domain" on page 92
3	Define resource groups and choice groups	"Defining groups" on page 94

Table 7. Steps for defining a new end-to-end automation policy (continued)

4	Define StartAfter, StopAfter, and ForcedDownBy relationships	"Defining StartAfter, StopAfter, and ForcedDownBy relationships" on page 98

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. Relationships

   You 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 the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.
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>APG</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>APG</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 96.

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 167. 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.

   _____

2. On Windows systems, change the directory to `EEZ_INSTALL_ROOT/bin`.

   _____

3. Issue the following command to start the tool:

   On Windows: `eezpolicychecker.bat <policy_file_name>`

   On AIX and Linux: `eezpolicychecker <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 15. 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 16. Using the command-line interface of the automation engine

This section describes how you use the command-line interface of the automation engine. For information about the end-to-end automation manager command shell, see Chapter 24, "Using the end-to-end automation manager command shell," on page 179.

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. On Windows systems, 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:

   ```
   eezdmn <option>
   ```

   For example:

   ```
   eezdmn -shutdown
   ```

   Table 8 on page 104 provides an overview of the available options. A detailed description is provided in the following sections of this chapter.

Note: If the automation engine is running on a Windows server, it will be 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 8. 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. For information about the end-to-end automation manager configuration dialog, see the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*.

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 9. 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	The automation engine is running. Before end-to-end automation can start, a policy must be activated. 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.
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 9. 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 end-to-end automation manager configuration dialog. (For information about the configuration dialog, see the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*.)
- 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: 2.2.0.051501, NO_APAR
Usage:

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>]") <DUMPFILE>
                       Dumps (all | specific) resources to a file

When no option is specified, start is used
```

Chapter 17. 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 110 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 110.
- End-to-end automation manager configuration dialog

  For information about starting the configuration dialog, see the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*.
- 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 16, "Using the command-line interface of the automation engine," on page 103.

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.

Starting and stopping the operations console on Windows

How you start and stop the Integrated Solutions Console server and the Eclipse Helps 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 16, “Using the command-line interface of the automation engine,” on page 103.

Chapter 18. 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 in the WebSphere Application Server administrative console:

1. Activate the Common Event Infrastructure (CEI) service when `server1` is started:
   a. Click **Servers --> Application servers --> server1** --> **Container Services --> Common Event Infractructure Service**.
   b. Select the check box **Enable service at server startup**.

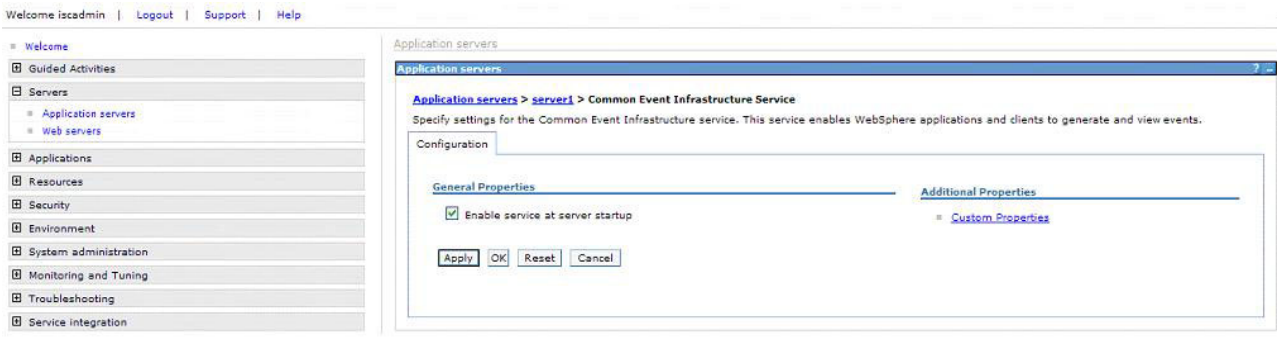

Figure 5. Common Event Infrastructure Service panel

   c. Save the Master configuration and restart the WebSphere Application Server `server1`.

Note: Alternatively, you can start both CEI services manually:
   a. In the administrative console, click **Applications** --> **Enterprise Applications**.
   b. Start the applications EventServer and EventServerMdb.

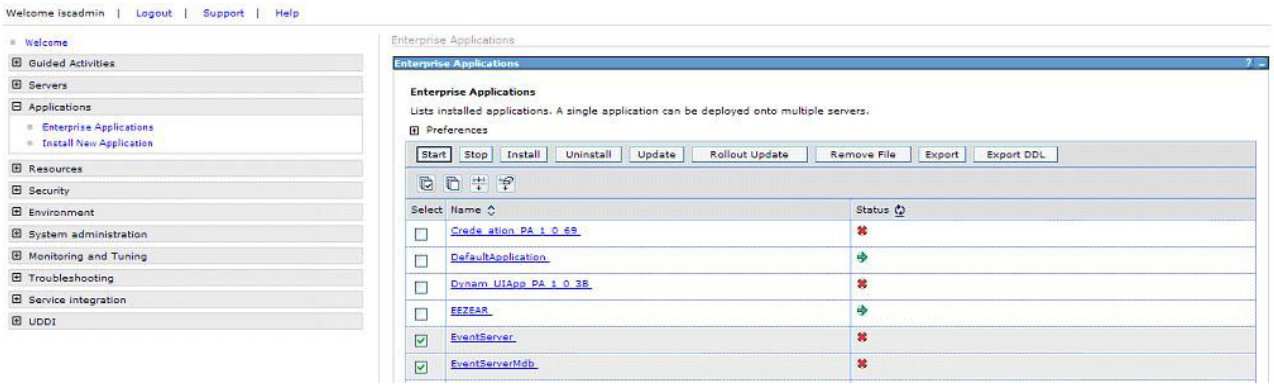

Figure 6. Enterprise Applications panel

2. Configure the Tivoli Enterprise Console and install the `baroc` file:
   a. In the administrative console, click **Resources --> JMS Providers --> Generic --> EIF JMS Provider --> JMS connection factories --> EEZTECSenderQCF --> Custom properties** (see Figure 7).
   b. Set **ServerLocation** to your TEC server name.
   c. Set **ServerPort** to the following values:
      - The TEC server runs on Windows: `5529`
      - The TEC server runs on AIX or Linux: `0`
   d. Install the file `SystemAutomation.baroc` on your TEC server.

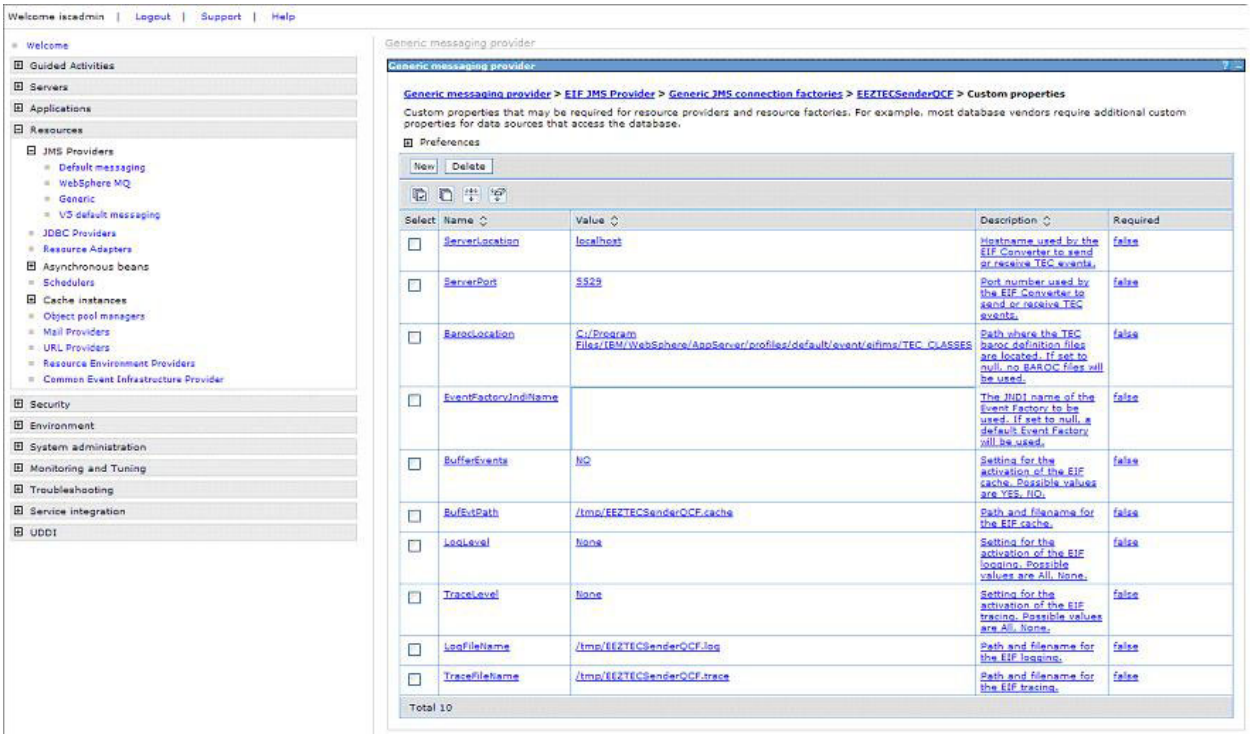

Figure 7. Custom properties panel

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/<profilname>/event/bin/emitevent.jacl
           -xml ../profiles/<profilname>/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:

   ```
   wtdumprl
   ```

   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.

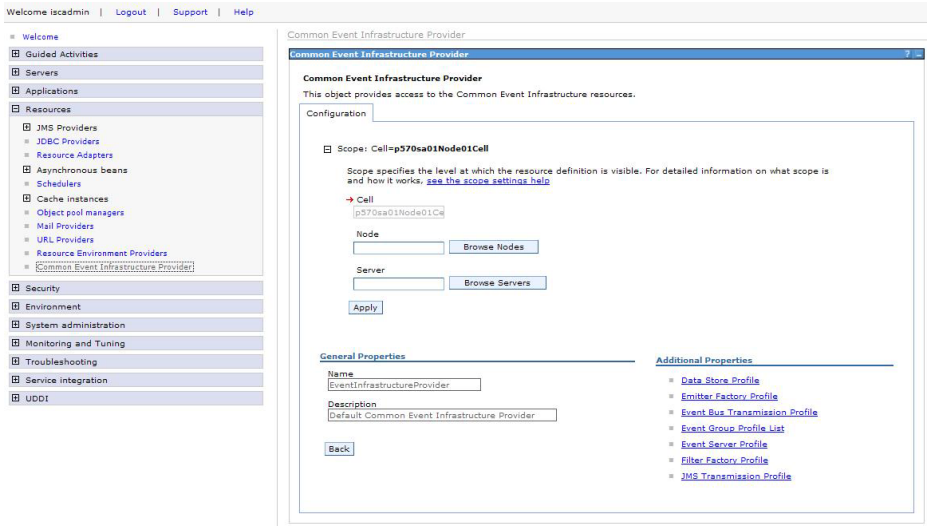

To activate the default event filter, perform the following steps:

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 10. 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.

Table 10. Valid XPath event selectors (continued)

XPath event selector	Description
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 4. Monitoring and managing automated resources

Chapter 19. 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 this part of the guide:

- The mode in which you are running the operations console and the capabilities of the automation domain you are working with determine which actions you can perform on the operations console. For an overview of the operations console modes, refer to Chapter 3, "Operations console modes," on page 15. The domain capabilities are outlined in Chapter 20, "Domain capabilities," on page 125.
- Chapter 21, "What you must know about the operations console," on page 127 gives you an overview of the information provided on the operations console and the actions you can perform.
- Chapter 22, "Monitoring resources," on page 143 describes how you can use the operations console to monitor resources and to analyze and resolve the problems that may occur.
- Chapter 23, "Managing resources," on page 167 describes how you start and stop resources from the operations console, suspend and resume automation for a resource, 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.
- Chapter 24, "Using the end-to-end automation manager command shell," on page 179 describes how to use the end-to-end automation manager command shell.

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 20. Domain capabilities

The actions you can perform on resources from the operations console are determined by the capabilities of the hosting automation domain, which in turn are determined by the capabilities of the automation product automating the domain.

The following table lists the domain capabilities by automation product:

Automation product	Request-driven (1) / Search resources with operator requests (2)	Search by name (3)	Search by class (4)	Suspend automation for resources (5)
SA MP End-to-end automation management component	Y	Y	Y	Y
SA MP Base component	Y	Y	Y	N
SA z/OS	Y	Y	N (depends on the SA z/OS release)	N
HACMP	N	Y	Y	N
Microsoft Server Clustering (MSCS)	N	Y	Y	N

Notes:

(1) Only request-driven automation domains maintain request lists for all resources, in which the requests are stored until they are canceled and which is analyzed to determine the winning request whenever a new request is submitted or an existing request is canceled.

On the operations console, request-related controls (for example, the buttons **Request online** and **Request offline** for starting and stopping resources, and the button **View requests** for displaying the request list) are only available for resources that are hosted by request-driven automation domains. For more information, refer to "Working with requests" on page 169.

For command-driven automation domains, the buttons **Bring online** and **Bring offline** are available for starting and stopping resources. For more information, refer to "Bringing resources online and offline" on page 172.

(2), (3), (4)
The entries in the table show which filter criteria can be specified on the Search panel. For example, some domains do not allow searching for resources against which operator requests were submitted. For more information, refer to "Searching for resources" on page 157.

(5) Suspending automation for a resource causes the automation manager not to react on observed state changes by issuing requests against the resource. For more information, refer to "Suspending and resuming automation for resources" on page 174.

Chapter 21. 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 you 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:

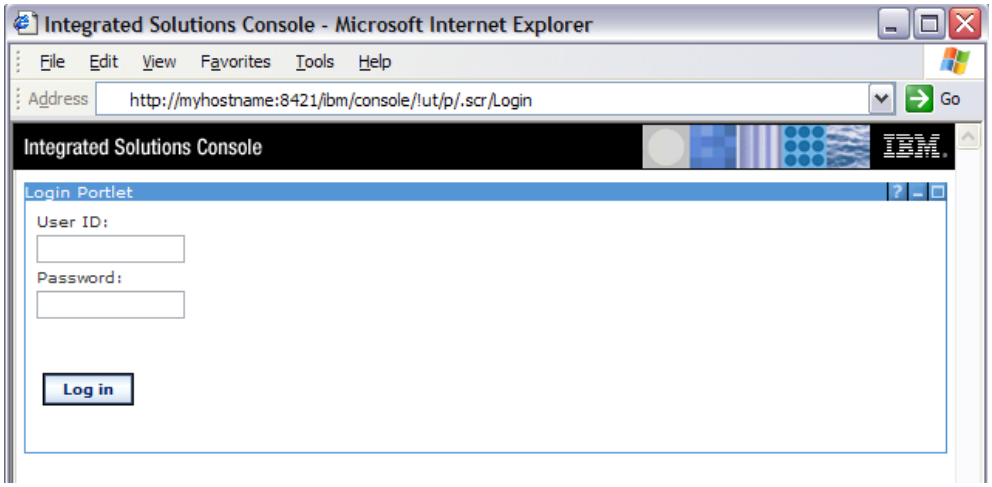

Figure 8. Log in panel of Integrated Solutions Console

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:

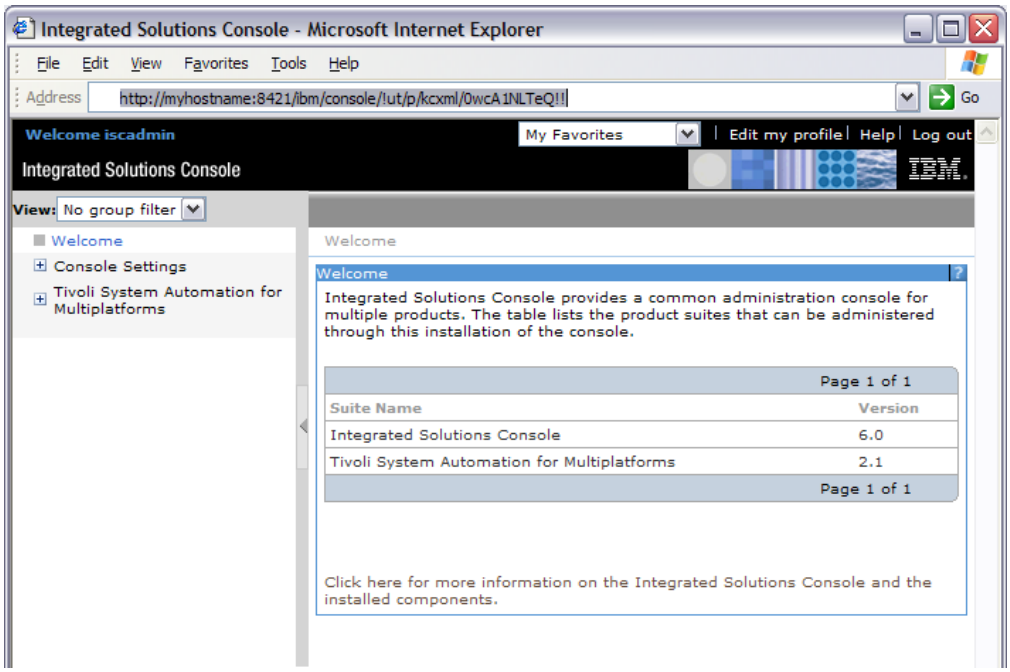

Figure 9. Welcome panel of Integrated Solutions Console

3. In the navigation tree on the left, expand the folder **Tivoli System Automation for Multiplatforms**:

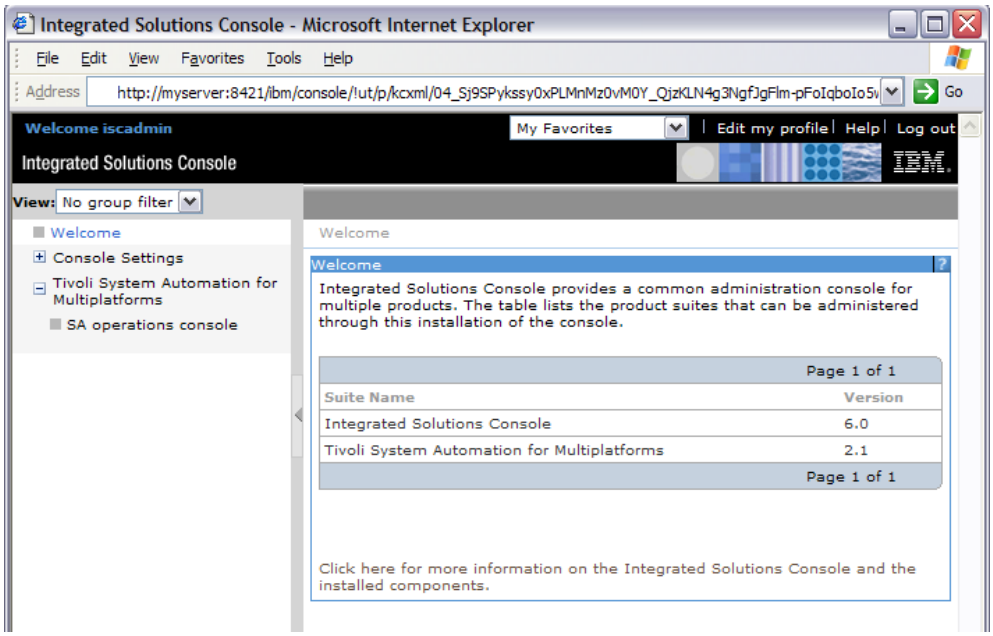

Figure 10. Operations console entry in the navigation tree

4. Click **SA operations console**.
   The Connect panel is displayed:

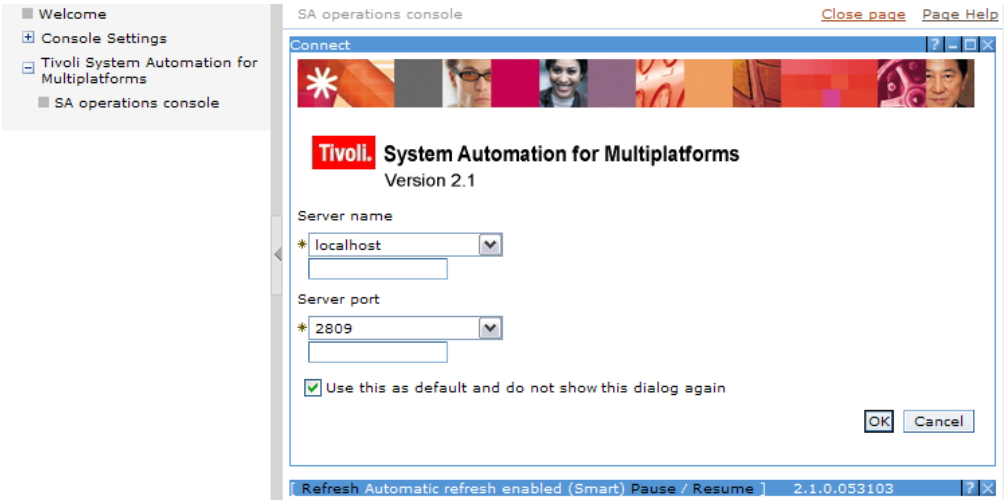

Figure 11. Connect panel

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 203.

   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:

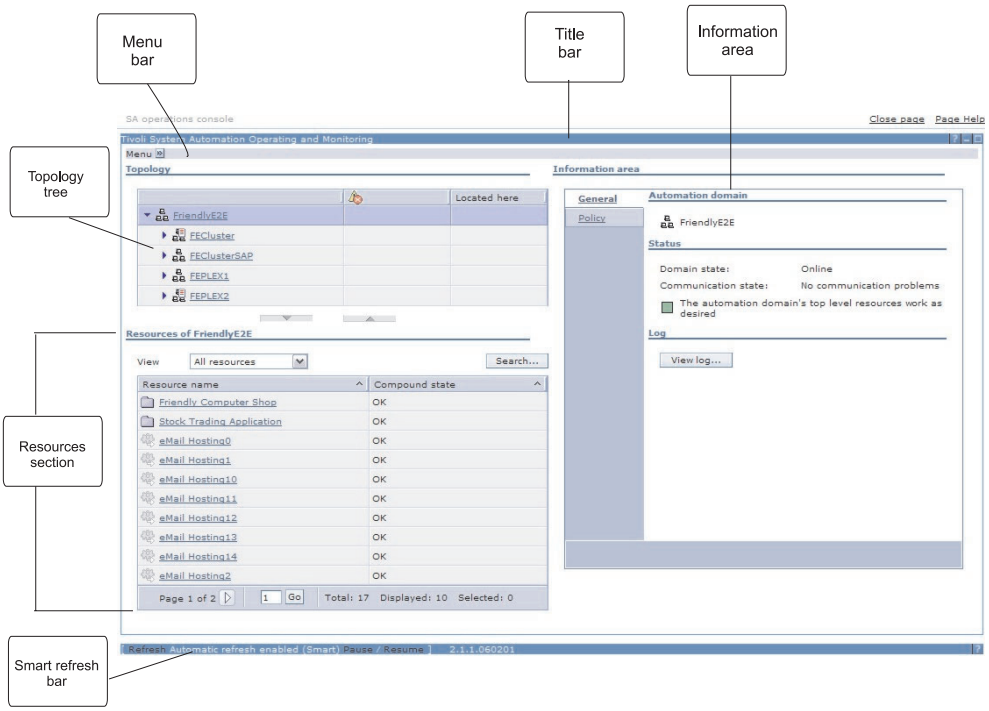

Figure 12. Main panel of the operations console

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 141.

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.

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 135.

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 139.

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 137.

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 140.

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 141.

What you must know about the topology tree

The following figure shows the topology tree and the resources section.

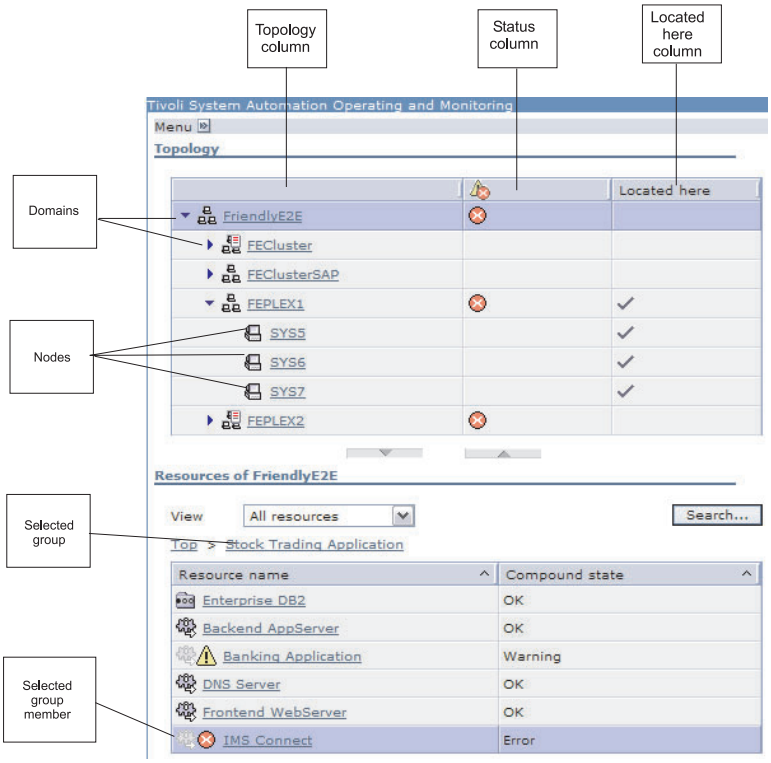

Figure 13. Topology tree and resources section

The topology tree is divided into three columns (see Figure 13):

- 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 134).
- The **Status** column shows the health status of the domain (see "What you can see in the Status column" on page 135).
- 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 135).

Figure 13 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 161).

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 11. 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. The complete list of icons is available in the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

Table 12. 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 12. 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 162).

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 13. 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 13 on page 133).

What you must know about the resources section

The following figure shows the layout of the resources section.

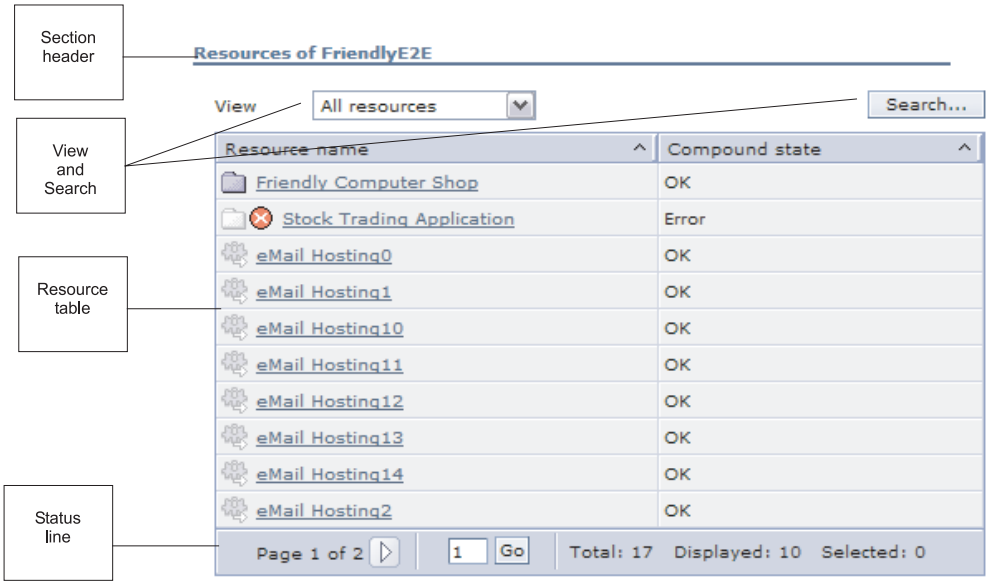

Figure 14. Layout of the resources section

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:

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

Allows you to display only resources that meet specific search criteria (see "Searching for resources" on page 157 for more information).

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 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.

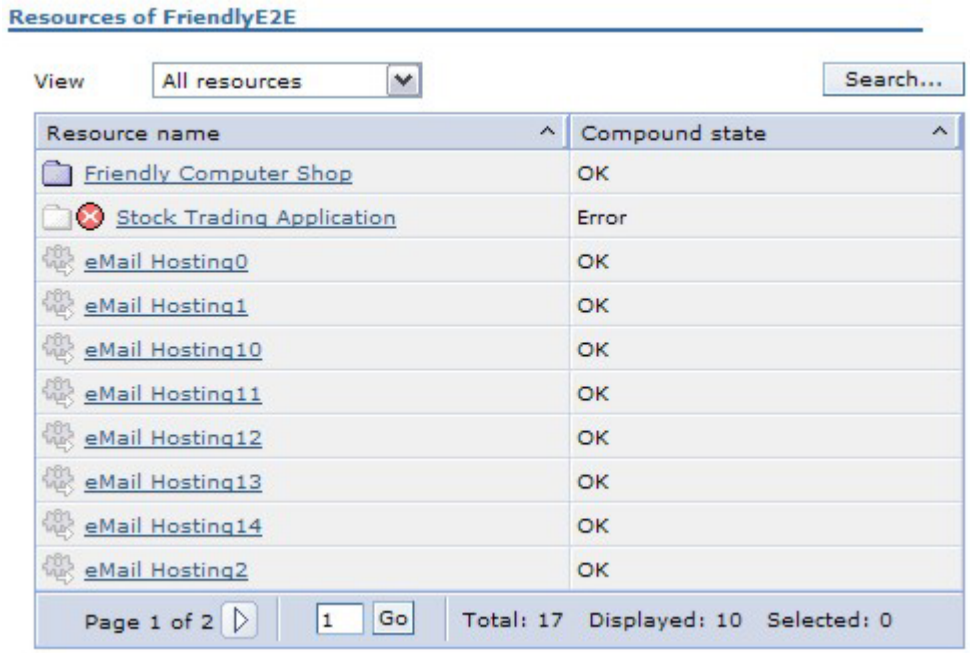

When you select a group in the resource table, the members of the group are displayed in the resource table. In the area above the table, a bread crumb trail appears. On the trail, the name of the group whose members are listed in the resource table is highlighted, indicating that the group is selected.

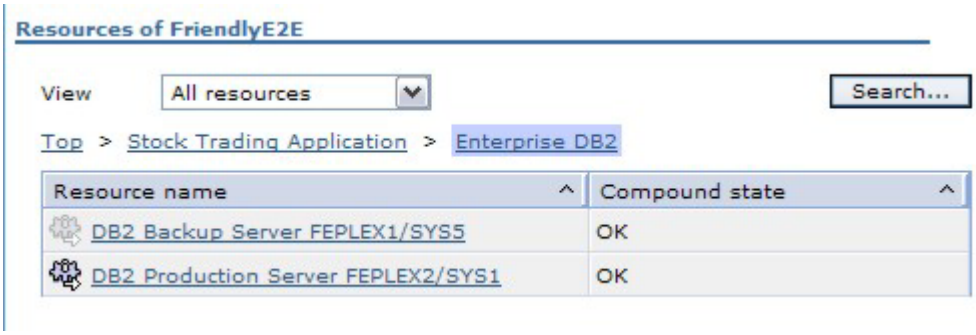

The bread crumb trail is useful for navigation and orientation:

- 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 column: The Resource 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
	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 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.

Icon	Description
	The black error icon indicates that the resource has encountered an unrecoverable error.

For the complete list of icons that appear in the operations console, refer to the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

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:

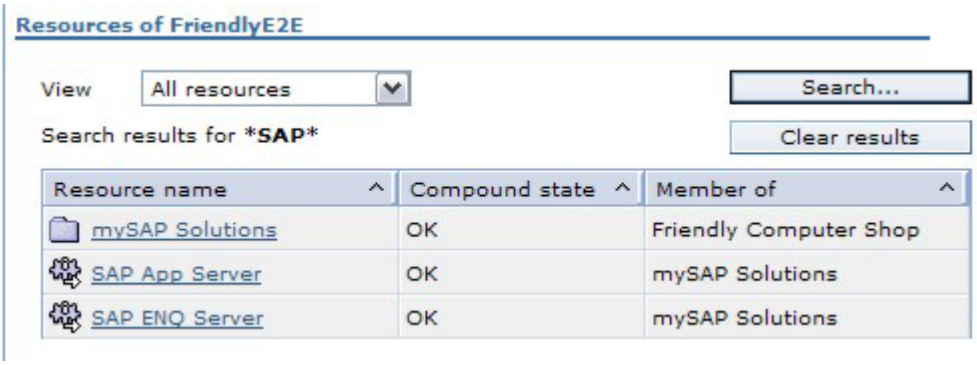

To limit the scope of resources that are currently displayed in the resource table to those that are in an error or warning state, you can additionally apply the **Errors and warnings** view that is provided in the **View** field.

Resource table columns: In the search results view, the resource table has three columns:

Resource 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	• General • Policy
a first-level automation domain in the topology tree	• General • Policy • Additional Info

When you select ...	...these pages are available
a resource or group in the resource table	• 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 the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*, appendix "Additional state-related information about end-to-end automation resources".

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 163.

What you must know about the main menu

The main menu is available on the menu bar of the operations console.

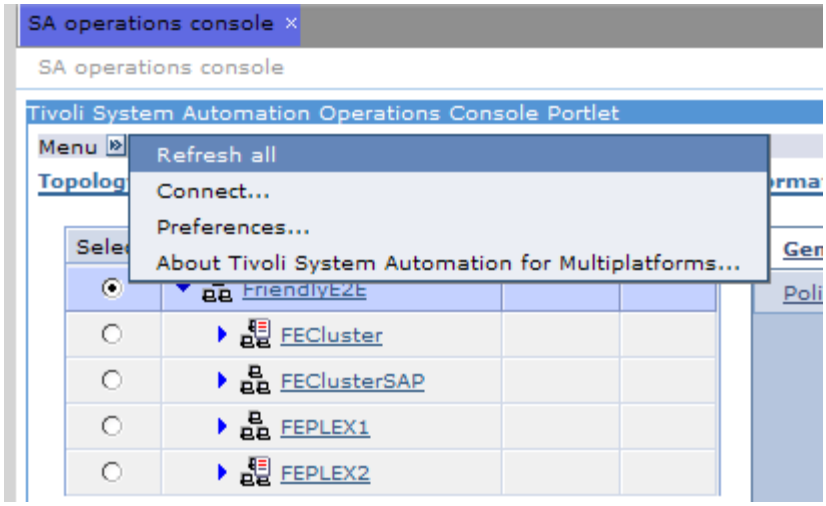

Figure 15. Main menu

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 164).
- 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 "Working with name filters" on page 158).

  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 162).
- 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 161).

View Allows you to adapt the view of the topology tree and the resource table to your screen resolution (see "Customizing the view").

About Displays information about the version of the component you are using.

Customizing the view

Use the View page (**Menu —> Preferences —> View**) to change the number of entries that are displayed in the topology tree and the resources section when both are visible. This is helpful when the current values are inadequate for your screen resolution. The number of entries that are displayed when one of the navigation sections is hidden is adapted automatically.

To change the values, perform the following steps:

1. In one or both fields on the View page, specify how many entries are to be displayed by default. The valid range is 5 through 100.
2. Click **OK** to save you changes.

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 22. 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 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 14. Compound state icons

Icon	Example	Description
		Compound state: **Warning** 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. 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.
		Compound state: **Error** 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.
		Compound state: **Fatal** 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. **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 173).

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:

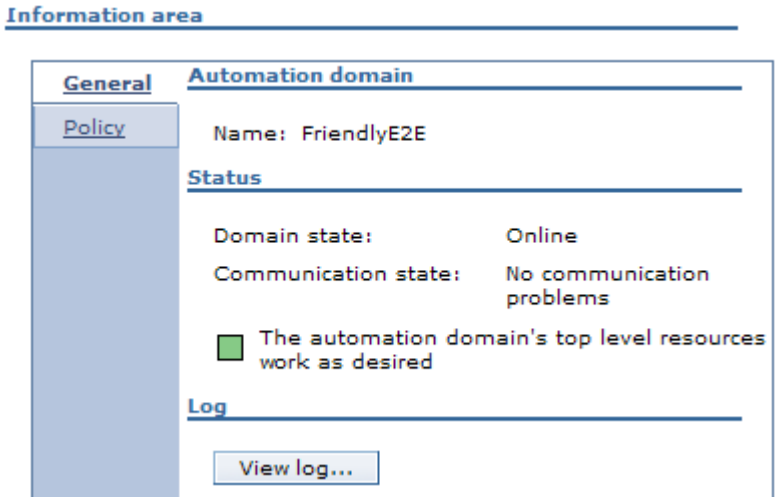

Operational state descriptions provided on the General page

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 143.

Table 15. 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 15. Operational state descriptions provided on the General page for a domain (continued)

Description on the General page	Troubleshooting
The domain contains top-level resources with errors. At least one resource matching the name filter <current domain filter> has an error.	**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. **What you can do:** Find out which resource is affected and analyze the problem, for example: • 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 15. Operational state descriptions provided on the General page for a domain (continued)

Description on the General page	Troubleshooting
The domain contains top-level resources with unrecoverable errors. At least one resource matching the name filter <current domain filter> has an unrecoverable error.	**What it means:** At least one of the resources you are using as domain health indicators has encountered an unrecoverable problem. **What you can do:** Find out which resource is affected and analyze the problem, for example: • 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. 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. After resolving the problem, you must use the Reset function to include the resource in automation again.

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 16. 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 icons is available in the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

Table 17. Communication state

Communication state icons and state descriptions	**What it means**
 No communication problems.	No action is required.
 Commands and queries can currently be issued against this domain, but at least one resource event was lost.	The state information provided for the domain's resources may be outdated. Perform a **Refresh all** to update the information. 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.
 No commands or queries can currently be issued against this domain, but resource events are still received.	One of the following problems may have occurred: • 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 17. Communication state (continued)

Communication state icons and state descriptions	What it means
 Commands and queries can currently be issued against this domain, but no resource events are received.	The state information provided for the resources that are hosted by the domain may be outdated. Perform a **Refresh all** to update the information. The configuration properties of the domain may need to be changed. Inform the system administrator of the domain.
 None of the communication paths to this domain are currently working.	No queries can be submitted, no events can be received. The resource state information may be outdated. No refresh is possible. Check if the adapter has failed.
 No commands or queries can currently be issued against this domain and at least one resource event was lost.	The state information provided for the resources of the domain may be outdated. View the log files manually for further information.
The automation adapter is currently not running.	The adapter may have been stopped intentionally by an administrator.

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 the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

Table 18. 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.

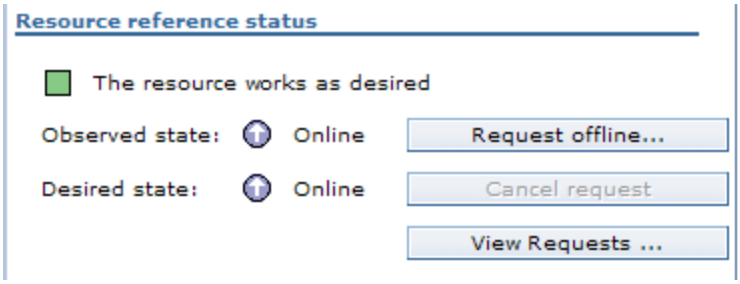

Figure 16. State information on the General page

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 143.

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 19. 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 19. 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 19. 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 16 on page 150). 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 13 on page 133).

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 200.

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

Use the Search panel to display only resources that meet specific search criteria. The resources will be displayed in the search results view of the resource table.

Submitting a search

To submit a search, perform the following steps:

1. Select a domain or node.

---

2. Click the **Search** button above the resource table. The Search panel is displayed.

---

3. Specify the search criteria for the resources you want to display.

---

4. Click **OK** to submit the search.

Results:
The resources that match the search criteria are displayed in the search results view of the resource table. If you specified a new search phrase in the Resource name section, the search phrase is saved as a name filter for the domain and becomes available in the Resource name drop-down list on the Search panel for the domain or any of its nodes.

Note: Search results are not refreshed automatically. To refresh, clear the search results and perform the search again.

Search panel sections and controls

The search criteria you can specify on the panel vary depending on the capabilities of the selected domain, or, if you selected a node, on the capabilities of the domain to which the node belongs. You can specify any, multiple, or all search criteria that are available.

Resource name section

Allows you to specify a search phrase to display only resources whose names contain the phrase.

You have the following options:

- To use an existing search phrase, select the search phrase from the drop-down list.
- To enter a new search phrase, select **Use entry from below** from the drop-down list and type the search phrase in the field below. Search phrases can have the following syntax:
  - Type the exact resource name to display a specific resource.

- Use the asterisk * as a wildcard to display all resources whose names contain the search phrase. The wildcard can appear in any position and, if necessary, more than once (for example, *DB2*), and can stand for 0..n characters.
- To display all the resources that contain at least one of several search phrases, type all phrases separated by a blank; the wildcard can be used in one or all phrases (for example, *DB2* SAP*).
- For resource names that may contain blanks, type the complete search phrase including the blank and enclose the phrase in single or double quotation marks, for example, "*SAP *Server". This ensures that it will be recognized as a single phrase.

Resource class section

- **Search for any resource class**
  This option is selected by default. If selected, the resource class is not used as search criterion.
- **Search only for selected resource classes**
  Allows you to search for resources by resource class type. To specify a class type, select the appropriate check box.
- **Search for resource classes matching the following search pattern**
  Allows you to specify a search phrase to display only resources whose class names contain the search phrase. Note that this option is not available for all automation domains, even if searching by resource class name is otherwise allowed.

  Search phrases can have the following syntax:
  - To search for resources of a specific resource class, type the exact class name.
  - Use the asterisk * as a wildcard to display resources whose class names contain the search phrase. The wildcard can appear in any position and, if necessary, more than once, and can stand for 0..n characters.
  - To display all the resources whose resource class names contain at least one of several search phrases, type all phrases separated by a blank; the wildcard can be used in one or all phrases.
  - For resource class names that may contain blanks, type the complete search phrase including the blank and enclose the phrase in single or double quotation marks. This ensures that it will be recognized as a single phrase.

Miscellaneous section

Select the check box to search for resources against which operator requests have been submitted.

Note that selecting the check box is only valid for request driven domains and has no effect for command-driven domains.

Working with name filters

Name filters are search phrases that you use to display only resources in the resource table whose name contains a search phrase. Typically, you specify these search phrases in the **Resource name** section on the Search panel, which appears when you click the Search button above the resource table. When you enter a search phrase on the Search panel and submit the query, the search phrase is saved as name filter and is from then on available for the domain and all of its nodes in the **Resource name** drop-down list on the Search panel until you delete it.

This topic describes how you define, edit, and delete name filters on the Name filters page, on the Preferences panel.

Defining a name filter

Before you begin:
You can define name filters in the following ways:

- You specify a search phrase in the Resource name section on the Search panel (for details, see "Searching for resources" on page 157)
- You specify a search phrase on the Name filters page, on the Preferences panel. Note that search phrases that you define there also become available for the domain and its nodes in **Resource name** drop-down list on the Search panel.

Perform the following steps to define a name filter on the Name filters page:

1. Open the Name filters page (**Menu —> Preferences —> Name filters**).

2. Select the domain for which you want to define a new filter.

3. Click **New**. The Name filters panel is displayed.

4. Specify the search phrase to define the name filter. You have the following options:
   - To display only one specific resource, type the exact resource name.
   - Use the asterisk * as wildcard to display all resources whose names contain the search phrase. The wildcard can appear in any position and, if necessary, more than once (for example, *DB2*), and can stand for 0..n characters.
   - To display all the resources that contain at least one of several search phrases, type all phrases separated by a blank; the wildcard can be used in one or all phrases (for example, *DB2* SAP*).
   - For resource names that may contain blanks, type the complete search phrase including the blank and enclose the search 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**.

Results:
The search phrase you specified is saved as a name filter for the domain. Note that the filter is domain-specific. If you want to use the same search phrase for a different domain and its nodes, you must specifically define an identical filter for that domain.

Applying an existing name filter

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 **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.

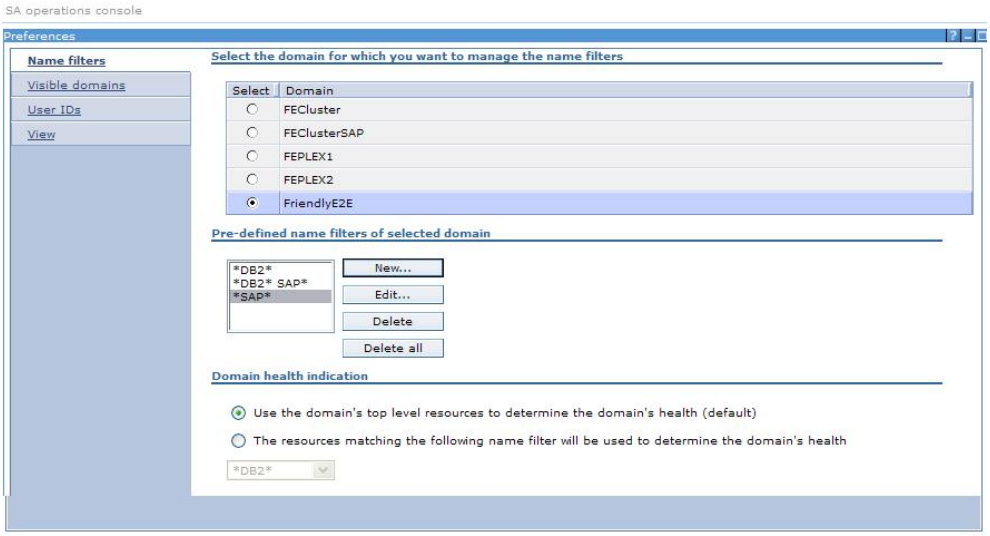

Figure 17. Name filters page on the Preferences panel

4. You use the buttons to perform the following tasks:

   New Opens the Name filters page on which you can specify a new name filter.

   Edit Opens the Name filters page on which you can edit the name filter you 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 are available 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 only lists 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.

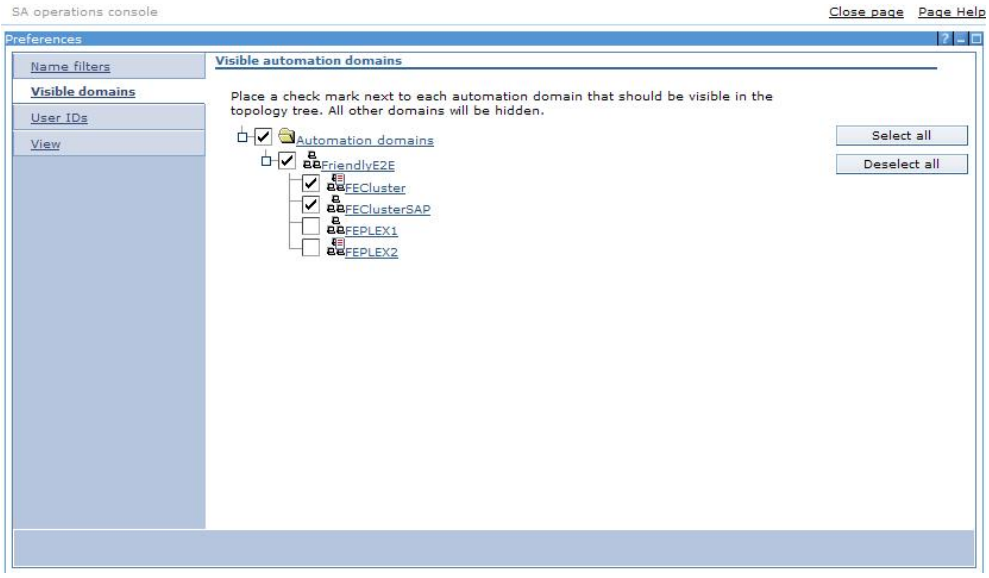

Figure 18. Visible automation domains page

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.

5. 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.

6. 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:

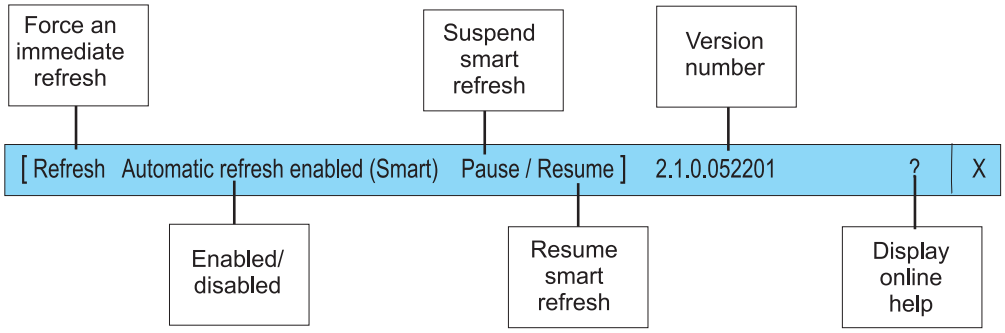

Note: A smart refresh only updates the information on the operations console that has changed since the last smart refresh. This usually guarantees that the information displayed on the console reflects the actual current state of all elements.

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 end-to-end automation manager configuration dialog. For more information about the configuration dialog, see the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*.

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 23. 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 following topics describe how you work with end-to-end automation policies on the operations console.

Note: You can also use the end-to-end automation manager command shell to activate or deactivate an end-to-end automation policy or to list the policies that are available in the policy pool. For information about the command shell, see Chapter 24, "Using the end-to-end automation manager command shell," on page 179. For information about the available command shell commands, see the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

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 where 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 167 and correct any errors.

To activate the policy, proceed as described in "Activating a policy" on page 167.

Working with requests

The tasks described in this topic are only available for resources that are hosted by request-driven automation domains.

Note: The topic describes how you perform the tasks on the operations console. You can also use the end-to-end automation manager command shell to work with end-to-end automation resources. For information about the command shell, see Chapter 24, "Using the end-to-end automation manager command shell," on page 179. For information about the available command shell commands, see the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

When an automation domain is request-driven, 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. On the Request Online panel, specify a comment in the entry field. The comment can later be viewed by displaying the request details.

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. On the Request Offline panel, specify a comment in the entry field. The comment can later be viewed by displaying the request details.

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 20. 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.

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:

- the requested action (Online, Offline, or Suspend)
- its source (for example, OPERATOR); if the request was submitted by an operator, the Source column also shows the user ID of the operator
- additional information about the request (in the Request info column). The information is generated by the automation manager that manages the resource
- 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.

Bringing resources online and offline

Perform this task to issue start or stop commands against resources that are hosted by command-driven automation domains, which do not maintain request lists for resources.

Before you begin:

Before issuing a start or stop command against a referenced first-level automation resource, you must suspend automation for the corresponding end-to-end automation resource reference, if the command will bring the referenced resource into a state that conflicts with the desired state of the resource reference.

If automation for the resource reference is not suspended in such a case, the end-to-end automation manager will issue a request against the referenced resource when it detects the state conflict, which will immediately bring the referenced resource into the desired state again that is defined for the resource reference. (see also "Suspending and resuming automation for resources" on page 174).

To bring a resource online or offline, perform the following steps:

1. Select the resource in the resource table.

---

2. On the general page, click Bring online or Bring offline.

   Note: The observed state of the resource determines which button is enabled. If the resource's observed state is Online, the **Bring offline** button is enabled, if its observed state is Offline, the **Bring online** button is enabled. If the resource's observed state is neither Online nor Offline, both buttons are enabled.

---

3. On the panel that appears, specify a comment. The comment is written to the log file for later reference.

---

4. Click **Submit** to submit the command.

Result: The resource is started or stopped.

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.

Note: This topic describes how you reset a resource from the operations console. You can also perform the task by using the command **resetres** in the end-to-end automation manager command shell. For information about the command shell, see Chapter 24, "Using the end-to-end automation manager command shell," on page 179. For information about the **resetres** command, see the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

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.

Suspending and resuming automation for resources

Suspending automation for a resource causes the automation manager not to react on observed state changes by issuing requests against the resource. The displayed observed state of a suspended resource only reflects the current situation. The compound state is still calculated in the usual way, by comparing the actual observed state of the resource to its desired state, but a state mismatch no longer triggers actions.

A state change of a suspended resource can still act as a trigger for state changes of other resources that have a relationship to the suspended resource. This includes that resources having relationships to the suspended resource may still be started or stopped by automation.

You can use the suspend function if you want to start or stop resources directly, without always having to interact with the automation manager for starting and stopping them. This may be required, for example, when you want to apply service to a resource. The service installation process (for example, an update installation program) may need to start and stop the application represented by the resource frequently, without knowing that the resource is automated and should only be started and stopped by automation.

In such a service scenario you should do this:

1. Suspend automation for the resource.
2. Apply service.
3. Resume automation for the resource.

Suspended resources show the following behavior:

An end-to-end automation resource group is suspended
: Automation is suspended for the group and all of its members.

An end-to-end automation choice group is suspended
: Automation is suspended for the group and all of its members. This means, for example, that the end-to-end automation manager will not stop

any alternative member whose observed state changes to online. Therefore, it is no longer ensured that only one member (the preferred member) is online at a time.

A suspended resource has relationships
A resource's relationships are still honored when automation is suspended:

- A suspended resource as the target of a forcedDownBy relationship can still cause the source resource to be stopped whenever the observed state of the suspended resource changes to offline.
- The observed state change of a suspended resource as target of a startAfter or stopAfter relationship still triggers the start or stop of the source resource.

States Suspending automation does not have an impact on the calculation of the operational and compound states of the resource, and a mismatch between the desired state and the observed state still causes the resource to go into a warning or error state, which is displayed on the operations console.

Operator requests can be submitted
Operator requests (Online, Offline, Cancel) are accepted although the resource is suspended. Depending on which action is performed, the requests are added to or removed from the request list and may trigger a change of the desired state. However, the automation manager will not take action to change the observed state should it conflict with the new desired state.

Suspended end-to-end automation resources can be reset
Suspended resources that are in operational state Unrecoverable Error or Reference Broken can be reset. A reset causes the observed state to change to Unknown, and the end-to-end automation manager will resubscribe for the referenced resource in order to retrieve the current observed state.

Steps for suspending automation for a resource

Perform the following steps:

1. Select the resource in the resource table.
2. On the general page, click **Suspend automation**.
3. On the panel that appears, specify a comment.
4. Click **Submit**.

Steps for resuming automation for a resource

Perform the following steps:

1. Select the resource in the resource table.
2. On the general page, click **Resume automation**.
3. On the panel that appears, specify a comment.
4. Click **Submit**.

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" on page 178).
- 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 178.

Note: This topic describes how you work with choice groups on the operations console. You can also change the preferred member of a choice group by using the command **chprefmbr** in the end-to-end automation manager command shell. For information about the command shell, see Chapter 24, "Using the end-to-end automation manager command shell," on page 179. For information about the **chprefmbr** command, see the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

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.

Chapter 24. Using the end-to-end automation manager command shell

You can use the end-to-end automation manager command shell to perform the following tasks by issuing commands to the end-to-end automation manager:

- List resources and resource groups and their states
- List resource group members
- List relationships
- Display, activate, and deactivate policies
- Change the preferred member of a choice group
- Issue online and offline requests against resources, and cancel requests
- Suspend and resume automation for resources
- Reset a resource from an unrecoverable error

The command shell can be used in two modes:

- **Line mode:** Allows you to issue a single command against the automation manager. When the command has been executed, the results are displayed on standard output and the command shell is closed. The output from a line mode command can be redirected to a file or to a tool that parses the results (for example, awk).
- **Shell mode:** Opens a subshell in interactive mode, allowing you to issue multiple commands against the automation manager successively. In shell mode, only one session is opened against the automation manager and you have to authenticate yourself only once. In shell mode, only automation manager commands are supported. In particular, it is not possible to redirect the output of a command to a file or to another command.

You cannot use the command shell to control the end-to-end automation engine, such as starting and stopping. For a description of the command-line interface of the automation engine, see Chapter 16, "Using the command-line interface of the automation engine," on page 103.

The following sections describe how to invoke and use the command shell in both modes. For a detailed description of the available commands, see the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

Using the command shell in shell mode

Before you begin:

- The end-to-end automation manager you want to connect to must be active (the WebSphere Application Server that the end-to-end automation manager uses (server1)). Otherwise, you will receive a message but the shell is not closed and you can issue a limited set of commands.

To access the command shell in shell mode, perform the following steps:

1. Log in to the server on which the end-to-end automation manager is running (using a Secure Shell, for example).
2. Issue the command **eezcs**.

3. Type your user credentials.

Results:

- If the command shell finds an active end-to-end automation domain:
  - The domain is selected as target for all commands you issue from the command shell.
  - A sub-shell opens and prompts you for input.

    Example:

    This is what you see in the command shell when an active end-to-end automation domain ("E2EDOM") was found:

    ```
    saxb05:/root # eezcs
    Connecting...
    Realm/Cell Name: null
    User Identity: iscadmin
    User Password:
    Using End-to-End Domain E2EDom
    EEZCS>_
    ```

  For a detailed description of the available commands, see the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.
- If no active end-to-end automation domain is found because no domain has joined or the domain is not online, a message is displayed but the connection is not closed and you can still issue the following commands at the command prompt:

  lseezdom
  Shows information about all domains that are currently known to the automation manager. The list of domains may contain first-level automation domains.

  help Displays the usage instructions for all shell commands or, when invoked with the command name as attribute, for a specific command.

  quit Closes the command shell.

Using the command shell in line mode

To issue a single command to an end-to-end automation manager, enter:

```
eezcs -c <command>
```

Results: When the command has been executed, the results are displayed, and the command shell is closed.

For a detailed description of the available commands, see the *IBM Tivoli System Automation for Multiplatforms End-to-End Automation Management Component Reference*.

Part 5. Working with the HACMP and MSCS adapters

Chapter 25. Working with the HACMP adapter and HACMP objects

The following sections describe how to work with the High Availability Cluster Multi-Processing (HACMP) adapter and HACMP objects.

Important notes:

1. The HACMP adapter can only be connected to an end-to-end automation management component V2R2 or later.
2. HACMP object names and their text fields, for example, group names, resource names, and descriptions, must not contain the following characters: " (double quotation mark), ' (single quotation mark), ; (semicolon), $ (dollar sign), / (slash)

Special considerations for the HACMP adapter

The following considerations apply to the system automation adapter for HACMP (HACMP adapter):

- HACMP clusters are not request- but command-driven. Commands for bringing resources and groups online or offline are performed but not retained as persistent goals. No list of previously issued commands is available, and commands previously issued against a group or resource cannot be canceled. The latest command issued against a group or resource determines whether it should be online or offline. Commands issued by operators have the same priority as commands issued by the end-to-end automation manager.
- HACMP resources and groups cannot be suspended from automation by an end-to-end automation operator.
- HACMP groups have no “real” desired state. HACMP performs online and offline commands on HACMP groups by propagating them to member resources. HACMP groups only act as containers and reflect the state of the contained HACMP resources. If some of the HACMP resources in a group were brought online and others offline, the group is in a mixed state - it is not clear whether the desired state of the group is online or offline.
- HACMP clusters do not have a policy concept as known by end-to-end automation. For this reason, the Policy Information page for HACMP domains does not show reasonable information.

Representation of HACMP objects and possible actions on the operations console

HACMP clusters
HACMP clusters are displayed as first-level domains on the operations console.

HACMP cluster nodes
HACMP cluster nodes are displayed on the operations console as nodes of their HACMP domain:

Figure 19. Two node HACMP cluster on the operations console

The nodes of an HACMP domain can be included in and excluded from automation:

- Excluding a node from automation: Stops the cluster services on the node.
- Including the node in automation: Starts the cluster services on the node.

HACMP resource groups and resources

HACMP resource groups are displayed as top-level resource groups. They can be brought online and offline from the operations console. Performing the actions on the operations console invokes the following command:

```
cIRGmove <resource_group>
```

HACMP resource groups are either move groups (if non-concurrent) or "collection" resource groups (if concurrent).

The following figure shows the single HACMP move group ("shop_rg") that is hosted by the domain "cl_hacmp".

Figure 20. HACMP top-level resource group

When you open the top-level resource group ("shop_rg"), you see that it comprises two resource groups. These resource groups are so-called "node instances" of the actual (top-level) resource group and are merely used as virtual containers for the constituents of the top-level resource group that can run on a specific node. As the HACMP sample domain depicted in the figures in this chapter consists of two nodes and the HACMP resource group can run on each of the nodes, the top-level resource group contains one virtual resource group for each node:

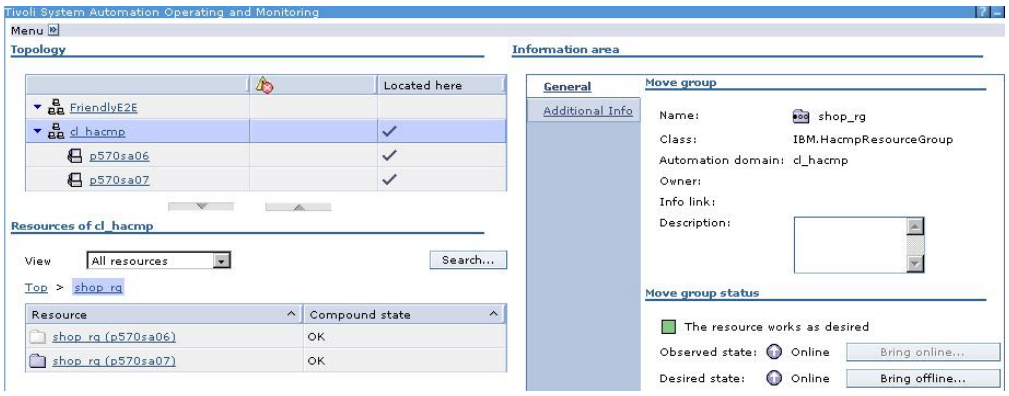

Figure 21. HACMP node instances of a resource group

As the top-level HACMP resource group is a so-called move group, which means that the group can only run on one node at a time, the node instance on node "p570sa07", which is currently running, appears in color, while the other node instance is grayed out.

When you open a node instance, the constituents of the top-level resource group that can run on the node are displayed. The sample node instance "shop_rg (p570sa07)" contains only a single member:

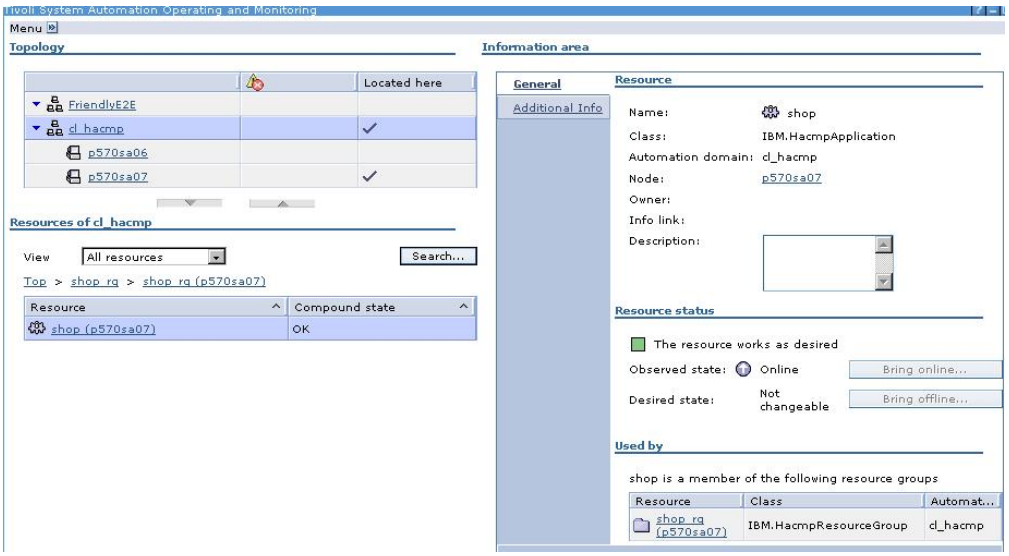

Figure 22. HACMP resource

Note that mountpoints, logical volumes, and volume groups that are automated by HACMP are not displayed in the resources section of the operations console.

HACMP relationships

On the operations console, only parent-child relationships between HACMP resource groups are reflected. The following HACMP resource group dependencies are not displayed on the operations console:

- "online on the same node" (collocation of resource groups)
- "online on different nodes" (anticollocation of resources groups)
- "online on same site" (site-collocation of resource groups)

Defining an end-to-end automation policy for HACMP resources

To include HACMP resources in an end-to-end automation policy, you create a resource reference for each of the HACMP resource groups that is to be managed by end-to-end automation management. You can use any of the end-to-end automation-specific relationships to specify dependencies between HACMP resource groups, or between resource groups that are managed by HACMP and resources that are managed by other first-level automation products.

When you define a resource reference for an HACMP resource group in an end-to-end automation policy, you must provide information about the HACMP resource group in the <ReferencedResource> subelement. You can easily obtain all the required information on the operations console by displaying the General page for the HACMP resource group (see "Gathering the required data for defining a policy" on page 86).

This is a sample end-to-end automation policy that references HACMP resources:

```
?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>E2E:shop->db2</PolicyName>
          <AutomationDomainName>FriendlyE2E</AutomationDomainName>
          <PolicyToken> 1.9.7</PolicyToken>
          <PolicyAuthor>Schawer</PolicyAuthor>
          <PolicyDescription>Demo policy shop(HACMP) depends-on db2(ITSAMP).</PolicyDescription>
      </PolicyInformation>

        <ResourceReference name="refha_shop">
                <!-- <DesiredState>Online</DesiredState> -->
                <Description>e2e ref to HACMP shop application.</Description>
                <Owner>Schawer ext:3704</Owner>
                <InfoLink>http://www.shop.com</InfoLink>
                <ReferencedResource>
                        <AutomationDomain>cl_hacmp/AutomationDomain>
                        <Name>shop_rg</Name>
                        <Class>IBM.HacmpResourceGroup</Class>
                </ReferencedResource>
        </ResourceReference>

        <ResourceReference name="refsa_db2">
                <!-- <DesiredState>Online</DesiredState> -->
                <Description>e2e ref to ITSAMP db2 application.</Description>
                <Owner>Schawer ext:3704</Owner>
                <InfoLink>http://w3.it-dep.com</InfoLink>
                <ReferencedResource>
                        <AutomationDomain>samp55078</AutomationDomain>
                        <Name>db2_rg</Name>
                        <Class>IBM.ResourceGroup</Class>
                </ReferencedResource>
        </ResourceReference>

        <Relationship>
                <Source>
                        <ResourceReference name="refha_shop"/>
                </Source>
                <Type>ForcedDownBy</Type>
                <Target>
                        <ResourceReference name="refsa_db2"/>
                </Target>
```

```
        </Relationship>
        <Relationship>
                <Source>
                        <ResourceReference name="refha_shop"/>
                </Source>
                <Type>StartAfter</Type>
                <Target>
                        <ResourceReference name="refsa_db2"/>
                </Target>
        </Relationship>
        <ResourceGroup name="E2E_shop_db2" >
                  <DesiredState> Online </DesiredState>
                  <Description>E2EGroup with DB2 and shop application</Description>
                  <Owner>schawer</Owner>
                  <InfoLink>http://www.shop.com</InfoLink>
                  <Members>
                     <ResourceReference name="refha_shop"/>
                     <ResourceReference name="refsa_db2"/>
                  </Members>
       </ResourceGroup>

</AutomationPolicy>
```

Starting, stopping, and querying the status of the HACMP adapter

Use the following commands in a command shell to start or stop the HACMP adapter and to query its state:

Table 21. HACMP adapter commands

Command	Description
hacadapter start	Starts the adapter
hacadapter stop	Stops the adapter
hacadapter status	Queries the adapter status

Chapter 26. Working with the MSCS adapter and Microsoft Server Clustering objects

The following sections describe how to work with Microsoft Server Clustering (MSCS) objects and the MSCS adapter.

Important notes:

1. The MSCS adapter can only be connected to an end-to-end automation management component V2R2 or later.
2. MSCS object names and their text fields, for example, group names, resource names, and descriptions, must not contain the following characters: " (double quotation mark), ' (single quotation mark), ; (semicolon)

Special considerations for the MSCS adapter

The following considerations apply to the MSCS adapter:

- MSCS clusters are not request- but command-driven. Commands for bringing resources and groups online or offline are performed but not retained as persistent goals. No list of previously issued commands is available, and commands previously issued against a group or resource cannot be canceled. The latest command issued against a group or resource determines whether it should be online or offline. Commands issued by operators have the same priority as commands issued by the end-to-end automation manager.
- MSCS resources and groups cannot be suspended from automation by an end-to-end automation operator.
- MSCS groups have no “real” desired state. MSCS performs online and offline commands on MSCS groups by propagating them to member resources. MSCS groups only act as containers and reflect the state of the contained MSCS resources. If some of the MSCS resources in a group were brought online and others offline, the group is in a mixed state - it is not clear whether the desired state of the group is online or offline.
- MSCS clusters do not have a policy concept as known by end-to-end automation. For this reason, the Policy Information page for MSCS domains does not show reasonable information.
- MSCS does not monitor resources which are not expected to be online.

  Example:

  A file share resource has two different cluster nodes as possible owners. If the file share is currently defined and working (that is, online) on the first node, MSCS does not monitor the state of the file share on the second cluster node. MSCS will not notice a manual definition of the file share on the second node.

  The MSCS adapter does *not* work around this monitoring approach and is thus not able to reliably report resources’ offline states.
- MSCS groups reject offline commands in the following cases:
  - The group contains the quorum resource.
  - The group contains the MSCS adapter service resource (if the adapter is made highly available).
- MSCS resources reject offline commands in the following cases:
  - The resource is the quorum resource and the quorum resource directly or indirectly depends on the resource to be taken offline.

  - The resource is the MSCS adapter service resource (if the adapter is made highly available).
  - If the MSCS adapter service resource directly or indirectly depends on the resource to be taken offline (if the adapter is made highly available).
- MSCS nodes reject exclude commands if the adapter is made highly available and the group that contains the MSCS adapter service resource is located on the node. In this case, message EEZZ0012E appears indicating that the group in question cannot be taken offline without impacting the MSCS adapter.

Representation of MSCS objects and possible actions on the operations console

MSCS clusters
MSCS clusters are displayed as first-level domains on the operations console.

Nodes MSCS cluster nodes are displayed on the operations console as nodes of their MSCS domain. They can be included in and excluded from automation:

- Excluding a node from automation: The MSCS node is suspended and all resources are moved away from the node.
- Including the node in automation: Resumes the MSCS node.

MSCS networks
MSCS networks are displayed as resource groups that contain MSCS network interfaces as group members. MSCS networks can only be monitored on the operations console.

MSCS network interfaces
MSCS networks interfaces are displayed as resources. An MSCS network interface is always a member of exactly one MSCS network. MSCS network interfaces can only be monitored on the operations console.

MSCS groups
MSCS groups are displayed as resource groups which contain MSCS resources as group members. MSCS groups can be brought online and taken offline. As MSCS is command-driven, no request lists are maintained by MSCS. MSCS propagates online and offline actions against a group to the member resources.

MSCS resources
MSCS resources are displayed as move groups which contain a set of member resources. One member resource ("fixed resource") is displayed for each MSCS node on which the MSCS resource is allowed to run. The move group representing an MSCS resource is always a member of exactly one MSCS group. Move groups representing MSCS resources can be brought online and taken offline. As MSCS is command-driven, no request lists are maintained by MSCS.

MSCS resource type objects
MSCS resource types are only displayed as additional information on the Additional Info page on the operations console.

MSCS relationships
The MSCS relationships hasMemberNetwork and hasMemberGroup are represented as group memberships. All other MSCS relationships are only displayed as additional or location information.

Defining an end-to-end automation policy for MSCS resources

All resources that are hosted by an MSCS cluster can be referenced in an end-to-end automation policy. However, online and offline commands are only supported for the following MSCS resources, which is why they are the recommended choice for referenced resources:

- Resource groups representing MSCS groups
- Move groups representing MSCS resources

For the following MSCS resources online and offline commands are not supported:

- Fixed resources representing MSCS resources
- Resource groups representing MSCS networks
- Resources representing MSCS network interfaces

When you define a resource reference in an end-to-end automation policy, you must provide information about the MSCS resource in the `<ReferencedResource>` subelement. You can easily obtain the required information on the operations console by displaying the General page for the MSCS resource (see "Gathering the required data for defining a policy" on page 86).

Referencing MSCS resources in an end-to-end automation policy

Use the following sections to learn what you must specify when you define resource references for MSCS resources in an end-to-end automation policy.

Referencing MSCS groups in an end-to-end automation policy

The following table shows what must be specified for an MSCS group in an end-to-end automation policy.

Table 22. Defining a resource reference for an MSCS group

ReferencedResource subelement	What to specify
AutomationDomain	Name of the MSCS domain
Name	Name of the group in the MSCS cluster. The name is displayed in the information area of the operations console.
Class	MSCS.Group
Node	The node element must be omitted

Example:

```
<ResourceReference name="Ref Calculator-rg">
      <Description>This is the reference to MSCS.Group </Description>
      <Owner>Bob Smith</Owner>
      <InfoLink>http://www.example.com/help/</InfoLink>
      <ReferencedResource>
          <AutomationDomain>saxbopt-kk</AutomationDomain>
          <Name>Calculator-rg</Name>
          <Class>MSCS.Group</Class>
      </ReferencedResource>
</ResourceReference>
```

Referencing move groups representing MSCS resources in an end-to-end automation policy

The following table shows what must be specified in an end-to-end automation policy for move groups representing MSCS resources.

Table 23. Defining a resource reference for a move group representing an MSCS resource

ReferencedResource subelement	What to specify
AutomationDomain	Name of the MSCS domain
Name	Name of the resource in the MSCS cluster. The name is displayed in the information area of the operations console.
Class	The MSCS resource type of the resource must be appended to the prefix `MSCS.MoveGroup`, for example, `MSCS.MoveGroup.Generic Service`
Node	The node element must be omitted

Example:

```
<ResourceReference name="Ref Calculator">
      <Description>This is the reference to MSCS.MoveGroup.Generic Application </Description>
      <Owner>Bob Smith</Owner>
      <InfoLink>http://www.example.com/help/</InfoLink>
      <ReferencedResource>
          <AutomationDomain>saxbopt-kk</AutomationDomain>
          <Name>Calculator</Name>
          <Class>MSCS.MoveGroup.Generic Application</Class>
      </ReferencedResource>
</ResourceReference>
```

Referencing fixed resources representing MSCS resources in an end-to-end automation policy

The following table shows what must be specified in an end-to-end automation policy for fixed resources representing MSCS resources.

Table 24. Defining a resource reference for a fixed resource representing an MSCS resource

ReferencedResource subelement	What to specify
AutomationDomain	Name of the MSCS domain
Name	Name of the resource in the MSCS cluster. The name is displayed in the information area of the operations console.
Class	The MSCS resource type of the resource must be appended to the prefix `MSCS.FixedResource`, for example, `MSCS.FixedResource.Generic Service`
Node	Name of the node to which the fixed resource is bound.

Referencing MSCS networks in an end-to-end automation policy

The following table shows what must be specified for an MSCS network in an end-to-end automation policy.

Table 25. Defining a resource reference for an MSCS network

ReferencedResource subelement	What to specify
AutomationDomain	Name of the MSCS domain
Name	Name of the network in the MSCS cluster. The name is displayed in the information area of the operations console.
Class	`MSCS.Network`
Node	Node element must be omitted.

Referencing MSCS network interfaces in an end-to-end automation policy

The following table shows what must be specified for an MSCS network interface in an end-to-end automation policy.

Table 26. Defining a resource reference for an MSCS network interface

ReferencedResource subelement	What to specify
AutomationDomain	Name of the MSCS domain
Name	Name of the network in the MSCS cluster. The name is displayed in the information area of the operations console.
Class	`MSCS.Network`
Node	Node element must be omitted.

Starting and stopping the MSCS adapter

How you start or stop an MSCS adapter depends on whether the adapter is highly available:

The adapter is made highly available using MSCS
: You start or stop the adapter by bringing the MSCS adapter group online or taking it offline in the Microsoft Cluster Administrator.

The adapter is not made highly available using MSCS
: You start or stop the adapter from the Services panel on the Microsoft Management Console.

Part 6. Appendixes

Appendix A. Policy definition worksheet

Use this worksheet to collect the information required for creating a resource reference for a first-level automation resource. The information you need about the first-level automation resource is available on the resource's General page (see “Gathering the required data for defining a policy” on page 86).

Table 27. 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	
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	

Table 27. Worksheet for defining an end-to-end automation policy (continued)

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	

Appendix B. 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
`<Tivoli_Common_Directory>/eez/logs`	message log files, trace files
`<Tivoli_Common_Directory>/eez/ffdc`	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 end-to-end automation manager configuration dialog. For information about the configuration dialog, refer to the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*. 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 2.2 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

Converting XML trace files to HTML format

The end-to-end automation engine and various adapters write traces and logs in an XML file format:

- The log files, which contain messages for administrators and operators, are automatically converted to HTML and can be viewed on the operations console by clicking the **View log** button for a domain.
- The trace files are only intended for use by IBM support. They are used, for example, to analyze the automation behavior or the startup or shutdown sequences of a component and may also contain additional information about exceptions that were generated by the automation engine or an automation adapter.

Trace files are hard to read because they are written in an XML dialect. However, you can easily convert them to HTML format to display them in a Web browser such as Mozilla or Microsoft Internet Explorer.

To convert the XML trace files to HTML, you use the log viewer tool that is shipped with the end-to-end automation management component. You find the log viewer tool in the following directory of the end-to-end automation management component archive:

```
<EEZ_INSTALL_ROOT>/install/logviewer214_basics.zip
```

You can unzip the file to any directory. For additional information about the tool, refer to the `readme.html` file, which becomes available in the directory to which you unzip the files.

To convert a trace file to HTML, perform the following steps:

1. Create a file named `stdtrace`.
2. Add the following single line to the file:

   ```
   select Time,SourceFile, SourceMethod,MessageId,LogText,Exception,Thread where (ProductId=SAMP)
   ```
3. Edit the file `viewer.bat` or `viewer.sh` and adjust the JAVA_PATH variable to point to the Java runtime environment shipped with WebSphere Application Server.
4. Use the `viewer.bat` or `viewer.sh` script to convert the trace or log file to HTML, for example:

   ```
   viewer -f stdtrace traceengine.xml > traceengine.html
   ```

Log files in a multilingual environment

In general, messages are generated according to the locale that best fits the language preference specified in the browser in which the operations console is displayed. Messages are presented on the operations console and written to one or multiple log files, depending on the SA for Multiplatforms subcomponent that generates the message.

If multiple browsers with different language preferences are used , the log files may contain messages in multiple languages. Additionally, some messages are written to the log files independent of any operator interaction. For example, when a SA for Multiplatforms subcomponent is started or stopped, it writes a message to its log file according to the locale in which it was started or stopped.

In case you need to understand the content of a message in the log file that is written in a language you do not know, refer to the message catalog provided in the *End-to-End Automation Management Component Reference* to find the message by message ID.

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 manager configuration dialog, verify that all parameters are set correctly.
2. Restart the automation engine.

For information about the configuration dialog, refer to the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*. For information on starting the automation engine, refer to Chapter 16, “Using the command-line interface of the automation engine,” on page 103.

A base component 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.

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 settings on the User credentials page of the configuration dialog.
For information about the configuration dialog, refer to the *IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide*. 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 210 describes how you change the environment variables on the administrative console of WebSphere Application Server.

Table 28. 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.

If a very large number of domains is to be monitored or the domain contains a very large number of resources and the value of **com.ibm.eez.aab.watchdog-interval-seconds** is not sufficiently large, the watchdog may not be able to contact all domains and receive their reply events within the given time. This results in incorrect communication state changes for the affected domains:

- In the WebSphere Application Server message log, pairs of messages EEZJ1003I can be found for each of these domains, indicating that the domain's communication state changed from ″OK″ to ″AsyncTimeout″ and back to ″OK″ within a short period of time.
- In addition, the operations console icons for the affected domains change accordingly for a short period of time from ″The domain is online″ to ″Resource events cannot be received″ and back to ″The domain is online″.

To resolve the problem, increase **com.ibm.eez.aab.watchdog-interval-seconds** to a value that is approximately double that of the number of domains. For example, if there are 200 domains, the value of **com.ibm.eez.aab.watchdog-interval-seconds** should be set to `400`.

If the number of resources to be monitored on the operations console is very large, increase the value of **com.ibm.eez.aab.watchdog-interval-seconds** in steps of 200 seconds until the result is satisfactory.

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 the *IBM Tivoli System Automation for Multiplatforms Base Component Administrator's and 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/db2tcpcm 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**:

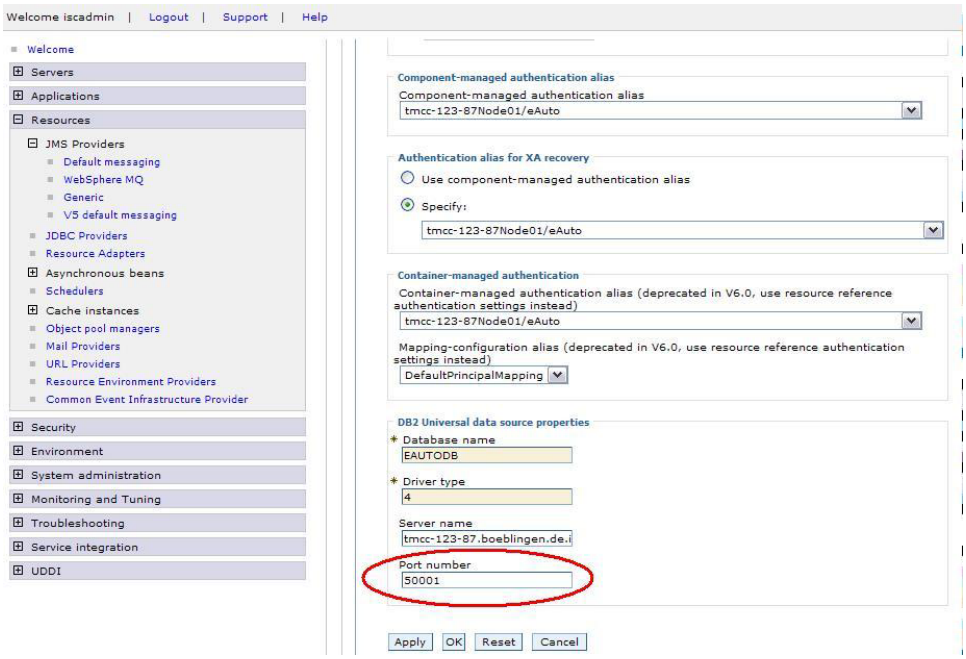

Critical exceptions in the WebSphere Application Server log file

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.

OutOfMemoryError in the WebSphere Application Server log file

An OutOfMemoryError may occur if a large amount of data is returned from a first-level automation domain. Depending on the situation, the error may become visible on the operations console or in the WebSphere Application Server message log file.

Perform the following steps to increase the JVM heap size:

1. Connect to the administrative console.
2. Go to **Servers —> Application Servers —> server1 —> Server Infrastructure —> Java and Process Management —> Process Definition —> Additional Properties —> Java Virtual Machine**
3. Increase the ″Maximum Heap Size″. The default value is 256 MB. If OutOfMemoryErrors occurred, it is recommended that you increase the value to 512 MB. Refer to the WebSphere Application Server online documentation for more information about how to determine the optimum value for the maximum heap size, depending on the available physical memory.
4. Save your changes. WebSphere Application Server must be restarted for the changes to take effect.

"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 EEUBus 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.

Troubleshooting command shell problems

AIX/Linux: Command shell hangs in shell mode - no input is possible

The command shell supports a command history function which can be exploited by using the scroll-up and scroll-down keys. On Windows this is a standard functionality provided by Java. On AIX/Linux this functionality is implemented by particular native input libraries. On some systems (depending on the distribution and version, and the shell used), this native code may lead to problems, for example:

- EEZCS fails with a javacore
- No input is possible (not even CTRL-C)

In order to circumvent this problem, you can disable the command history function by setting the HISTORY value to "false" in the file `<EEZ_INSTALL_ROOT>/bin/eezcs.sh`. This is the default setting in `eezcs.sh` which you need to change:

```
# Set HISTORY to false if you experience input problems
HISTORY=true
```

Troubleshooting automation engine problems

eezdmn command hangs during startup or shutdown

If the **eezdmn** command is hung during startup or shutdown of the automation engine, for example, because of an extreme load on the automation manager, you receive a timeout message after 60 seconds.

You can adjust the timeout value by adding the parameter EEZDMNCLIREADTIMEOUT to the script file `eezdmn.sh` (AIX/Linux) or to the batch file `eezdmn.bat` (Windows) and setting it to an appropriate value. The timeout value must be specified in milliseconds. For example, to receive a timeout message after 30 seconds, set the value of the parameter to 30000.

Troubleshooting HACMP adapter problems

Use this section for troubleshooting problems you experience when working with the HACMP adapter.

HACMP adapter log files

Increasing the trace logging level

If your trace is not detailed enough to analyze a problem and the problem can be recreated, it may be useful to increase the trace logging level:

1. Invoke the adapter configuration dialog using **cfghacadapter**.
2. On the main panel of the configuration dialog, click **Configure**.
3. Select the Logger tab.
4. Set the **Trace logging level** to **Maximum**.
5. Click **Apply**. The new setting takes effect immediately.

For more information about the HACMP adapter configuration dialog, see the *Installation and Configuration Guide*.

Log file locations

The HACMP adapter log files are located in the Tivoli Common Directory:

- Default location: `/var/ibm/tivoli/common`
- HACMP adapter-specific subdirectory structure in the Tivoli Common Directory:
  - `eez/ffdc` – Contains the First Failure Data Capture files (if the FFDC recording level is not set to Off in the adapter configuration dialog)
  - `eez/logs` – Contains the HACMP adapter trace file:
    - `traceFlatAdapter.log`

HACMP adapter does not start

Possible causes:

- HACMP level is lower than 5.3.0.5

  To check, use: `lslpp –l cluster.es.server.utils`
- Cluster services have not been started

  Start the services using smitty: **hacmp** —> C-SPOC —> Manage...

HACMP adapter terminates

Cluster services terminated while the HACMP adapter was running

If the adapter is automated, it should restart automatically on next priority node where cluster services run.

Adapter attempts to start but terminates again

This may indicate that the adapter has not been configured correctly. For information about configuring the adapter, see the *Tivoli System Automation for Multiplatforms Installation and Configuration Guide*.

HACMP adapter does not connect to the host

Make sure the firewall allows connections in both directions.

Check with **netstart**:

- whether the adapter listens on the request port (default port is `2001`)
- whether the end-to-end automation manager listens on the event port (default port is `2002`)

HACMP resource groups cannot be started or stopped

To bring HACMP resource groups online or offline, the HACMP Cluster-SubState must be STABLE. If the Cluster-SubState is UNSTABLE, which is typically the case during resource state transitions, Bring online and Bring offline actions against resource groups are not accepted. You can view the Cluster-Substate on the Additional Info page for the HACMP cluster. The information on the page is not updated automatically. To see if the Cluster-SubState has changed, use **Menu** —> **Refresh all** from the Menu bar of the operations console. When the Cluster-SubState has changed to STABLE, Bring online and Bring offline actions against resource groups can again be perform.

The following figure shows the Additional Info page for the HACMP cluster "cl_hacmp":

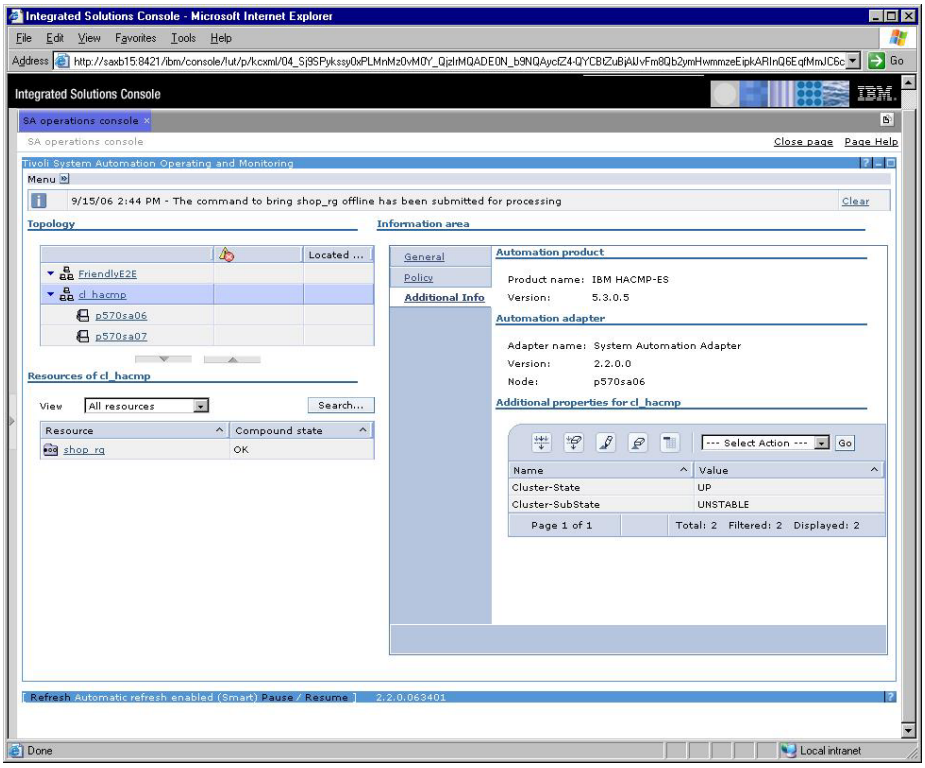

Figure 23. Additional Info page for an HACMP cluster

Troubleshooting MSCS adapter problems

Use this section for troubleshooting problems you experience when working with the MSCS adapter.

MSCS adapter log files

This is where the adapter log files are located:

- Tivoli Common Directory

  Default location: `C:\Program Files\IBM\tivoli\common`

  MSCS adapter-specific subdirectory structure in Tivoli Common Directory:

  – `eez\ffdc` – Contains the First Failure Data Capture files (if the FFDC recording level is not set to Off in the adapter configuration dialog)

  – `eez\logs` – Contains the MSCS adapter log files:

    - `msgMSCSAdapter.log`
    - `traceMSCSAdapter.log` (if trace logging level is not set to Off)
    - `eventMSCSAdapter.log` (if trace logging level is not set to Off)

- The default adapter installation directory is `C:\Program Files\IBM\tsamp\eez\mscs`.

  Subdirectories and files used for troubleshooting:

  – The file `data\eez.release.information.txt` is created in the adapter installation directory when the MSCS adapter is started. It contains information about service applied to the MSCS adapter and about the configuration settings used.

  – The installation log files are located in the subdirectory `_inst_logs`.

Adapter configuration dialog problems occur

A problem occurs using the adapter configuration dialog

Problem determination:

- The file `cfgmscsadapter.bat` contains a command for launching the configuration dialog
- The file contains a duplicate of this command which enables diagnostic output (option **-DEBUG**)

The Apply button on the Logger page cannot be clicked

Possible cause: The MSCS adapter is not running.

Configuration files cannot be replicated

Possible causes:

- The MSCS cluster is not available.
- The cluster contains only a single node.

Replication fails with the message "Login on target node failed"

Possible cause: The domain user ID was not specified in the correct format, which is `<user_ID>@<domain_name>`.

MSCS adapter does not start

MSCS adapter does not start

Problem determination:

- The application event log should contain the message “The service SA MP MSCS Adapter has been started.”
- In the configuration file `cfg\mscs.service.properties`, uncomment the property `service-log-file`, restart the service, and investigate the resulting file.

  Ensure to comment the property again before returning to normal operation.

The SA MP Adapter Service reports the status Started for some seconds and stops again

- Startup should be completed within 60 seconds.
- Refresh the view to see the actual status.

Problem determination:

- Investigate the MSCS adapter log file `msgMSCSAdapter.log`.
- If no error messages can be found, increase the trace logging level to Maximum and provide all logs to IBM support.

The file msgMSCSAdapter.log contains the message EEZA0061E indicating that the adapter failed to bind to a socket

Possible reason if the MSCS adapter service is made highly available using MSCS:

- The network name or virtual IP address used for the “Automation adapter host” is not available during adapter startup

Possible solution:

- Check the spelling of the network name or virtual IP address in the adapter configuration dialog.
- Check that there are appropriate “Network Name” / “IP Address” resources defined in MSCS and that they are working properly.

- Check that the MSCS adapter service resource has a dependency on the "Network Name" / "IP Address" resources in MSCS.

MSCS adapter terminates

The MSCS adapter services stops and the log files contain no related error messages. In particular, message "EEZA0104I" does not appear in the MSCS adapter log file `msgMSCSAdapter.log`. The message indicates that the MSCS adapter was successfully stopped.

Problem determination:

1. Search for `javacore.*.txt` files in the subdirectory `lib`.
2. Use Windows tool `drwtsn32` to configure dump capturing. Use the following settings:

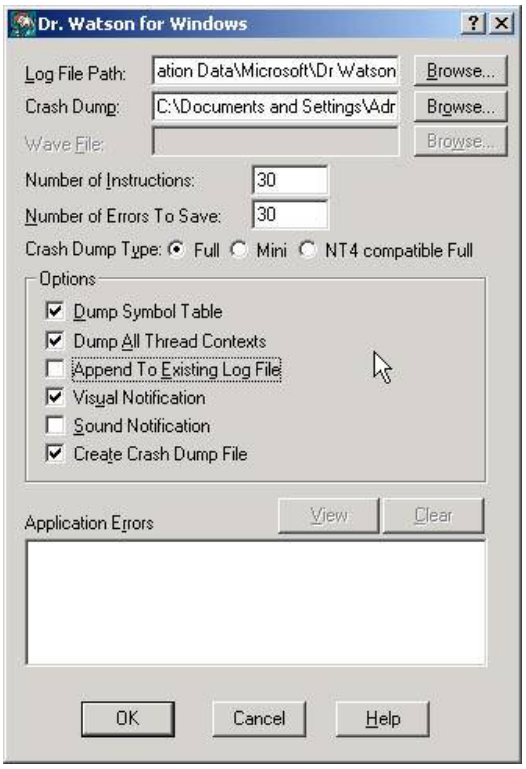

3. Try to recreate the MSCS adapter termination.
4. Provide the data to IBM support.

MSCS domain does not join

The MSCS domain does not join within two minutes and the MSCS adapter service is no longer running

Problem determination:

- Investigate the MSCS adapter log file `msgMSCSAdapter.log`.
- If no problems can be found, increase the trace logging level to "Maximum" and provide all logs to IBM support.

The MSCS domain does not join within two minutes but the MSCS adapter service is running

Problem determination and possible causes:

- An invalid host name or IP address is specified for the end-to-end automation management server.

- The end-to-end automation management server cannot be reached from the system running the MSCS adapter. To check, use ping, telnet, and tracert commands.
- Determine the network name / IP address the MSCS adapter sends to the end-to-end automation management server:
  - Increase the trace logging level at least to “Minimum”, restart the MSCS adapter, investigate the log file `eventMSCSAdapter.log`.
  - Locate the latest adapter join event (“EVT_RSN=domainAdapterJoin”). The event contains the required information.
- The system running the MSCS adapter cannot be reached from the end-to-end server. To check, use ping, telnet, and tracert commands.

Appendix C. 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, HACMP, 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.

Index

J

L

M

N

O

P

R

S

T

U

V

W

X

Readers' Comments — We'd Like to Hear from You

System Automation for Multiplatforms
End-to-End Automation Management Component
Administrator's and User's Guide
Version 2.2

Publication No. SC33-8275-00

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

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. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:
- Send your comments to the address on the reverse side of this form.
- Send a fax to the following number: FAX (Other Countries): (+49)+7031-16-3456
- Send your comments via e-mail to: eservdoc@de.ibm.com

If you would like a response from IBM, please fill in the following information:

Name	Address
Company or Organization	
Phone No.	E-mail address

Readers' Comments — We'd Like to Hear from You
SC33-8275-00

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

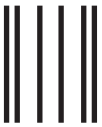

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

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

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

Program Number: 5724-M00

Printed in USA

SC33-8275-00