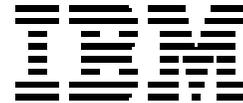


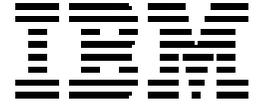
IBM Tivoli Intelligent Orchestrator and Tivoli Provisioning Manager



Operator's Guide

Version 1.1.1

IBM Tivoli Intelligent Orchestrator and Tivoli Provisioning Manager



Operator's Guide

Version 1.1.1

Note

Before using this information and the product it supports, read the information under "Notices".

Second Edition (December 2003)

© Copyright International Business Machines Corporation 2003. All rights reserved.

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

Contents

Figures	vii	Configuring firewalls	34
Tables	ix	Configuring virtual IPs	34
Preface	xi	Configuring servers	35
Who should read this guide	xi	Configuring boot servers	47
What this guide contains	xi	Configuring blade servers	50
Publications	xii	Configuring terminal servers	53
Contacting software support	xiii	Configuring ACLs	57
Conventions used in this book	xiii	Configuring power units	58
Operating system-dependent variables and paths	xiv	Configuring license pools	62
Chapter 1. Introduction to the IBM Tivoli Intelligent Orchestrator	1	Configuring software products	65
What is the Tivoli Intelligent Orchestrator?	1	Configuring software patches	67
What is the Provisioning Manager?	2	Configure software stacks	70
The main components of the Tivoli Intelligent Orchestrator	3	Configuring resource pools	75
The user interfaces	4	Adding a new resource pool	75
The Web-based interface	5	Editing the resource pool properties	75
The command-line interface	8	Editing the user-defined variables for the resource pool	76
The operating modes	9	Assigning workflows to a resource pool	76
Application Topologies	10	Assigning a software stack to the resource pool	76
Device driver support for workflow integration	10	Removing a resource pool	76
Interactive user interface	12	Chapter 3. Managing the Tivoli Intelligent Orchestrator	77
Software stack	13	Changing the operating mode of an application	77
SOAP initiation of device functions	13	Using the Web-based interface	77
Chapter 2. Configuring the Tivoli Intelligent Orchestrator	15	Using the command-line interface	78
Signing on to and signing off from the system	16	Changing the operating mode of a cluster	78
Signing on to the system	16	Using the Web-based interface	79
Signing off from the system	16	Using the command-line interface	79
Configuring users	16	Managing customer accounts	80
Adding a new user	16	Setting up a new customer account	80
Editing the user profile	18	Editing the properties of the customer account	80
Removing the user	19	Removing the customer account	81
Changing the global operating mode	21	Managing applications and clusters	81
Using the Web-based interface	21	Setting up a new application	81
Using the command-line interface	22	Editing the properties of the application	82
Configuring inventory resources	23	Adding a new cluster to the application	82
Configuring subnets	23	Setting up the routing properties for the cluster	83
Configuring VLANs	25	Assigning workflows to the application cluster	83
Configuring switches	26	Assigning a software stack to the application cluster	83
Configuring switch fabrics	29	Adding a new dedicated server to the application cluster	84
Configuring load balancers	31	Bringing an application online	84
Configuring routers	33	Manually moving the application to maintenance	85
		Removing an application from a customer account	85
		Removing a cluster from an application	85

Managing servers	85
Manually adding overflow servers to an application cluster	86
Manually freeing overflow servers	86
Removing a dedicated server from an application cluster	87
Moving available servers between resource pools ..	87
Initializing a server within a resource pool	87
Cleaning up a server within a resource pool	87
Performing management tasks on a server	88
Accessing the software stack for the server	90
Performing deployment tasks from the interactive views	90
Creating a VLAN	92
Removing the VLAN from the switch fabric	93
Turning a port on	93
Turning a port off	93
Moving a port to a VLAN	93
Creating a router connection	94
Removing a router connection	94
Rebooting a device	94
Initializing a device	94
Creating a virtual IP	95
Removing a virtual IP	95
Adding a server	95
Removing a server	95
Monitoring the system	96
Monitoring customers and resource pools	96
Monitoring overall server performance for an application or resource pool	97
Displaying cluster properties for an application ..	97
Monitoring the server performance	98
Displaying deployment requests	98

Chapter 4. Configuring and using workflows..... 103

About the Deployment Engine's workflows	103
Configuring and using workflows	104
Displaying workflow data	105
Displaying workflow references	106
Adding a new workflow	106
Editing the workflow properties	107
Setting up the workflow transitions	107
Setting up the workflow variables	109
Executing a workflow	111
Cloning a workflow	111
Exporting a workflow	112
Importing a workflow	113
Displaying the execution history for all workflows	113
Displaying the execution history for an individual workflow	113
Displaying execution details for a workflow	114
Displaying execution details for an embedded	

workflow	114
Deleting the execution history of a workflow	114
Configuring device drivers	115
Adding a new device driver	116
Editing the properties of the device driver	116
Assigning workflows to the device drivers	116
Removing the device driver	116
Assigning workflows	116
Assigning workflows to an application cluster ...	117
Assigning workflows to a resource pool	118
Assigning workflows to a data center device	119
Assigning workflows to a device driver	119
Managing the Deployment Engine	120

Chapter 5. Configuring Java plug-ins, simple commands, and logical device operations..... 121

Configuring Java plug-ins	121
Displaying Java plug-in data in simple view and detail view modes	122
Removing the Java plug-in	122
Importing the Java plug-in	123
Configuring simple commands	124
Displaying simple command data in simple view and detail view modes	124
Adding a new simple command	125
Editing the simple command properties	126
Setting up the command's variables	126
Removing the simple command	127
Cloning the simple command	128
Executing the simple command	128
Exporting the simple command	129
Importing the simple command	130
Configuring logical device operations	131
Displaying logical device operations in simple view and detail view modes	131
Adding a new logical device operation	132
Editing the properties of the logical device operation	133
Setting up variables for the logical device operation	134
Executing the logical device operation	135
Removing the logical device operation	135
Exporting the logical device operation	136
Importing the logical device operation	137

Chapter 6. Generating reports..... 139

Displaying reports	139
Report layout	140
List of available reports	141
Workflows	142
Customer Applications	145
Inventory	148

Usage	154	Appendix A. SOAP command syntax	
Chapter 7. SOAP commands available through command line	161	reference	175
SOAP scripts available through the command-line interface	161	OperationsModeService.wsdl	175
SOAP commands available through the Web-based interface	163	getApplicationMode	176
FaultManagementService.wsdl	163	getClusterMode	177
OperationsModeService.wsdl	164	getEffectiveApplicationMode	177
RecommendationsService.wsdl	165	getEffectiveClusterMode	178
SOAP commands initiating device functions	165	getGlobalMode	179
BootServerComponentService.wsdl	166	setApplicationMaintenanceMode	179
ClusterComponentService.wsdl	167	setApplicationMode	180
DeviceComponentService.wsdl	167	setClusterMode	180
FileRepositoryComponentService.wsdl	168	setGlobalMode	181
FirewallComponentService.wsdl	168	setMaintenanceMode	181
IpSystemComponentService.wsdl	169	Appendix B. Defining the report design in an XML file: a case study	183
LoadBalancerComponentService	169	Defining the report design in the XML file	183
OperatingSystemComponentService.wsdl	170	Displaying the generated report	190
PowerUnitComponentService.wsdl	171	Notices	191
RouterComponentService.wsdl	171	Trademarks	192
SoftwareComponentService.wsdl	171	Glossary	195
SoftwareStackComponentService.wsdl	172	Tivoli Intelligent Orchestrator-specific terms	195
SparePoolComponentService.wsdl	172	General terms	196
SwitchComponentService.wsdl	173	Index	203
SwitchFabricComponentService.wsdl	173		

Figures

1. The high-level architecture of the Tivoli Intelligent Orchestrator	4
2. An overview of the Tivoli Intelligent Orchestrator's screen and its features	7
3. A 3D graphical overview of the data center	8
4. The command-line interface of the Tivoli Intelligent Orchestrator	10
5. The VLAN interactive view	27
6. The switch interactive view	30
7. The switch fabric interactive view	35
8. The router interactive view	43
9. The firewall interactive view	44
10. The NIC interactive view for a server	50
11. The Tivoli Intelligent Orchestrator monitoring performance parameters	130
12. Overall server performance for a customer application	131
13. Overall server performance for a resource pool	132
14. Individual server performance	133
15. Deployment requests for an application	134
16. The generic structure of a nested workflow	138
17. A simple command seen as a wrapper for a Java plug-in	169
18. The Tivoli Intelligent Orchestrator displaying a report	187
19. The report layout	188
20. The Inventory report generated based on the design defined in the XML file	238

Tables

1. User credentials for accessing the Tivoli Intelligent Orchestrator	21
2. SOAP scripts available through the command-line interface	210
3. WSDL locations available through the Web-based interface	211
4. SOAP commands for FaultManagementService.wsdl	211
5. SOAP commands for OperationsModeService.wsdl	212
6. SOAP commands for RecommendationsService.wsdl	213
7. WSDL locations for initiating device functions	213
8. SOAP commands for BootServerComponentService.wsdl	214
9. SOAP commands for ClusterComponentService.wsdl	215
10. SOAP commands for DeviceComponentService.wsdl	215
11. SOAP commands for FileRepositoryComponentService.wsdl	216
12. SOAP commands for FirewallComponentService.wsdl	216
13. SOAP commands for IpSystemComponentService.wsdl	217
14. SOAP commands for LoadBalancerComponentService.wsdl	217
15. SOAP commands for OperatingSystemComponentService.wsdl	218
16. SOAP commands for PowerUnitComponentService.wsdl	219
17. SOAP commands for RouterComponentService.wsdl	219
18. SOAP commands for SoftwareComponentService.wsdl	219
19. SOAP commands for SoftwareStackComponentService.wsdl	220
20. SOAP commands for SparePoolComponentService.wsdl	220
21. SOAP commands for SwitchComponentService.wsdl	221
22. SOAP commands for SwitchFabricComponentService.wsdl	221

Preface

This guide describes the IBM Tivoli Intelligent Orchestrator.

It contains a general overview of the architecture and functionality of the Tivoli Intelligent Orchestrator and describes the procedures to configure and operate the system.

Who should read this guide

This guide is intended for corporate and Internet data center system administrators who will operate and maintain the Tivoli Intelligent Orchestrator. You should read this guide if you perform one or more of the following roles:

- Application owners, who will manage specific configurations of the system
- Data center operators, who will monitor, deploy, configure applications and infrastructure, and manage resources using the Tivoli Intelligent Orchestrator

It is assumed that you have an understanding of Internet browser navigation and computer operating systems. Prior knowledge of n-tier applications and data center infrastructure would also be very helpful.

A certain level of familiarity with a number of specific terms is also required. You may find it helpful to review the “Glossary” before reading other chapters in this guide.

What this guide contains

This guide contains the following chapters:

- Chapter 1, “Introduction to the IBM Tivoli Intelligent Orchestrator”, provides an overview of the system’s basic functionality, describes the available user interfaces, the operating modes, and the concepts introduced by *Application Topologies v1.0*.
- Chapter 2, “Configuring Tivoli Intelligent Orchestrator”, describes the procedures performed by a data center system administrator to configure users, servers, networking infrastructure devices, license pools, software products, stacks, and patches, and resource pools. Also described is the procedure to modify the Tivoli Intelligent Orchestrator’s global operating mode.
- Chapter 3, “Managing Tivoli Intelligent Orchestrator”, describes the procedures performed by a data center system administrator to manage customer accounts, applications, clusters, and servers, and to monitor the Tivoli Intelligent Orchestrator. Also described are the procedures to toggle between operating modes at application and cluster levels.
- Chapter 4, “Configuring and using workflows”, offers an introduction to the system’s workflows, and provides the procedures to configure and use workflows. Also provided are the procedures to configure device drivers and manage the Deployment Engine through the Web-based user interface.
- Chapter 5, “Configuring Java plug-ins, simple commands, and logical device operations”, describes the procedures to configure the Java plug-ins, simple commands, and logical device operations that are required for customizing and successfully executing various workflows.

- Chapter 6, “Generating reports”, describes the reports that the Tivoli Intelligent Orchestrator is able to generate and convert into PDF files or send to the printer.
- Chapter 7, “SOAP commands available through command line”, provides a reference to the SOAP scripts and commands that are available using the Web-based interface and the command-line interface.
- Appendix A, “SOAP command syntax reference”, provides a reference to the command syntax and examples for some of the SOAP commands that can be used through the Web-based user interface.
- Appendix B, “Defining the report design in an XML file: a case study”, provides a practical example of a report being generated based on the design defined in an XML file.
- The “Glossary” provides a list of some of the most commonly used terms in this guide.

Publications

This section lists publications in the Tivoli Intelligent Orchestrator library and related documents. It also describes how to access Tivoli publications online.

Tivoli Intelligent Orchestrator library

The following manuals are available in the Tivoli Intelligent Orchestrator library for Version 1.1.0:

- This manual, the *IBM Tivoli Intelligent Orchestrator: Operator’s Guide*, contains a general overview of the architecture and functionality of the Tivoli Intelligent Orchestrator and describes the procedures to configure and operate the system.
- The *IBM Tivoli Intelligent Orchestrator: Installation Guide* provides all the necessary information to install the Tivoli Intelligent Orchestrator and its third-party components on Windows® 2000.
- The *IBM Tivoli Intelligent Orchestrator: Overview Guide* provides an architectural overview of the Tivoli Intelligent Orchestrator, describes the main components of the system and their interaction, and details the concepts and the mathematical background that the system functionality is based on.

Accessing publications online

The *Release Notes* and publication for the IBM Tivoli Tivoli Intelligent Orchestrator product are available at the Tivoli Software Information Center Web site:

<http://www.ibm.com/software/tivoli/library/>

Click the Product manuals link to access the Tivoli Software Information Center.

Click the IBM Tivoli Tivoli Intelligent Orchestrator link to access the product library.

Note: If you print PDF documents on other than letter-sized paper, select the **Fit to page** check box in the Adobe Acrobat Print dialog. This option is available when you click **File > Print**. **Fit to page** ensures that the full dimensions of a letter-sized page print on the paper that you are using.

Contacting software support

If you have a problem with any Tivoli product, refer to the following IBM Software Support Web site:

<http://www-3.ibm.com/software/sysmgmt/products/support/>

If you want to contact software support, see the IBM Software Support Guide at the following Web site:

<http://techsupport.services.ibm.com\guides\handbook.html>

The guide provides information about how to contact IBM Software Support, depending on the severity of your problem, and the following information:

- Registration and eligibility
- Telephone numbers and e-mail addresses, depending on the country in which you are located
- Information you must have before contacting IBM Software Support

Conventions used in this book

This guide uses the following typeface conventions:

Bold

- Lowercase commands and mixed case commands that are otherwise difficult to distinguish from surrounding text
- Interface controls (check boxes, push buttons, radio buttons, spin buttons, fields, folders, icons, list boxes, items inside list boxes, multicolumn lists, containers, menu choices, menu names, tabs, property sheets), labels (such as Tip:, and Operating system considerations:)
- Column headings in a table
- Keywords and parameters in text

Italic

- Citations (titles of books, diskettes, and CDs)
- Words defined in text
- Emphasis of words (words as words)
- Letters as letters
- New terms in text
- Variables and values you must provide

Monospace

- Examples and code examples
- File names, programming keywords, and other elements that are difficult to distinguish from surrounding text
- Message text and prompts addressed to the user
- Text that the user must type
- Values for arguments or command options

<text>

Indicates a variable in a path name. For example, in the path `<INSTALL-DIR>\t1m`, `INSTALL-DIR` depends on the location where you have installed the component, while `\t1m` is constant.

Operating system-dependent variables and paths

This guide uses the Windows convention for specifying environment variables and for directory notation.

When using the UNIX command line, replace *%variable%* with *\$variable* for environment variables and replace each backslash (\) with a forward slash (/) in directory paths.

Note: If you are using the bash shell on a Windows[®] system, you can use the UNIX[®] conventions.

Chapter 1. Introduction to the IBM Tivoli Intelligent Orchestrator

This chapter provides the information you need to familiarize yourself with the look and basic functionality of the Tivoli Intelligent Orchestrator. The following information is included:

- Orchestrated Provisioning Overview
- Main components
- User Interface Options
- The Operating Modes
- Application Topologies

Orchestrated Provisioning Overview

Tivoli Intelligent Orchestrator is an automated resource management solution for corporate and Internet data centers. Through orchestrated provisioning it provides the ability to manipulate the IT environment in real time, according to defined business policies, to achieve the desired business goals.

The Tivoli Intelligent Orchestrator proactively configures resources among applications in a multi-application environment to balance end-user traffic demands, excess capacity, and the contractual obligations of service level agreements. Using an adaptive control technology, the system accurately predicts capacity fluctuations, and facilitates dynamic infrastructure reallocation.

The Tivoli Intelligent Orchestrator can:

- Gather information about the performance of all your application clusters and build a workload model that can predict future resource requirements
- Manage resources across all your application clusters to optimize business-aligned service level delivery
- Automate the deployment of the optimal computing resources to each application environment

The Tivoli Intelligent Orchestrator enables you to create, customize, and store a large variety of workflows that automate all data center processes, and makes it possible to build a powerful library of processes that can be assembled to meet any data center process requirement. The Tivoli Intelligent Orchestrator's workflows automate various data center processes, from configuring and allocating servers, to installing, configuring, and patching software, and can be either large and complex or as simple as a single command.

The Tivoli Intelligent Orchestrator includes *Application Topologies*. Three key concepts are at the core of this foundation: the *device driver*, the *switch fabric*, and the *software stack*.

Note: For further information on *Application Topologies* and the key concepts it introduces, refer to the “Application Topologies” section, later in this chapter.

Application Topologies requires an abstraction of the logical structure of an application, as well as the physical application deployment. The same application is deployed into different environments, such as testing, staging, or production environments. As a first step, the necessary infrastructure has been created.

Application software, as well as installing, configuring, patching software are very complex. By modeling all the software dependencies and identifying software device drivers, the relationship between workflows and software products is very much simplified. Through *Application Topologies*, the Deployment Engine workflows become device drivers as part of the Tivoli Intelligent Orchestrator's support for various hardware and software products. The system extends the existing Data Center Model and Deployment Engine components, enabling a new way of viewing workflows.

Note: For additional workflow information, refer to the “Configuring and using workflows” chapter, later in this guide.

What is the Provisioning Manager?

Within the IBM Tivoli Intelligent Orchestrator, the IBM Tivoli Provisioning Manager is a stand-alone product that can be purchased separately, based on your data center business needs. The Provisioning Manager automates the manual provisioning process of infrastructure deployment. It captures and automates the execution of your best practices for managing data center resources by building a framework to compile and then repeatedly execute these best practices consistently and efficiently. The Provisioning Manager comes with predefined best practices for standard products from all major infrastructure vendors.

If you have only purchased the Provisioning Manager, refer to the following chapters:

- Introduction to the IBM Tivoli Intelligent Orchestrator
- Configuring and using workflows
- Configuring Java plug-ins, simple commands, and logical device operations

Main components

The main components of the Tivoli Intelligent Orchestrator are:

Deployment Engine

This component is responsible for the creation, the storage, and the execution of repeatable workflows that automate the server configuration and allocation in the system. A workflow can represent either an entire reconfiguration process affecting multiple servers, or a single step in a larger reconfiguration process.

Data Center Model

This component includes a representation of all of the *physical* and *logical assets* under the Tivoli Intelligent Orchestrator's management, such as servers, switches, load balancers, application software, VLANs, security policies, service level agreements, and so on. It keeps track of the data center hardware and associated allocations to customer sites.

Data Acquisition Engine

This component is responsible for acquiring and pre-processing performance data from each managed application environment. Data is captured from the application, operating system, and infrastructure layers. This component uses a subscribing mechanism to distribute signals to other components of the Tivoli Intelligent Orchestrator, and performs filtering of raw signals.

Application Controller

An instance of the Application Controller is created for each application environment under management. Based on the application's workload model and predictions, as well as on real-time performance data, this component determines the resource requirements of the application.

Global Resource Manager

This component receives requirements for servers or network devices from all the application controllers, and manages the overall optimization. It has two primary responsibilities: it makes optimal resource allocation decisions and it ensures a stable control over the application infrastructure. The Global Resource Manager considers the different server requirements for each application environment and then determines where the servers are to be allocated.

Management Interface

This component provides an overview of the state of all *physical* and *logical assets* in the data center infrastructure, offering information about the servers and their allocation, and generating configurations and allocations. It can also be used to create application environments. It includes two user interfaces: a *Web-based interface* and a *command-line interface*.

Note: For more information on the specific terms that are used in this guide, refer to the "Glossary" chapter.

The following figure illustrates the general architecture of the Tivoli Intelligent Orchestrator.

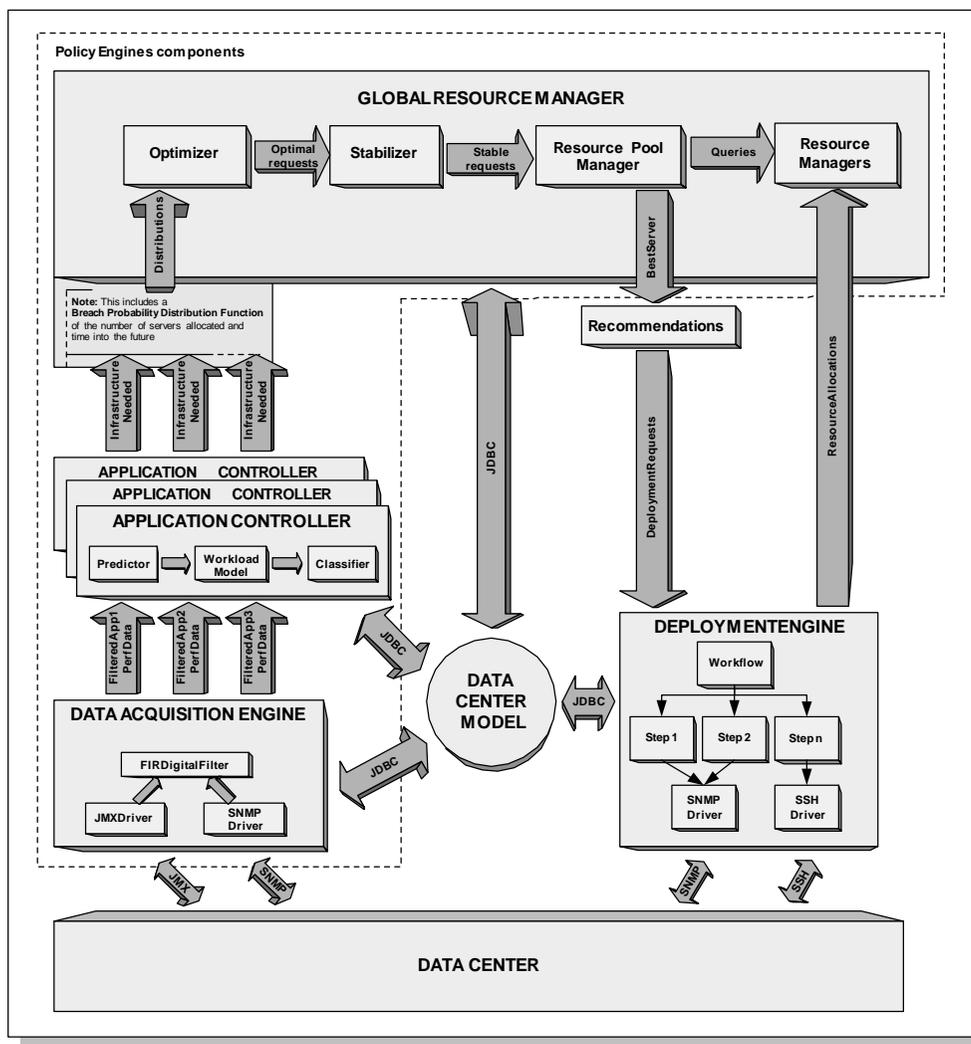


Figure 1 The high-level architecture of the Tivoli Intelligent Orchestrator

User Interface Options

The management interface of the Tivoli Intelligent Orchestrator provides two user interface options:

- The *Web-based interface*, which is designed to be used by application owners and data center operators, offers an intuitive way to display information about applications and components
- The *command-line interface* is designed to be used by data center operators who prefer to access the system properties and operations using the command line

Note: For more information on how to configure and operate the Tivoli Intelligent Orchestrator using these interfaces, refer to the “Configuring Tivoli Intelligent Orchestrator” and “Managing Tivoli Intelligent Orchestrator” chapters described later in this guide.

The following sections provide details on how to use each of these interfaces.

The Web-based interface

The Web-based user interface offers real-time access to Tivoli Intelligent Orchestrator and displays information about the deployed application and components. It allows you to monitor and control the resource management and application performance.

The following image provides an overview of the Web-based interface and its features.



Figure 2 An overview of the Tivoli Intelligent Orchestrator's screen and its features

Header

The header displays the Tivoli Intelligent Orchestrator logo, name, and version number, your user name, and the **Help** and **Logout** buttons. The path to the current page in the content area is also displayed.

Navigation trees

Navigation trees, displayed as tabs in the left navigation frame, provide access to all the system configuration and management options. The following navigation trees are available:

- Data center assets and resources
- System configuration and workflow management
- Customer applications
- Real-time performance monitoring
- Reports

Additional information on each of the tabs is available later in this section.

Content area

When you select an object on one of the tabs in the navigation tree, this area displays detailed information about the object. This area also displays additional tabs containing various configuration data for the selected object.

Command box

This editable box enables you to perform two operations:

- Execute SOAP commands available in the command-line interface
- Find various objects within the system without browsing the navigation trees

Also available in the bottom pane are the following buttons:

- **History**, which keeps track of all of the successful and failed commands per session
- **Overview**, which offers a 3D graphical overview of your data center

Navigation tree Overview

Using the navigation trees, you can configure and manage typical n-tier application architectures, as well as hardware assets and other system resources.

A brief description of the navigation trees follows:

Data center assets and resources



This navigation tree enables you to configure and manage all of your system's hardware assets and resources. From this tree, you can:

- Configure blade servers, boot servers, terminal servers, ACLs, firewalls, load balancers, power units, routers, servers, switches, switch fabrics, subnetworks, license pools, software products, software patches, and software stacks.
- Manage and configure all the available resource pools, and the resources that are currently under maintenance (applications, overflow servers, hardware resources).

System configuration and workflow management



This navigation tree enables you to configure and manage user, logical device, workflow and simple commands. From this tree, you can:

- Manage and configure all of the users that are currently registered with the system.
- Configure all logical device operations and device drivers defined in the system, as well as all available workflows, Java plug-ins, and simple commands.
- Add new workflows, copy or customize, execute, or delete any of the existing workflows, manage variables and transitions within workflows. Workflows are grouped by author and category.
- Determine the current state of a data center component configuration by reviewing the execution history of the workflows applied to it.

Customer applications



This navigation tree enables you to configure and manage typical n-tier application architectures. The in-depth structure of the **Customer applications** navigation tree includes *customer accounts, applications, clusters, and servers*, and makes it possible to define, configure, and manage all of these objects.

Real-time performance monitoring



This navigation tree enables you to monitor the system's health. From this tree, you can:

- List all the servers that are available in the resource pools.
- Monitor the overall performance of the customers, applications, clusters, and servers defined in the system, and also evaluate the performance of individual servers.

Reports



This navigation tree enables you to generate various reports on the status and usage of system resources. From this tree, you can:

- Generate reports on the availability of resources: which resources are mostly used, the utilization over a specified period of time.
- Predict future shortfalls of physical resources.
- Report such status as the current allocation of resources, the current status of resource pools.
- Obtain usage reports per cluster, application or customer account, detailing the type of resources, number of times used, total time used, and so on.
- Generate SLM reports showing the measured adherence to SLOs of each managed application.
- Generate network administration reports, showing all the commands executed on various network devices.

The command-line interface

The command-line interface is designed to be used by data center operators who prefer to access the system's internal properties and operations using the command line. The following figure shows an example of Tivoli Intelligent Orchestrator's command-line interface, illustrating the way a number of commands are implemented and executed.

```
D:\thinkcontrol\tools>D:\thinkcontrol\tools>soapcli tdtest tdtest http://localhost:9088/tn/soap/wsd/om/ibm/tn/soap/OperationsModeService.wsdl getGlobalMode
manual

D:\thinkcontrol\tools>soapcli tdtest tdtest http://localhost:9088/tn/soap/wsd/om/ibm/tn/soap/OperationsModeService.wsdl get@applicationMode 1342
default

D:\thinkcontrol\tools>soapcli tdtest tdtest http://localhost:9088/tn/soap/wsd/om/ibm/tn/soap/OperationsModeService.wsdl set@applicationMode 1342 manual

D:\thinkcontrol\tools>soapcli tdtest tdtest http://localhost:9088/tn/soap/wsd/om/ibm/tn/soap/OperationsModeService.wsdl get@applicationMode 1342
manual

D:\thinkcontrol\tools>
```

Figure 3 The command-line interface of Tivoli Intelligent Orchestrator

In this example, the operator uses the command-line interface to get the current global operating mode, to get the current operating mode for a certain application, and then to set the operating mode of the same application to **manual**.

Note: For additional information on the SOAP commands that can be executed using the Tivoli Intelligent Orchestrator's command-line interface, refer to the "SOAP commands available through command line" chapter, described later in this guide. The "SOAP command syntax reference" appendix provides syntax details and examples for some of the most commonly used SOAP commands.

The Operating Modes

The operating modes can be applied on two levels:

- To the entire Tivoli Intelligent Orchestrator application, known as the global operating mode.
- To individual applications or clusters under the system's management.

There are three main operating modes:

Automatic mode

In this mode, Tivoli Intelligent Orchestrator makes all of the resource change decisions, and instantly executes the necessary allocations. Recommendations are automatically authorized by the system and are subsequently acted upon by the Deployment Engine. For stability reasons, additional deployment requests for a given cluster will not be initiated while other deployment requests are pending.

Semiautomatic mode

In this mode, Tivoli Intelligent Orchestrator identifies the need for a resource change and issues a request. A human operator can view pending requests and must authorize a request before the Deployment Engine will act upon it.

Manual mode

In this mode, Tivoli Intelligent Orchestrator identifies the need for a resource change. Recommendations are generated but not acted upon by the Deployment Engine. The operator can review the recommendations but must explicitly execute the resource reallocations manually.

Changing the global operating mode or the current operating mode of an application or a cluster can be done using either the Web-based interface or the command-line interface.

Note: To learn how to change the global operating mode, refer to the “Changing the global operating mode” section in the “Configuring Tivoli Intelligent Orchestrator” chapter of this guide. For information on how to toggle between operating modes at application and cluster levels, refer to the “Changing the operating mode of an application” and “Changing the operating mode of a cluster” sections, in the “Managing Tivoli Intelligent Orchestrator” chapter.

The global operating mode can be inherited by all the customer applications managed by the system, provided the operating mode for each application was initially set to **default**. The global operating mode does not override any application mode other than **default**.

Similarly, the operating mode of an application can be inherited by all the clusters within that application, provided the clusters' operating mode was initially set to **default**. The application operating mode does not override any cluster mode other than **default**.

Application Topologies

Applications have both logical structures and physical deployments. Physical devices run elements of one or more applications. Tivoli Intelligent Orchestrator includes *Application Topologies v1.0*, which requires an abstraction of logical structures and physical deployments for applications. The same application is deployed into different environments, such as testing, staging, or production. Along with the device driver abstraction, the application structure makes it possible to automatically create the infrastructure that entire applications are based on. Likewise, the structure provides enough information to automatically generate appropriate workload models.

Three key concepts are at the core of this foundation:

Device driver

The *device driver* abstraction represents each different manufacturer's products and versions, and each device driver has a make and model number, tying the Deployment Engine workflows to a higher-level abstraction for all data center assets. This abstraction also models the common functionality of logical devices. For example, three different switch drivers can all perform the same high-level function, for example, assigning a port to a VLAN, although they will require different workflows. Also, a device driver for a load balancer also performs the same functions as a layer-3 switch. All of these are treated as logical device support.

Switch fabric

The *switch fabric* represents a group of switches and routers combined together to form a huge logical multi-layer switch. The *switch fabric* can be programmed to create ad hoc networking topologies.

Software stack

The *software stack* concept is introduced to simplify the relationship between workflows and software products, by modeling all the software dependencies and identifying software products. For example, for each server in the data center, you can interactively add and remove software by automatically invoking the appropriate workflows as defined by the *device driver*.

Through application topologies, the Deployment Engine workflows become device drivers as part of Tivoli Intelligent Orchestrator's support for various hardware and software products. The system extends the existing Data Center Model and Deployment Engine components, enabling a new way of viewing workflows.

Note: For details on the Deployment Engine's workflows, refer to the "Configuring and using workflows" chapter, later in this guide.

By means of the *device driver* workflow support, you can create new workflows that embed device actions as commands. The system determines which actual workflows must be executed based on the Data Center Model data for that device. Using SOAP, you can also initiate any of the device driver workflows.

Device driver support for workflow integration

The logical devices supported by Tivoli Intelligent Orchestrator are:

- Application
- Boot server

- Cluster
- Device
- File repository
- Firewall
- IP system
- Load balancer
- Power unit
- Router
- Service access point
- Software
- Software stack
- Spare pool
- Switch

The following device driver support is implemented in this release, organized into the following hierarchy of device driver categories:

- **Boot Servers**
 - JumpStart Server
 - Rembo Boot Server
- **Clusters**
 - Default Lab Cluster
- **Custom**
- **File Repositories**
 - CVS File Repository
 - UNIX File System Repository
 - Windows File System Repository
- Image Software Stack
 - Image Software Stack
- Network Appliances
 - Firewalls
 - Cisco PIX
 - Load balancers
 - Alteon Load Balancer
 - Cisco CSS 11000
 - Routers
 - Cisco 2621
 - Switches
 - Cisco 3548
 - Cisco 6500 Switches Hybrid Mode
 - Cisco 6500 Switches Native IOS Mode
- Regular Software Stack
- Servers
 - Virtual Server
- Service Access Points

- SNMP V1 Service Access Point
- SSH Service Access Point
- Software Patches
 - Copy, UnTar, and Install
 - Copy, UnZip, and Install
 - Self Extracting *.exe
 - Windows Hot Fix
- Software Products
 - MS Internet Information Server App
 - RedHat Linux Operating System
 - ThinkUtils for Windows Family
 - Windows Operating System
- Software Stacks
 - Virtual Server Installable Software Stack
- Storage Systems
- [other]
 - Default Device Driver

Each data center device can either inherit all of its workflows from the associated device driver or can override one or more workflows with its own custom workflows.

Interactive user interface

An interactive user interface enables you to access views to various devices in your data center and perform deployment tasks. In the background, workflows are initiated for each user interface manipulation. The interactive user interface is focused around the following entities: switch fabric, layer-2 switch, load balancer, router, firewall, NIC, server, cluster, and spare pool. Interactive views are available for switch fabrics, switches, VLANs,

routers, firewalls, and NICs on servers, and can be easily accessed in a few different ways. Some interactive views include a toolbar, which provides further access to other views. The toggle buttons provided on this toolbar are:



Devices Legend ()

Use this button to display/hide the Devices Legend window, which lists all the servers and network devices currently available in the system. Click any device icon in this list to access its interactive view.

VLANs Legend ()

Use this button to display/hide the VLANs Legend window, which lists all the VLANs currently available in the system. Click any VLAN icon in this list to access its interactive view.

Refresh ()

Use this button to refresh the current view. It updates the displayed information with current data.

Note: For additional information on all accessible interactive views, refer to the “Configuring Tivoli Intelligent Orchestrator” chapter.

Software stack

For each server in the data center, a **Software** tab lists all the software products installed on that server. You can interactively add and remove software by automatically invoking the appropriate workflows as defined by the device driver. However, you have to ensure that the correct software products are compatible and could be stacked together as the software will not enforce dependencies.

SOAP initiation of device functions

You can initiate any of the device driver workflows using a number of SOAP commands available through the command line.

Note: For additional information, refer to the “SOAP commands initiating device functions” section in the “SOAP commands available through command line” chapter, later in this guide.

Chapter 2. Configuring Tivoli Intelligent Orchestrator

Various configurations can be performed with Tivoli Intelligent Orchestrator: some affect the entire system and others affect only certain system components and individual data center devices. This chapter describes the following:

- How to set up the operating mode at global, application, and cluster levels.
- How to configure software products and stacks.
- How to configure servers, assets and resources, such as switches, load balancers, routers, firewalls, switch fabrics, subnets, VLANs, virtual IPs, boot servers, and blade servers.

The configuration procedures affecting the entire system are available from the **System configuration and workflow management** navigation tree. The configuration procedures affecting software products, stacks, and individual devices are available from the **Data center assets and resources** navigation tree.

Most of the device configuration procedures can be performed only in the maintenance area, for example, editing the properties of a device, setting up management interfaces for certain devices, creating new devices, or removing them. None of the required configuration procedures is possible while the data center device is functioning in a customer application environment.

Maintenance is designed to serve as a staging area for devices that are not under the system's control. A device in maintenance can be either:

- A faulty overflow server or device that has been marked as “failed,” and temporarily placed into maintenance to repair or reconfigure, or
- An entire customer application that needs to be reconfigured. For example, to be able to edit the properties of a customer application or modify some of its configuration settings, you must move it to maintenance first, and then make the required changes.
- A newly created data center device that has not yet been allocated to a customer application or resource pool

After being repaired or reconfigured, a data center device in maintenance must be manually moved back to the customer application or resource pool it belongs to. This procedure is generically called “bringing (the device) online”.

Optional configurations, such as setting up the routing configuration of specific devices are also available, and may be necessary, for example, for creating and executing various workflows.

This chapter includes information on the following:

- Signing on and Signing off
- Configuring users
- Changing the global operating mode
- Configuring inventory resources
- Configuring resource pools

These procedures are described in detail in the following sections.

Note: For more information on how to set up and perform various tasks with applications, clusters, servers, inventory resources, and so on, refer to the corresponding sections in the “Managing Tivoli Intelligent Orchestrator” chapter, described later in this guide.

Signing on and Signing off

The procedures described in this section assume that Tivoli Intelligent Orchestrator is already successfully installed.

Test

Signing on to the system

Before you sign on to the system, ensure you have:

- The appropriate Internet browser, such as Microsoft Internet Explorer v.5.0 or higher.
- The IP address of the server where the Tivoli Intelligent Orchestrator is installed. Contact your installation team to obtain this IP address.
- Your user name and password for Tivoli Intelligent Orchestrator. Contact your installation team to register as a new user with the system.

To sign on to Tivoli Intelligent Orchestrator:

1. Open your Internet browser. In the **Address** box, type the following address:
`http://<Tivoli Intelligent Orchestrator_hostname>:9080/tcWebUI/login.jsp`
where *<Tivoli Intelligent Orchestrator_hostname>* is the full DNS name of the server. The Sign On screen is displayed.
- Note:** You must use the full DNS name of the server. Sign on will fail if only the IP address is used.
2. Enter your user name and password, and then click **OK**.

Note: Concurrent logins using the same user name and password are allowed.

Signing off from the system

To sign off from the system, click the **Logoff** button, which is located on the right side of Tivoli Intelligent Orchestrator’s header. The Sign Off screen is displayed.

You will be logged out of the system automatically after thirty minutes of session inactivity.

Configuring users

This section describes the following procedures:

- Adding a new user
- Editing a user profile
- Removing the user

Adding a new user

To add a new user:

1. Click **System configuration and workflow management > Users**. The **Users** tab displays all of the users that are currently registered with the system.



2. Click **Add user** . The User Data dialog box is displayed.
3. Enter the required user data in the appropriate text boxes. Under **Access**, select the appropriate check boxes to grant the relevant access privileges to the new user.

Note: For details on access privileges for users, refer to the following subsection, “User Access Privilege Descriptions”.

4. At the bottom of the window, enter the password for the new user account.
5. Click **Save**.
6. Click the **System configuration and workflow management** navigation tree to refresh the displayed data.

User Access Privilege Descriptions

The table below lists the currently defined access privileges that can be assigned to users accessing Tivoli Intelligent Orchestrator. User access to certain functionality in the Web-based user interface, as well as access to a number of SOAP commands available through command line can be allowed or restricted, based on the assigned access privileges.

Note: For additional information on all the SOAP commands that are available through the command line, refer to the “SOAP commands available through command line” chapter, later in this guide.

Table 1. User credentials for accessing Tivoli Intelligent Orchestrator

Access privilege	Description
Tivoli Intelligent Orchestrator login	Grants basic access to Tivoli Intelligent Orchestrator, data center monitoring, and operation management.
user defined variables	Grants access to user-defined object variables. It requires access to Tivoli Intelligent Orchestrator.
user administration	Grants access to user creation and user privilege control through the Web-based interface. It requires access to Tivoli Intelligent Orchestrator.
operating mode query	Grants read-only access to operating modes, through the OperationsModeService.wsdl.
operating mode change	Grants write access to operating modes, through the OperationsModeService.wsdl.
recommendation change	Grants write access to recommendations, through the RecommendationsService.wsdl.
cluster management	Grants access to the bookkeeping of servers added to or removed from a cluster.
ip address management	Grants access to IP address allocation.
spare pool management	Grants access to spare pool management processes such as choosing a server for an application, adding new servers to or removing new servers from the pool.

Table 1. User credentials for accessing Tivoli Intelligent Orchestrator

Access privilege	Description
object lock object unlock object lock check	Grants access to persistent object lock management processes.
fault management	Grants access to fault management processes such as setting or clearing the fail flag for servers and other data center devices. For details, refer to the FaultManagementService.wsdl.
Boot Server interaction	Grants access, through the command line, to logical commands for deployments on boot servers. For details, refer to the BootServerComponentService.wsdl.
Cluster interaction	Grants command line access to logical commands for deployments on clusters. For details, refer to the ClusterComponentService.wsdl.
Device interaction	Grants command line access to logical commands for initiating devices, through the DeviceComponentService.wsdl.
File Repository interaction	Grants command line access to logical commands for deployments on file repository systems. For details, refer to the FileRepositoryComponentService.wsdl.
Firewall interaction	Grants command line access to logical commands for setting up firewall configurations. For details, refer to the FirewallComponentService.wsdl.
IP System interaction	Grants command line access to logical commands for configuring IP systems. For details, refer to the IpSystemComponentService.wsdl.
Load Balancer interaction	Grants command line access to logical commands for deployments on load balancers. For details, refer to the LoadBalancerComponentService.
Power Unit interaction	Grants command line access to logical commands for controlling power units. For details, refer to the PowerUnitComponentService.wsdl.
Router interaction	Grants command line access to logical commands for routers. For details, refer to the RouterComponentService.wsdl.
Software interaction	Grants command line access to logical commands for software deployments on managed devices. For details, refer to the SoftwareComponentService.wsdl.
Software Stack interaction	Grants command line access to logical commands for software stacks. For details, refer to the SoftwareStackComponentService.wsdl.
Spare Pool interaction	Grants command line access to logical commands for deployments on spare pools. For details, refer to the SparePoolComponentService.wsdl.
Switch interaction	Grants command line access to logical commands for deployments on switches. For details, refer to the SwitchComponentService.wsdl.
Switch Fabric interaction	Grants command line access to logical commands for deployments on switch fabrics. For details, refer to the SwitchFabricComponentService.wsdl.

Editing a user profile

To edit the profile of an existing user:

1. Click **System configuration and workflow management > Users**. The user list is displayed.
2. Find the user whose profile you want to edit. Click the **Edit user**  button. The User Data dialog box is displayed.
3. Modify the user data as required.
4. Click **Save**.

Removing the user

To remove a user:

1. Click **System configuration and workflow management** > **Users**. The users list is displayed.
2. Locate the user profile to be removed. Click the **Delete user** button .
3. When prompted, click **OK**.

Changing default passwords

Tivoli Intelligent Orchestrator requires that a default set of user IDs and default passwords are created and used during installation and configuration. A command line tool is provided to change the passwords after Tivoli Intelligent Orchestrator has been installed. Using the command line tool, you can change the passwords for the following user IDs:

- tioldap
- wasadmin
- tioappadmin
- tiointernal
- root
- tiodb

To change the password for one of the default user IDs, do the following:

1. Ensure the following variables are currently defined: `WAS_HOME`, `JAVA_HOME`, and `TC_HOME`. These environment variables should be defined after the Tivoli Intelligent Orchestrator or Tivoli Provisioning Manager installation process is complete.
2. Ensure the WebSphere(R) Application Server is started.
3. Log on as `tioadmin`.
4. Switch to the `TIO_installdir/tools` directory, where `TIO_installdir` is the directory where Tivoli Intelligent Orchestrator or Tivoli Provisioning Manager is installed.
5. Run one of the following commands, depending on which platform you are on:
Windows: `changePassword.cmd <user_ID> <new_password> <current_was_password>`
AIX: `changePassword.sh <user_ID> <new_password> <current_was_password>`
Linux: `changePassword.sh <user_ID> <new_password> <current_was_password>`

where the variables are defines as follows:

Variable	Description
<code><user_ID></code>	The user ID whose password you want to change.
<code><new_password></code>	The new password you want to use for the user ID.
<code><current_was_password></code>	The current password for the wasadmin user ID. If you have not yet changed the wasadmin password, use the default value as defined in the table below.

You can only change the password for one user ID at a time.

Table 2: Default user names and passwords

User name	Default password	Description
tiadmin	<user defined>	<ul style="list-style-type: none"> Defined in the operating system. Used to log onto the operating system. used to install Tivoli Intelligent Orchestrator or Tivoli Provisioning Manager.
tioldap	tioldap	<ul style="list-style-type: none"> Defined in IBM Directory Server Used by WebSphere to connect to the Directory Server This account must be given rights to search user accounts.
root	think4me	<ul style="list-style-type: none"> Defined in IBM Directory Server This is the user ID prompted for during Tivoli Intelligent Orchestrator and on the Directory Server panel.
wasadmin	wasadmin	<ul style="list-style-type: none"> Defined in IBM Directory Server Used by WebSphere as the administrator account Used to start, stop, and manage WebSphere
tioappadmin	tioappadmin	<ul style="list-style-type: none"> Defined in IBM Directory Server Used to log into the Tivoli Intelligent Orchestrator and Tivoli Provision in Manager web console This is the initial application user which should have all rights
tiointernal	internal	<ul style="list-style-type: none"> Defined in IBM Directory Server Used by the Tivoli Intelligent Orchestrator for system initiated actions
tiodb	think4me	<ul style="list-style-type: none"> Defined in the database server This is the database user

- Except when changing the password for user ID `tioappadmin`, you must restart the WebSphere Application Server and Tivoli Intelligent Orchestrator or Tivoli Provisioning Manager, after each password change for the change to take effect. Refer to Appendix C of the Installation Guide for information on stopping and starting Tivoli Intelligent Orchestrator or Tivoli Provisioning Manager.

Note: Stopping and starting the TIO or TPM server requires you to enter the `wasadmin` user name and password. After using the `changePassword` command to change the `wasadmin` password, stop the server and enter the old `wasadmin` password. When starting the server, with the `tio.sh start` command, use the new `wasadmin` password.

- If you are changing the passwords for `root` and `tiodb`, the password change made by the command line tool is only registered within Tivoli Intelligent Orchestrator or Tivoli Provisioning Manager. Additionally, you must also make the password change

within the IBM Directory Server for user ID root, and within DB2 Universal Database for user ID tiodb. Refer to the respective product documentation for instructions on changing passwords for those user IDs.

Changing the global operating mode

The global operating mode is the mode that the entire Tivoli Intelligent Orchestrator system is running in. The available global operating modes, **automatic**, **semiautomatic**, and **manual**, are described in detail in the “The Operating Modes” section, in the preceding “Introduction to the IBM Tivoli Intelligent Orchestrator” chapter.

You can change the Tivoli Intelligent Orchestrator’s global operating mode using either the system’s Web-based interface or the command-line interface, provided you are granted the *user administration* access privilege.

Note: For more information on access privileges, refer to the “User Access Privilege Descriptions” section, later in this chapter.

You are also provided with the possibility to change the operating mode at the application or cluster level. The global operating mode can be inherited by all the customer applications managed by the system, provided the operating mode for each application was initially set to **default**. The global operating mode does not override any application mode other than **default**. Similarly, an application mode can be inherited by all the clusters within that application, provided the operating mode for each cluster was initially set to **default**. The application operating mode does not override any cluster mode other than **default**.

Note: For more information on how to switch operating modes at the application level, refer to the “Changing the operating mode of an application” section in the “Managing Tivoli Intelligent Orchestrator” chapter.

The procedures to change the global operating mode using either the Web-based interface or the command-line interface are described in the following sections.

Using the Web-based interface

You can change the current global operating mode in two different ways using the Web-based interface:

- Using the System configuration and workflow management menu
- Using the Command box

Using the System configuration and workflow management menu

To change the global operating mode using the **System configuration and workflow management** menu:

1. Click **System configuration and workflow management > Configuration**. The **Configuration** tab is displayed.
2. In the **Change Mode to** list, click the global operating mode you want to switch to. The default mode is **automatic**.
3. Click **Apply**.

Using the Command box

To change the operating mode of the system using the **Command** box:

1. In the **Command** box, type the appropriate SOAP command to switch to the operating mode you want using the following format:

```
soapcli.sh <username> <password>  
http://localhost:9080/tcSoap/wsdl/com/ibm/tc/soap/OperationsModeService.w  
sdl setSetGlobalMode <new-mode>
```

where <new-mode> is one of the following values:

manual

In this mode, Tivoli Intelligent Orchestrator identifies the need for a resource change. The operator can review the recommendations but must explicitly execute the resource reallocations manually.

automatic

In this mode, Tivoli Intelligent Orchestrator makes all of the resource change decisions, and instantly executes the necessary allocations.

semi-automatic

In this mode, Tivoli Intelligent Orchestrator identifies the need for a resource change and issues a request. A human operator can view pending requests and must authorize a request before it is executed.

2. Press Enter.

Note: For additional information on all the SOAP commands that are available through the Tivoli Intelligent Orchestrator's Web-based interface, refer to the "SOAP commands available through command line" chapter, described later in this guide.

Using the command-line interface

To change the current global operating mode using the command-line interface, use the following SOAP command:

```
soapcli.sh <username> <password>  
http://localhost:9080/tcSoap/wsdl/com/ibm/tc/soap/OperationsModeService.w  
sdl setSetGlobalMode <new-mode>
```

where <new-mode> is one of the following values:

manual

In this mode, Tivoli Intelligent Orchestrator identifies the need for a resource change. The operator can review the recommendations but must explicitly execute the resource reallocations manually.

automatic

In this mode, Tivoli Intelligent Orchestrator makes all of the resource change decisions, and instantly executes the necessary allocations.

semi-automatic

In this mode, Tivoli Intelligent Orchestrator identifies the need for a resource change and issues a request. A human operator can view pending requests and must authorize a request before it is executed.

Note: For additional information on all the OperationsModeService SOAP commands that are available through the Tivoli Intelligent Orchestrator's command-line interface, refer to the "SOAP commands available through command line"

chapter, described later in this guide.

Configuring inventory resources

The following hardware assets and resources are grouped under the **Data center assets and resources > Inventory** menu:

- Blade servers
- Boot servers
- Terminal servers
- ACLs
- Firewalls
- Load Balancers
- Power Units
- Routers
- Servers
- Switch Fabrics
- Subnets
- License Pools
- Software Products
- Software Patches
- Software Stacks

To illustrate the interdependency that exists between the hardware resources managed by the system, the following procedures are presented in a top-down approach. For example, to be able to create a new virtual LAN (VLAN), you need to create a subnet first. Also, to be able to create and configure a resource pool or an application cluster, you must have a VLAN. Therefore, the logical order to create and configure new assets and resources in Tivoli Intelligent Orchestrator is as follows:

Switch Fabric > Subnet > VLAN > Switch > Load Balancer > Resource Pools > Servers.

The following sections describe the procedures to configure these resources:

- Configuring switch fabrics
- Configuring subnets
- Configuring VLANs
- Configuring switches
- Configuring load balancers
- Configuring routers
- Configuring firewalls
- Configuring Virtual IPs (VIPs)
- Configuring servers
- Configuring boot servers
- Configuring blade servers
- Configuring terminal servers
- Configuring ACLs
- Configuring power units
- Configuring license pools
- Configuring software products
- Configuring software patches
- Configuring software stacks

Configuring subnets

A newly added subnet will be displayed in both of the following menus:

- **Data center assets and resources > Inventory > Subnets**
- **Data center assets and resources > Maintenance > Subnets.**

While the subnet is in maintenance, you can edit its properties, assign workflows to it, set up blocking IP addresses for it, or remove it. To be managed by Tivoli Intelligent Orchestrator, the subnet must be removed from the maintenance area.

Adding a new subnet

To add a new subnet:

1. Click **Data center assets and resources > Inventory > Subnets**. The Subnet Inventory screen displays all the subnets that are currently available.



Subnet IP Address	Subnet Name
10.1.1.0	10.1.1.0/24
10.1.1.10	10.1.1.10/24
10.1.1.22	10.1.1.22/24
10.1.1.40	10.1.1.40/24
10.1.1.64	10.1.1.64/24
10.1.2.10	10.1.2.10/24
10.1.2.11	10.1.2.11/24
10.1.3.10	10.1.3.10/24
10.1.3.40	10.1.3.40/24
10.1.3.64	10.1.3.64/24
10.1.3.80	10.1.3.80/24
10.1.3.95	10.1.3.95/24
10.1.4.10	10.1.4.10/24

2. Click **Edit > Add subnetwork**. The New Subnetwork dialog box is displayed.
3. Type a network address and a mask for the new subnet.
4. Click **Save**.

Editing the subnet properties

You can only edit the properties of a subnet in maintenance. To edit the properties of a subnet:

1. Click **Data center assets and resources > Maintenance > Subnets > <Subnet name>**.
2. Click **Edit > Properties**. The following dialog box is displayed

Note: You can also select a subnet to access by clicking **Data center assets and resources > Inventory > Subnets** to display a list of all available subnets. In the list, find the offline subnet whose properties you want to edit. Click its **More** button, and then choose the **Edit** menu option.

3. Edit the subnet's properties as required.
4. Click **Save**.

Assigning workflows to the subnet

To learn how to assign workflows to a subnet, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Setting up blocking IP addresses for the subnet

You can only set up blocking IP addresses for a subnet in maintenance. To set up the blocking IP addresses for a subnet:

1. Click **Data center assets and resources > Maintenance > Subnets > <Subnet name>**. Click the **Blocked IP's** tab.

Note: You can also access the subnet by clicking **Data center assets and resources > Inventory > Subnets** to display a list of all available subnets. In the list, find the offline subnet for which you want to set up the blocking IPs. Click the **Blocked IP's** tab.

2. Type the blocking IP addresses as required.
3. Click **Add**.

Removing the subnet

You can only remove a subnet that is in maintenance.

Attention: Subnets are used by the entire Tivoli Intelligent Orchestrator system. Removing a subnet may affect more than one customer application.

To remove a subnet:

1. Click **Data center assets and resources > Maintenance > Subnets**. In the list, find the subnet you want to remove, and then click its  **Delete subnet** button.

Note: You can also access the subnet by click **Data center assets and resources > Inventory > Subnets** to display a list of all available subnets. In the list, find the offline subnet you want to remove. Click its **More**  button, and then select the  **Delete** menu option.

2. When prompted, click **OK**.

Configuring VLANs

Note: This procedure requires that you have already created and configured a subnet. If you have not configured a subnet, refer to the previous section, “Configuring subnets”.

A new VLAN is added to the list of VLANs that are currently in maintenance. You can edit the properties of a VLAN in maintenance, or remove it. You can also access the VLAN view, which offers details on the VLAN’s topology and gives further access to other interactive views.

Adding a new VLAN

To add a new VLAN:

1. Click **Data center assets and resources > Inventory > Switch Fabrics > Default Fabric > VLANs**. The **VLANs** tab is displayed.
2. Click **Edit** >  **Add VLAN**. The New VLAN dialog box is displayed.
3. Type a name for the new VLAN.
4. In the **VLAN Number** text box, type a number for the new VLAN.
5. In the corresponding lists, click the switch fabric and the subnetwork for the new VLAN.
6. Click **Save**.

Note: For additional information on how to add a VLAN to a switch fabric by initiating this deployment task from an interactive view, refer to the “Creating a VLAN” section in the “Managing Tivoli Intelligent Orchestrator” chapter.

Editing the VLAN properties

To edit the properties for a VLAN in maintenance:

1. Click **Data center assets and resources > Maintenance > VLANs > <VLAN name>**. Click **Edit >  Properties**.

Note: You can also access the VLAN by clicking **Data center assets and resources > Inventory > Switch Fabrics > Default Fabric > VLANs** to display a list of all VLANs. In the list, find the offline VLAN you want to edit. Click its **More ** button, and then select the ** Edit** menu option.

2. Edit the VLAN properties as required.
3. Click **Save**.

Assigning workflows to the VLAN

To learn how to assign to a VLAN, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Removing the VLAN

A VLAN must be in maintenance mode before it can be removed. To remove a VLAN:

1. Click **Data center assets and resources > Maintenance > VLANs**. In the list, find the VLAN to be removed, and then click its ** Delete VLAN** button.

Note: You can also access the VLAN by clicking **Data center assets and resources > Inventory > Switch Fabrics > Default Fabric > VLANs** to display a list of all VLANs. In this list, find the offline VLAN you want to remove, click its **More ** button, and then select the ** Delete** menu option.

2. When prompted, click **OK**.

Note: For additional information on how to remove a VLAN from a switch fabric by initiating this deployment task from an interactive view, refer to the “Removing the VLAN from the switch fabric” section in the “Managing Tivoli Intelligent Orchestrator” chapter.

Configuring switches

Note: This procedure requires that you have already created and configured a switch fabric. If you have not configured a VLAN, refer to the previous section, “Configuring VLANs”.

A new switch is added to the list of switches that are currently in maintenance. You can edit the properties of a switch in maintenance, edit its user-defined variables, set up its ports and network interface, set up its routing configuration, or remove it.

Adding a new switch

To add a new switch:

1. Click **Data center assets and resources > Inventory > Switch Fabrics** and select **Default Fabric**.
2. Click **Edit** and select **Add Switch**.
3. Complete the fields on the Switch Data dialog box as follows:
 - a. Type a name for the new switch.
 - b. In the **Fabric** list, select a fabric for the switch. If required, enter the TCP port number for the switch.
 - c. Clear the **Failed** check box.
4. Click **Save**.

Editing the switch properties

A switch must be in maintenance before you can edit its properties. To edit the properties of a switch:

1. Click **Data center assets and resources > Maintenance > Switches > <Switch name>**. Click **Edit >  Properties**. The Switch Data dialog box is displayed.
2. Edit the switch properties as required.
3. Click **Save**.

Accessing a switch

You can access a switch by clicking **Data center assets and resources > Inventory > Switch Fabrics > Default Fabric** and selecting a switch. If you select, **View >  Icon View**.

The displayed devices are grouped under three main categories: **VLANs**, **Ports**, and **Devices** (servers). If the **Show All** check box under each category name is selected, all the devices available for that category are displayed. All the links between the displayed devices are also visible.

Showing or hiding devices, links, and linked objects

An entire category of devices can be shown in or hidden from this view by selecting or clearing its corresponding **Show All** check box. You can also control the way individual devices and their links and linked objects are shown on the screen. For example, to hide a switch, right-click the device and select **Hide Switch**.

For every device type, the pop-up menu options are:

- **Hide <Device type>**, which hides the selected device. This is the only menu option available if the selected device is not linked to other objects.
- **Hide All Links**, which hides all of the links to the selected device.
- **Hide All Linked Objects**, which hides all of the objects linked to the selected device. The **Show All** check box corresponding to the hidden objects is automatically cleared. To redisplay the hidden objects, select the **Show All** check box.
- **Show Linked Objects Only**, which shows only the objects that have links to other objects, and hides all of the unlinked objects.

In addition to the menu options above, menu options that initiate deployment tasks are also available for some devices. For example, the **Hardware Reboot** option is available for servers and you can turn the port associated with a switch on and off.

Assigning workflows to the switch

To learn how to assign workflows to a switch, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Adding a new switch port

To add a new switch port to the switch:

1. Click **Data center assets and resources > Inventory > Switch Fabrics** > *<Switch Fabric name>* <Switch name>.
2. Click **Edit** >  **Add Port**. The New Port dialog box is displayed.
3. Type the port number and the module for the new switch port you intend to add. In the **VLAN** list, select a VLAN for the new switch port. The NIC for the switch port should be set up on the **NICs** page of the corresponding server.
4. Click **Save**. The new switch port is added to the list of available ports for that switch.

Setting up a network interface for the switch

To set up a network interface for a switch:

1. Click **Data center assets and resources > Inventory > Switch Fabrics** > *<Switch Fabric name>* <Switch name>. Click the **Info** tab.
2. Click **Edit** >  **Add Interface**. The New Network Interface dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **Name** text box, type a name for the network interface.
 - b. In the **IP Address** text box, type an IP address for the switch.
 - c. In the **Subnetwork** list, select a subnet.
 - d. In the **ACL** list, select an Access Control List (ACL) to map it to the new network interface. Clear the **Failed** check box.
 - e. Select the **Managed** check box if the network interface is intended for management purposes.
 - f. Select **Perform using dcm.interaction** to invoke the IPSystem.Add IP Address logical operation. This adds the specified IP address (network interface) to the specified IP network device.
4. Click **Save**.
5. If the switch has multiple IP addresses, repeat steps 2 through 4 for each of them, so that these IP addresses will not be used for other devices.

Setting up the routing configuration for the switch

To set up the routing configuration for the switch:

1. Click **Data center assets and resources > Inventory > Switch Fabrics** > *<Switch Fabric name>* <Switch name>. Click the **Info** tab.
 - a. Select the **router** check box to enable the switch to act as a router.
 - b. Select the **firewall** check box to enable the router to act as firewall.

2. To display further menu options on the screen, click the **More**  button, and then select the **Add Route** menu option. The New Route dialog box is displayed.
3. Under **Destination**, specify the IP address of the destination subnet, and then enter the gateway for the enabled router. Click **Save**.
4. Repeat steps 3 and 4 to define more routes for the router.
5. If applicable, instead of defining one route after another through steps 3 and 4, you can click the **More**  button, and then select the **Apply Routing Table** menu option. The Apply Routing Table dialog box is displayed.
6. In the **Routing Table** list, click the routing table you want to apply to the router. Select the **Remove existing routes** check box if you want to remove all the individual routes that have been previously set up for this device. Click **Save**. This will invoke the **dcm.interaction** operation.

Removing the switch

A switch must be in maintenance before you can remove it. To remove a switch:

1. Click **Data center assets and resources > Maintenance > Switches** to display a list of all the switches that are currently in maintenance. Find the switch you want to remove, and then click the **Delete switch**  button.
2. When prompted, click **OK**.

Configuring switch fabrics

This section describes the following procedures:

- Accessing the switch fabric view
- Adding a new switch to the default fabric

Accessing the switch fabric view

You can access the switch fabric view in either of the following ways:

- Using the **Icon view**  button on the Switch Fabric Info screen, or
- From any other interactive view by clicking the switch fabric icon.

To access the switch fabric view using the **Icon view** button:

1. Click **Data center assets and resources > Inventory > Switch Fabrics > Default Fabric**. The Info screen is displayed.
2. Click **Edit >  Icon view**. The “Overview: Default Fabric” screen is displayed.

The displayed devices are grouped under three main categories: **Switches**, **Network Topology** (VLANs and routers), and **Devices**. If the **Show All** check box under each category name is selected, all of the devices available for that category are displayed. All of the links between the displayed devices are also visible.

Accessing other interactive views from the switch fabric view

From the switch fabric view, you can quickly access the following views:

- *Server Info view*, by clicking the server icon
- *VLAN view* by clicking the VLAN icon
- *Router view* by clicking the router icon attached to a VLAN
- *All Routes* by hovering over a VLAN’s upper corner router icon

When selecting a VLAN, all connected devices should also be selected

When selecting a device, all associated VLANs should also be selected

Showing or hiding devices, links, and linked objects

An entire category of devices can be shown in or hidden from the switch fabrics view by selecting or clearing its corresponding **Show All** check box. You can control the way individual devices and their links and linked objects are shown on the screen. For example, to hide a device such as a VLAN, right-click the device and then select **Hide VLAN** from the displayed pop-up menu.

For every device type, the pop-up menu options are:

- **Hide <Device type>**, which hides the selected device. This is the only menu option available if the selected device is not linked to other objects.
- **Hide All Links**, which hides or shows all of the links to the selected device.
- **Hide All Linked Objects**, which hides all of the objects linked to the selected device. The **Show All** check box corresponding to the hidden objects is automatically cleared. To redisplay the hidden objects, select the **Show All** check box.
- **Show Linked Objects Only**, which shows only the objects that have links to other objects, and hides all of the unlinked objects.

In addition to the menu options above, menu options that initiate deployment tasks are also available for some devices. For example, the **Turn Port On** and **Turn Port Off** option is available for ports, and the **Hardware Reboot** option is available for servers.

Note: For step-by-step instructions on how perform deployment tasks from the switch fabric view, refer to the “Performing deployment tasks from the interactive views” section, in the “Managing Tivoli Intelligent Orchestrator” chapter.

Adding a new switch to the default fabric

To add a new switch to the default fabric:

1. Click **Data center assets and resources > Inventory > Switch Fabrics > Default Fabric**. The “Switch Fabric: Default Fabric” screen displays the switches available for the default fabric.
2. Click **Edit >  Add switch**. The Switch Data dialog box is displayed.
3. On the Switch Data dialog box, fill in the fields as follows:
 - a. Type a name for the new switch.
 - b. In the **Fabric** list, click **Default Fabric**.
 - c. If required, type the port numbers for the new switch. Clear the **Failed** check box.
4. Click **Save**.

Click the **Data center assets and resources** navigation tree to refresh the displayed data. The new switch is also added to the **Switches** list

Note: For additional information on how to edit the switch’s properties and user variables, and on how to set up its network interfaces and routing configuration, refer to the “Configuring switches” section, earlier in this chapter.

Configuring load balancers

A new load balancer is added to the list of load balancers that are currently in maintenance. You can edit the properties of a load balancer, edit its user-defined variables, set up its network interface and routing configuration, or remove it.

Adding a new load balancer

To add a new load balancer:

1. Click **Data center assets and resources > Inventory > Load Balancers**. The Load Balancer Inventory screen displays all the load balancers that are currently available.
2. Click **Add load balancer** . The Load Balancer Data dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **Load Balancer Name** text box, type a name for the new load balancer.
 - b. In the **Driver ClassName** list, click a driver class name for the new load balancer. Clear the **Failed** check box.
4. Click **Save**.

Configuring the switch functionality of the load balancer

To configure the switch functionality of the load balancer:

1. Click **Data center assets and resources > Inventory > Switch Fabrics > <Switch Fabric name>**. The Switch Fabric screen displays the available switches.
2. Click **Edit** and select **Add switch**. The Switch Data dialog box is displayed.
3. Complete the fields as follows:
 - a. For the switch name, type the same name as the load balancer. In the **Fabric** list, click a fabric for the new switch.
 - b. If required, type the TCP port number for the new switch. Clear the **Failed** check box.
 - c. Click **Save**. The new switch is added to the **Switches** list.
4. Click **Data center assets and resources > Inventory > Load Balancers > <New load balancer name>** and click the **Switch** tab to display the switch configuration for the newly created load balancer.

Editing the load balancer properties

The load balancer must be in maintenance before you can edit its properties. To edit the properties of a load balancer:

1. Click **Data center assets and resources > Maintenance > Load Balancers** and select the load balancer whose properties you want to edit. Click the **Properties** tab.

Note: You can also access the load balancer properties by clicking **Data center assets and resources > Inventory > Load Balancers** to display all the load balancers that are currently available. Find the offline load balancer whose properties you want to edit, and then click its **Edit load balancer** button . The **Properties** tab is displayed.

2. Edit the load balancer's properties as required.
3. Click **Save**.

Assigning workflows to the load balancer

To learn how to assign workflows to a load balancer, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, later in this guide.

Adding a VIP to the load balancer

To add a virtual IP (VIP) to the load balancer:

1. Click **Data center assets and resources > Inventory > Load Balancers > <Load balancer name>**. The **Info** tab is displayed.
2. Click **Add Virtual IP** . The **Virtual IP Data** dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **Virtual IP Address** text box, type the address of the new VIP.
 - b. In the **Load Balancer Name** list, click a load balancer for the new VIP.
 - c. In the **Cluster** list, click the load balanced cluster for the new VIP.
 - d. Type the balancing algorithm based on the specific device driver.
 - e. In the appropriate text boxes, specify the output TCP port number and the range of ports for which the algorithm will be performed.
4. Click **Save**. This will invoke the **dcm.interaction** operation.

Note: For more details on how to add a VIP to a load balancer by initiating this deployment task from the NIC interactive view, refer to the “Creating a virtual IP” section in the “Managing Tivoli Intelligent Orchestrator” chapter.

Setting up a network interface for the load balancer

To set up a network interface for the load balancer:

1. Click **Data center assets and resources > Inventory > Load Balancers > <Load balancer name>**. Click the **Networking** tab.

Note: You can also access the network interface by clicking **Data center assets and resources > Inventory > Load Balancers** to display all the load balancers that are currently available. In the list, find the load balancer for which you want to set up a network interface, and then click its **Edit load balancer** button . Click the **Networking** tab.

2. Click **Edit > Add Interface**. The New Network Interface dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **Name** text box, type a name for the network interface.
 - b. In the **IP Address** text box, type an IP address for the load balancer.
 - c. In the **Subnetwork** list, click a subnet.
 - d. In the **ACL** list, click an ACL to map it to the new network interface. Clear the **Failed** check box.
 - e. Select the **Managed** check box if the network interface is intended for management purposes.
4. Click **Save**.
5. If the load balancer has multiple IP addresses, repeat steps 2 through 4 for each of them, so that these IP addresses will not be used for other devices.

Setting up the routing configuration for the load balancer

To set up the load balancer's routing configuration:

1. Click **Data center assets and resources > Inventory > Load Balancers > <Load balancer name>**. Click the **Networking** tab.

Note: You can also access the routing configuration by clicking **Data center assets and resources > Inventory > Load Balancers** to display all the available load balancers. Find the load balancer whose routing configuration you want to set up, and then click its **Edit load balancer** button . Click the **Networking** tab.

2. On the Networking tab, do the following:
 - a. Under **Services Enabled**, select the **router** check box to enable the load balancer to act as router.
 - b. Select the **firewall** check box to enable the router to act as firewall as well.
 - c. Click **Save** to save the routing configuration options.
3. To display further menu options on the screen, click the **More**  button, and then select the **Add Route** menu option. The New Route dialog box is displayed.
4. Under **Destination**, specify the IP address of the destination subnet, and then the gateway for the enabled router. Click **Save**.
5. Repeat steps 3 and 4 if you need to define more routes for the router.
6. If applicable, instead of defining one route after another through steps 3 and 4, you can select the **Apply Routing Table** menu option. The Apply Routing Table dialog box is displayed.
7. In the **Routing Table** list, click the routing table you want to apply to the router. Select the **Remove existing routes** check box if you want to remove all the individual routes that have been previously set up for this device. Click **Save**.

Removing the load balancer

Before you can remove a load balancer, it must be in maintenance. To remove a load balancer:

1. Click **Data center assets and resources > Maintenance > Load Balancers** to list all the load balancers that are currently in maintenance. Find the load balancer you want to remove.

Note: You can also access the load balancer by clicking **Data center assets and resources > Inventory > Load Balancers** to display the load balancers list. In the list, find the offline load balancer to be removed.

2. Click the **Delete load balancer** button .
3. When prompted, click **OK**.

Configuring routers

Configuring routers with Tivoli Intelligent Orchestrator involves the following steps:

- Editing the user-defined variables for the router
- Assigning workflows to the router
- Setting up a management interface for the router
- Setting up the routing configuration for the router

For a step-by-step description of the editing user-defined variables, assigning workflows, setting up a management interface, and setting up the routing configuration procedures, refer to the corresponding subsections in the previous “Configuring switches” or “Configuring load balancers” sections.

Configuring firewalls

Configuring firewalls with the Tivoli Intelligent Orchestrator assumes that you are enabling firewall capabilities at the device level. Configuring firewalls with Tivoli Intelligent Orchestrator involves the following steps:

- Editing the firewall’s user-defined variables
- Assigning workflows to the firewall
- Setting up a management interface for the firewall
- Setting up the routing configuration for the firewall

For step-by-step procedures for assigning workflows, setting up management interfaces, and setting up routing configurations for firewalls, refer to the corresponding subsections in the previous “Configuring switches” section.

Configuring Virtual IPs (VIPs)

This section describes the following procedures:

- Adding a new VIP to a load balancer
- Editing the properties of the VIP
- Assigning workflows to the VIP
- Removing the VIP

A newly added VIP is displayed in the **Data center assets and resources > Inventory > Load Balancers > Virtual IPs** menu. You can edit the VIP properties, assign workflows to it, or remove it. These procedures are described in the following sections.

Adding a new VIP to a load balancer

For information on how to add a VIP to a load balancer, refer to the “Adding a VIP to the load balancer” section, earlier in this chapter.

Editing the properties of the VIP

To edit the properties of the VIP:

1. Click **Data center assets and resources > Inventory > Load Balancers > <Load balancer name> > Virtual IP**. Click the **Modify Virtual IP** button  corresponding to the VIP whose properties you want to edit. The Virtual IP Data dialog box is displayed.

Note: You can also edit the VIP properties by clicking **Data center assets and resources > Inventory > Load Balancers > <Load balancer name> <VIP name>**. Click the **Properties** tab. The Virtual IP Data dialog box is displayed.

2. Edit the VIP properties as required, and then click **Save**.

Adding a real IP to a VIP

To associate the real IP address to the virtual IP address of a particular load balancer:

1. Click **Data center assets and resources > Inventory > Load Balancers > <Load balancer name>**.
2. Select the VIP that you want to modify.
3. Click **Add real IP to virtual IP**.
4. Select the real IP and click **Save**.

Assigning workflows to the VIP

To learn how to assign workflows to a VIP, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Removing the VIP

To remove a VIP:

1. Click **Data center assets and resources > Inventory > Load Balancers > <Load balancer name> > Virtual IP**. Click the **Delete Virtual IP** button  corresponding to the VIP you want to remove.
2. When prompted, click **OK**.

Note: For details on how to remove a VIP from a load balancer by initiating this deployment task from the NIC view, refer to the “Removing a virtual IP” section in the “Managing Tivoli Intelligent Orchestrator” chapter.

Configuring servers

The servers managed by Tivoli Intelligent Orchestrator are either currently functioning in an application cluster or residing in a resource pool awaiting to be allocated to an application cluster. The maintenance area may contain a number of faulty servers, which have been marked by the system as “failed” and temporarily moved to maintenance to be repaired or reconfigured, as well as servers that have not yet been assigned to either an application cluster or a resource pool.

The servers currently running in an application cluster or residing in a resource pool can be found in different states. In the list, the current status of a server is indicated by the color of the icon preceding the server name. A legend, located underneath the server list, explains the meaning of each icon color.

Servers are listed by customer, application, cluster, resource pool, or by resource pool. If a server is listed as *Unknown*, an owner has not yet been defined.

Dedicated

This is a server that is a permanent part of an application cluster.

Assigned

An overflow server is a server residing in a resource pool, awaiting to be allocated to any application cluster. An assigned server is an overflow server that has already been allocated to an application cluster in need. It appears marked as “assigned” both in the resource pool and in the application cluster it was allocated

to.



In transition

This is an overflow server that is being moved from a resource pool to an application cluster, or vice versa. The system marks its transitional status until the moving process is finalized.



Available

This is an overflow server that hasn't yet been allocated to an application cluster. It is available for being allocated to any application cluster in need.



In maintenance

This is a server that has been temporarily put into maintenance, to be repaired or reconfigured. Although it is no longer automatically managed by the Tivoli Intelligent Orchestrator, the server in maintenance is not removed from its cluster.



Failed

This is a faulty server whose malfunction has just been detected by the system, and has been marked as “failed” before being automatically removed from its application cluster.



Unknown

This is a server whose status is temporarily unavailable, due to a system component failure, or no owner is defined.

A new server is added to the list of servers that are currently in maintenance. You can add either a new *dedicated* server or a new *overflow* server.

Note: For more information on how to allocate overflow servers to application clusters, refer to the “Manually adding overflow servers to an application cluster” section in the “Managing Tivoli Intelligent Orchestrator” chapter.

While the new server is in maintenance, you can edit its properties, edit its user-defined variables, set up a NIC, a network interface, and a routing configuration for it, or remove it. You can also access the NIC view for the server, which offers details on the topology of the server and gives further access to other interactive views, configure the server's software stack, and set up credentials for the server.

Adding a new server

To add a new server:

1. Click **Data center assets and resources > Inventory > Servers**.
2. Click **Add server** . The Server Data dialog box is displayed.

3. Based on the type of server you are adding, enter the required data as described below.
 - a. For a new dedicated server:
 1. In the **Name** box, type the name of the new server.
 2. In the **Belongs to** list, click the application cluster this server will be assigned to.
 3. Click **Save**.
 - b. For an overflow server:
 1. In the **Name** box, type the name of the new server.
 2. In the **Belongs to** list, click the resource pool this server will be assigned to.
 3. Click **Save**.

Searching for a server

You can search for a server or servers by name, cluster, or pool. To do this:

1. Click **Data center assets and resources > Inventory > Servers**.
2. On the Server Inventory page, enter the information to use for your search.
3. Click **Search**.

Editing the server properties

To edit the properties of the server:

1. Click **Data center assets and resources > Inventory > Servers** to display all the available servers. Find the server whose properties you want to edit, and then click its **More** button. From the displayed pop-up menu, select **Edit**. The Server Data dialog box is displayed.

Note: You can also edit the server properties by clicking **Data center assets and resources > Inventory > Servers > <Server name>**. On the **Info** tab, click **Edit > Properties**. The Server Data dialog box is displayed.

2. Edit the server properties as required.
3. Click **Save**.

Editing the user-defined variables for the server

To edit the user-defined variables for a server:

1. Click **Data center assets and resources > Maintenance > Servers** and select the server whose user-defined variables you want to edit. Click the **Variables** tab.

Note: You can also edit the user-defined variables by clicking **Data center assets and resources > Inventory > Servers** to list all the available servers. Find the server whose properties you want to edit, and then click its **More** button. From the displayed pop-up menu, select **Edit**. Click the **Variables** tab.

2. Complete the fields as follows:
 - a. In the **Key** box, type the name of the SNMP community that specifies the access rights for the server:
snmp.community.read to query the server, or
snmp.community.write to change it.
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.

3. Click **Add**.

Assigning workflows to the server

To learn how to assign workflows to a server, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Setting up a NIC for the server

A server must be in maintenance before you can set up a network interface card (NIC). To set up a NIC for the server:

1. Click **Data center assets and resources > Maintenance > Servers** and select the server for which you want to set up the NIC. On the **Info** tab, click **Edit > Add NIC**. The New Network Interface Card dialog box is displayed.

Note: You can also set up the NIC by clicking **Data center assets and resources > Inventory > Servers** to list all the available servers. Find the offline server for which you want to set up the NIC, and then click its **Edit server** button . On the **Info** tab, click **Edit > Add NIC** . The New Network Interface Card dialog box is displayed.

2. Complete the fields as follows:
 - a. In the **MAC** box, type the MAC address of the NIC.
 - b. In the **Switch** list, click a switch.
 - c. In the **Port** box, type the physical port connected to the NIC. If required, specify the module for the port as well.
 - d. In the **VLAN** list, click a VLAN.
 - e. Select the **Managed** check box if the NIC is intended for management purposes. Clear the **Failed** check box.
 - f. Select the **Netboot** check box if you want the NIC to be boot-enabled.
3. Click **Add**.

Accessing the NIC view for the server

You can quickly access the NIC view for a server in a few different ways, as described in the following sections:

- Accessing the **View > Icon View**  menu on the Server Info screen, or
- From the Devices Legend window, or
- From any other interactive view by clicking the NIC icon

To access the NIC view for a server from the **View > Icon View**  menu:

1. Click **Data center assets and resources > Inventory > Servers** and click the server you are interested in. The **Info** tab is displayed.

Or:

In the **Customer applications** navigation tree, select the server you are interested in. The **Info** tab is displayed.

2. Click **View > Icon View**  menu. The “Overview: <Server name>” screen is displayed.

Accessing other interactive views from the NIC view

From the NIC view, you can quickly access the following views:

- *Server Info view*, by clicking the **Info** tab
- *Switch Fabric view*, by clicking the switch fabric icon
- *VLAN view* by clicking the VLAN icon in the VLAN Legend window

Note: For step-by-step procedures on how to perform various deployment tasks from the NIC view, refer to the “Performing deployment tasks from the interactive views” section, described in the “Managing Tivoli Intelligent Orchestrator” chapter.

Setting up a network interface for the server

To set up a network interface for the server:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**.
2. On the **Info** tab, identify the NIC to which you want to add a new interface, and then click the **More**  button.
3. From the pop-up menu, select **Add Interface** . The New Network Interface dialog box is displayed.
4. Complete the fields as follows:
 - a. In the **Name** box, type a name for the network interface.
 - b. In the **IP** box, type the IP address for the server.
 - c. Select a subnet, and then choose an ACL for the network interface.
 - d. Clear the **Failed** check box if it is selected.
 - e. Select the **Managed** check box if the network interface is intended for management purposes.
 - f. Type a default gateway IP address for the network interface.
 - g. Select **Perform using dcm.interaction** to invoke the IPSystem.Add IP Address logical operation. This adds the specified IP address (network interface) to the specified IP network device.
5. Click **Save**.

Setting up the routing configuration for the network interface

A server with multiple NICs can act as a router. To set up the routing configuration of its network interfaces:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**. The **Info** tab is displayed.
2. In the list, select the network interface that requires a new routing configuration and click the **More**  button.
3. Select the **Add Route**  menu option. The New Route dialog box is displayed.
4. Select the IP address of the destination subnet from the **Destination** drop down list, and specify the gateway for the enabled router.
5. Click **Save**.
6. Repeat steps 2 through 4 to define more routes for the router.

If applicable, instead of defining one route after another in steps 2 through 4, you can select the **Apply Routing Table** menu option. The Apply Routing Table dialog box is displayed.

In the **Routing Table** list, click the routing table you want to apply to the router. Select the **Remove existing routes** check box if you want to remove all the individual routes that have been previously set up for this device. Click **Save**.

Configuring the software stack

For each server in the data center, a **Software Stack** tab lists all the software products that should be installed on that server. You can interactively add and remove software by automatically invoking the appropriate workflows, as defined by the device driver.

Note: For additional information on how to configure the software products, refer to the “Configuring software products” section, later in this chapter.

Assigning a software stack to a server

To assign a software stack to a server:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**. Click the **Software Stack** tab.

Note: You can also access the software stack by clicking **Customer applications**, and drilling down to the server you are interested in. Click the **Software Stack** tab.

2. In the **Software Stack Configuration** list, click the stack you want to assign to the server. It specifies what should be installed on that server. Click **Save**.

After you select one of the stacks in the list, its content is displayed on the screen without allowing you to add other entries to the stack list.

Adding a new entry to the server's stack

A new entry augmenting the server's stack can be either a software product, a software patch, or another nested stack. Augmenting the server's stack can be done only in **Automatic** mode, when the stack inherited by the server is managed by workflows, and further entries can be manually added to the stack list.

To manually add a new entry to the server's stack:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**. Click the **Software Stack** tab.

Note: You can also access the software stack by clicking **Customer applications**, and drilling down to the server you are interested in. Click the **Software Stack** tab.

2. In the **Software Stack Configuration** list, click **Automatic**, and then click **Save**.
3. Under **Software Stack Entries**, select one of the following options:
 - Select **Install Software**, and then use the corresponding list to select the software product you want to add to the server's stack. If the product you want to install requires licensing, click the corresponding license pool in the **License Pool** list
 - Select **Install Patch**, and then select the software patch you want to add to the server's stack

- Select **Install Stack**, and then select the nested stack you want to add to the server's stack.
4. Click **Install**. The selected product, patch or stack is added to the server's stack.
 5. Repeat step 3 and 4 to add more entries to the server's stack. The order the entries are listed in the stack is important, indicating the order in which the products will be installed on the server. If necessary, use the **Move up** and **Move down** buttons to rearrange the entries in the list.

Editing the configuration of the server's stack entries

To edit the configuration of the server's stack entries:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**. Click the **Software Stack** tab.

Note: You can also access the stack entries by clicking **Customer applications**, and locate the server you are interested in. Click the **Software Stack** tab.

2. In the **Software Stack Entries** list, identify the entry whose configuration you want to edit, and then click its **Edit product/stack configuration**  button.
 - If the entry is a software product, you are sent to the **Properties** tab for that software product.
 - If the entry is a software stack, you are sent to the **Properties** tab for that software stack.
3. Edit the properties of the software product or stack as required, and then click **Save**.

Removing an entry from the server's stack

To remove an entry from the server's stack:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**. Click the **Software Stack** tab.

Note: You can also access the software stack tab by clicking **Customer applications**, and selecting the server you are interested in. Click the **Software Stack** tab.

2. In the **Software Stack Entries** list, identify the entry that you want to remove, and then click its **Remove Software Product**  button.
3. When prompted, click **OK**.

Installing a software product on the server

For each server in the data center, a **Software** tab lists all the software products installed on it. You can interactively add software products by automatically invoking the appropriate workflows, as defined by the device driver.

To install a software product on the server:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**.
2. Click the **Software** tab. The Server: <Server name> screen displays all the software products installed or available to be installed on the server.
3. In the **Install Software** list, click the product that you want to install on the server.
4. Click **Install**. The workflow triggered by the *Software.Install* logical device operation is executed. If the execution is successful, the software product is installed on the server.

Note: Tivoli Intelligent Orchestrator does not manage software product licenses or patches which were installed as a single software product (or single software patch). To have Tivoli Intelligent Orchestrator manage software product licenses or patches, install software products and patches as a part of software stack. For information on how to do this, see “Adding a new entry to the server’s stack” on page 40.

Uninstalling a software product from the server

You can interactively remove software products by automatically invoking the appropriate workflows, as defined by the device driver.

To uninstall a software product on the server:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**. Click the **Software** tab.
2. In the **Software** list, identify the software product you want to uninstall from the server.
3. Click the **Delete Software**  button. The workflow triggered by the *Software.Uninstall* logical device operation is executed. If the execution is successful, the software product is uninstalled from the server.

Refreshing the status of the software product

To refresh the status of a software product installed on the server:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**. Click the **Software** tab.
2. In the **Software** list, identify the software product on which you want to obtain updated data.
3. Click the **Refresh Software Status** button .

Configuring credentials for the server

Tivoli Intelligent Orchestrator provides a way to streamline the access of managed servers to various services using credentials. For a more efficient use, credentials obtained from clients are stored in a credential repository. Services are provided by host servers through service access points (SAPs). A service access point is a channel that is used by a client server to access a service.

A number of protocols are available to be used for accessing services, such as IPv4, IPv6, IPX, fiber-channel, and so on. For the IPv4 protocol type, a number of specific protocol subtypes are available, as follows:

- FTP, the standard File Transfer Protocol
- ICMP, for Ping
- SCP, for secure remote transfer of a file from one server to another
- SNMP, for queries and configurations
- SSH, for secure remote access
- Telnet, for remote access

The following procedures are described:

- Adding a new service access point
- Editing the properties of the service access point

- Assigning workflows to a service access point
- Setting up password credentials for the SAP
- Setting up RSA credentials for the SAP
- Setting up RSA credentials for the SAP
- Setting up SNMP credentials for the SAP
- Removing the service access point

Adding a new service access point

To add a new service access point (SAP) to the server:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**.
2. Click the **Credentials** tab.
3. On the **Credentials** tab, click **Edit > Add Access Point**. The New Access Point dialog box is displayed.
4. Complete the fields as follows:
 - a. In the **Type** list, click the type of protocol you will be using to connect to that service.
 - b. Select the **Host** check box if you want this server to be the host providing the service through SAPs. Optionally, you can enter the SAP's domain and context. The context field can be used to define a unique identifier for this SAP.
 - c. Type the port number to use for this service access point. Based on the selected type of protocol, the default port for the SAP is displayed. That port number should be set to **0** for a client SAP (**Host** option not selected).
 - d. Select the **Authentication** check box if the access to that service will require authentication. If authentication is required, specific credentials must be defined, corresponding to the chosen type of protocol.
 - e. Select **Perform using dcm.interaction** to invoke the Device.Create Service Access Point logical operation.
5. Click **Save**.

Editing the properties of the service access point

To edit the properties of the service access point (SAP):

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**. Click the **Credentials** tab. In the list, identify the SAP whose properties you want to edit, click its **More**  button, and then select  **Edit**.

Note: You can also access the SAP properties by clicking **Data center assets and resources > Inventory > Servers > <Server name>**. Click the **Credentials** tab. In the list, click the name of the SAP whose properties you want to edit. On the **Info** tab, click **Edit > Properties**.

The properties dialog box is displayed.

2. Edit the properties of the SAP as required.
3. Click **Save**.

Assigning workflows to a service access point

A number of predefined workflows triggered by logical device operations are available to be assigned to service access points (SAPs). In order for a workflow to be executed in the specified SAP, you need to perform the following procedures:

- Tagging the workflow with a logical device operation
- Assigning workflows to the service access point

These procedures are described in the following subsections.

Tagging the workflow with a logical device operation

To learn how to tag a workflow with a logical operation, and how to define a new logical operation if required, refer to the “Tagging the workflow with a logical device operation” section in the “Configuring and using workflows” chapter.

Assigning workflows to the service access point

To assign workflows to a service access point (SAP):

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**.
2. Click the **Credentials** tab.
3. In the list, click the name of the SAP which you want to assign the workflows to.
4. Click the **Workflows** tab.
5. In the **Logical Operations** list, click the logical operation for the workflow you are interested in.
6. In the **Workflow** list underneath, click the workflow to be assigned to the specified SAP.
7. Click **Add Selected Workflow**. The workflow is added to the list of all workflows that are currently assigned to that SAP.
8. If necessary, repeat steps 5 through 7 to assign other workflows to this SAP.

Note: Some device drivers are pre-defined for SAPs that are running certain application protocols. For example, SNMP v1 and SSH both have pre-defined drivers.

Setting up password credentials for the SAP

If the type of protocol you will be using to access the service is Telnet or FTP, password credentials for the service access point (SAP) may be required. To set up the password credentials for the SAP:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**.
2. Click the **Credentials** tab.
3. In the list, identify the SAP for which you want to set up the password credentials, click its **More**  button.
4. Select the  **Add Password Credentials** option. The New Password Credentials dialog box is displayed.
5. In the **Search Key** text box, type the alias that uniquely identifies the user name-password pair you will be using to access the service. Optionally, you can use the **Enable Password** text box to type the specific password that enables you to perform special administrator operations on the server.
6. Type the key fingerprint for the server.
7. Click **Save**.

Setting up RSA credentials for the SAP

If the type of protocol you will be using to access the service is SSH or SCP, RSA credentials for the service access point (SAP) will be required. To set up the RSA credentials for the SAP:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**.
2. Click the **Credentials** tab.
3. In the list, identify the SAP to set up the RSA credentials for and click its **More**  button.
4. Select the  **Add RSA Credentials** option. The New RSA Credentials dialog box is displayed.
5. Complete the fields as follows:
 - a. In the **Search Key** text box, type the name you will be using to access the service.
 - b. Type the user name.
 - c. Type the passphrase.
 - d. Specify the public and private keys for the server.
6. Click **Save**.

Setting up SNMP credentials for the SAP

If the type of protocol you will be using to access the service is SNMP, SNMP credentials for the server will be required. To set up the SNMP credentials for the SAP:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**.
2. Click the **Credentials** tab.
3. In the list, identify the SAP for which you want to set up the SNMP credentials and click its **More**  button.
4. Select the  **Add SNMP Credentials** option. The New SNMP Credentials dialog box is displayed.
5. Type the search key you will be using to access the service, and then type the community name for the server.

Removing the service access point

To remove a service access point (SAP):

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**.
2. Click the **Credentials** tab.
3. In the list, identify the service access point to be removed and click its **More**  button.
4. Select the  **Delete** option.
5. When prompted, click **OK**.

Removing the server

Only a server in maintenance can be removed. To remove the server:

1. Click **Data center assets and resources > Maintenance > Servers** to list all the servers that are currently in maintenance. Find the server you want to remove, and then click its **Delete server** button .

Note: You can also access the server by clicking **Data center assets and resources > Inventory > Servers** to list all the available servers. In the list, find the server you

want to remove, and then click its **Delete server** button .

2. When prompted, click **OK**.

Configuring boot servers

This section describes the following procedures:

- Adding a new boot server
- Editing the boot server properties
- Editing the user-defined variables for the boot server
- Assigning workflows to the boot server
- Setting up a NIC for the boot server
- Adding an interface card to the boot server
- Setting up a network interface for the boot server
- Setting up the routing configuration for the network interface
- Configuring credentials for the boot server
- Removing the boot server

These procedures are described in the following sections.

Adding a new boot server

To add a new boot server:

1. Click **Data center assets and resources > Inventory > Boot Servers**. The Boot Server Inventory screen lists all the boot servers currently available in the system.
2. Click the  **Add Boot Server** button. The Boot Server Configuration dialog box is displayed.
3. Type the name of the new boot server.
4. Click **Save**.

You have finished adding a new boot server. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the boot server properties

To edit the boot server's properties:

1. Access the Boot Server Properties tab in one of the following ways:
 - Click **Data center assets and resources > Inventory > Boot Servers**. The Boot Server Inventory screen lists the available boot servers. In the list, identify the boot server whose properties you want to edit, and then click its **Edit Boot Server**  button. The **Properties** tab is displayed.
 - Click **Data center assets and resources > Inventory > Boot Servers > <Boot server name>**. On the **Info** tab, click **Edit -> Properties** . The **Properties** tab is displayed.
 - Click **Data center assets and resources -> Inventory -> Boot Servers -> <Boot server name>**. Click the **Properties** tab.
2. Edit the name of the boot server as required, and then click **Save**.

You have finished editing the boot server's properties. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the user-defined variables for the boot server

To edit the user-defined variables for the boot server:

1. Click **Data center assets and resources > Inventory > Boot Servers > <Boot server name>**.
2. Click the **Variables** tab.
3. Complete the fields as follows:
 - a. In the **Key** text box, type the name of the SNMP community that specifies the access rights for the boot server:
snmp.community.read to query the server
snmp.community.write to change it.
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
4. Click **Add**.

Assigning workflows to the boot server

To learn how to assign workflows to the boot server, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Setting up a NIC for the boot server

To set up a NIC for the boot server:

1. Click **Data center assets and resources > Inventory > Boot Servers > <Boot server name>**.
2. On the **Info** tab, click **Edit -> Add NIC** . The New Network Interface Card dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **MAC** text box, type the MAC address of the NIC.
 - b. In the **Switch** list, click a switch. In the **Port** box, type the physical port connected to the NIC. If required, specify the module for the port as well.
 - c. In the **VLAN** list, click a VLAN.
 - d. Clear the **Failed** check box.
 - e. Select the **Managed** check box if the NIC is intended for management purposes.
 - f. Select the **Netboot Enabled** check box if you want the NIC to be boot-enabled.
4. Click **Save**.

Adding an interface card to the boot server

To add a serial interface card to the boot server:

1. Click **Data center assets and resources > Inventory > Boot Servers > <Boot server name>**.
2. On the **Info** tab, click **Edit -> Add serial IC** . The **New Interface Card** dialog box is displayed.
3. Type the name of the new interface card.
4. Click **Save**.

Setting up a network interface for the boot server

To set up a network interface for the boot server:

1. Click **Data center assets and resources > Inventory > Boot Servers > <Boot server name>**.
2. On the **Info** tab, click **Edit > Add Interface** . The **New Network Interface** dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **Name** text box, type a name for the network interface.
 - b. In the **IP** text box, type the IP address for the server.
 - c. In the corresponding lists, click a subnetwork and an ACL for the boot server.
 - d. Select the **Managed** check box if the network interface is intended for management purposes.
 - e. Clear the **Failed** check box.
4. Click **Save**.

Setting up the routing configuration for the network interface

A boot server with multiple NICs can act as a router. To set up the routing configuration of its network interfaces:

1. Click **Data center assets and resources > Inventory > Boot Servers > <Boot server name>**. The **Info** tab is displayed.
2. In the list, identify the network interface for which you want to set up the routing configuration and click its **More**  button.
3. Select the **Add Route**  menu option. The **New Route** dialog box is displayed.
4. Under **Destination**, specify the IP address of the destination subnet, and then the gateway for the enabled router.
5. Click **Save**.
6. Repeat steps 2 through 5 to define more routes for the router.
7. If applicable, instead of defining one route after another, you can select the **Apply Routing Table** menu option. The **Apply Routing Table** dialog box is displayed.
8. In the **Routing Table** list, do the following:
 - a. Click the routing table you want to apply to the router.
 - b. Select the **Remove existing routes** check box if you want to remove all the individual routes that have been previously set up for this device.
 - c. Click **Save**.

Configuring credentials for the boot server

For information on how to define the service access points and configure credentials for the boot server, refer to the “Configuring credentials for the server” section, earlier in this chapter.

Removing the boot server

To remove the boot server:

1. Click **Data center assets and resources > Inventory > Boot Servers** to list all the available boot servers. In the list, find the boot server you want to remove, and then click its **Delete Boot Server** button .
2. When prompted, click **OK**.

Configuring blade servers

This section describes the following procedures:

- Adding a new blade server chassis
- Editing the properties of the blade server chassis
- Editing the user-defined variables for the blade server chassis
- Assigning workflows to the blade server chassis
- Adding a new blade server to the blade server chassis
- Editing the properties of the blade server
- Setting up a NIC for the blade server chassis
- Setting up a network interface for the blade server chassis
- Setting up the routing configuration for the network interface
- Removing a blade server from the blade server chassis
- Removing the blade server chassis

Adding a new blade server chassis

To add a new blade server chassis:

1. Click **Data center assets and resources > Inventory > Blade Servers**. The **Blade Server Chassis** tab lists all of the blade server chassis currently available in the system.
2. Click the **Add Blade Server Chassis**  button. The Blade Chassis Configuration dialog box is displayed.
3. Type the name of the new blade chassis, and then click **Save**.

You have finished adding the new blade server chassis. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the properties of the blade server chassis

To rename the blade server chassis:

1. Click **Data center assets and resources > Inventory > Blade Servers**. The **Blade Server Chassis** tab lists the available blade server chassis. In the list, identify the blade chassis admin whose properties you want to edit, and then click its **Edit blade admin**  button. The **Properties** tab is displayed.
2. Edit the name of the blade chassis admin as required, and then click **Save**.

You have finished editing the blade server chassis properties. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the user-defined variables for the blade server chassis

To edit the user-defined variables for the blade server chassis:

1. Click **Data center assets and resources > Inventory > Blade Servers > <Blade server chassis name>**.
2. Click the **Variables** tab.
3. Complete the fields as follows:
 - a. In the **Key** text box, type the name of the SNMP community that specifies the access rights for the blade server chassis:
snmp.community.read to query the server
snmp.community.write to change it.
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
4. Click **Add**.

Assigning workflows to the blade server chassis

To learn how to assign workflows to the blade server chassis, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, later in this guide.

Adding a new blade server to the blade server chassis

To add a new blade server to a blade server chassis:

1. Click **Data center assets and resources > Inventory > Blade Servers > <Blade Server Chassis name>**.
2. On the **Info** tab, click **Edit > Add Blade Server** . The Server Data dialog box is displayed.
3. Complete the fields as follows:
 - a. Type the name of the new blade server.
 - b. In the **Belongs to** list, click the cluster or resource pool the blade server will belong to.
 - c. Specify the blade slot for the new server.
 - d. Under **SNMP community names**, specify the access rights for the new server:
Read to query the server
Write to change the server.
4. Click **Save**.

You have finished adding the new blade server to the current blade admin chassis. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the properties of the blade server

A blade server must be in maintenance before you can edit its properties. To manually send a blade server to maintenance, identify it first in the list of blade servers available for a certain blade chassis admin, click its **More**  button, and then select the **To maintenance** menu option.

To edit the properties of the blade server:

1. Click **Data center assets and resources > Inventory > Blade Servers > <Blade Server Chassis name>**.
2. The **Info** tab lists all the blade servers available in the current blade chassis admin. Identify the offline blade server whose properties you want to edit, click its **More**  button, and then select the **Edit**  option. The Server Data dialog box is displayed.
3. Edit the blade server properties as required.
4. Click **Save**.

Setting up a NIC for the blade server chassis

To set up a NIC for the blade server chassis:

1. Click **Data center assets and resources > Inventory > Blade Servers > <Blade Server Chassis name>**.
2. On the **Info** tab, click **Edit > Add NIC** . The New Network Interface Card dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **MAC** text box, type the MAC address of the NIC.
 - b. In the **Switch** list, click a switch.
 - c. In the **Port** box, type the number of the physical port connected to the NIC. If required, specify the module for the port as well.
 - d. In the **VLAN** list, click a VLAN.
 - e. Select the **Managed** check box if the NIC is intended for management purposes.
 - f. Clear the **Failed** check box.
 - g. Select the **Netboot Enabled** check box if you want the NIC to be boot-enabled.
4. Click **Save**.

Setting up a network interface for the blade server chassis

To set up a network interface for the blade server chassis:

1. Click **Data center assets and resources > Inventory > Blade Servers > <Blade Server Chassis name>**.
2. On the **Info** tab, click **Edit -> Add Interface** . The New Network Interface dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **Name** text box, type a name for the new network interface.
 - b. In the **IP Address** box, type the IP address for the blade server chassis.
 - c. In the corresponding lists, click a subnet and an ACL for the blade server chassis.
 - d. Select the **Managed** check box if the network interface is intended for management purposes.
 - e. Clear the **Failed** check box.
4. Click **Save**.

Setting up the routing configuration for the network interface

A blade server chassis with multiple NICs can act as a router. To set up the routing configuration of its network interfaces:

1. Click **Data center assets and resources > Inventory > Blade Servers > <Blade Server Chassis name>**. The **Info** tab is displayed.
2. In the list, identify the network interface for which you want to set up the routing configuration, and click its **More** button.
3. Select the **Add Route** menu option. The New Route dialog box is displayed.
4. Under **Destination**, specify the IP address of the destination subnet, and then the gateway for the enabled router. Click **Save**.
5. Repeat steps 2 through 4 to define more routes for the router.
6. If applicable, instead of defining one route after another by following steps 2 through 4, you can select the **Apply Routing Table** menu option. The Apply Routing Table dialog box is displayed.
7. In the **Routing Table** list, click the routing table you want to apply to the router. Select the **Remove existing routes** check box if you want to remove all the individual routes that have been previously set up for this device. Click **Save**.

Removing a blade server from the blade server chassis

You can only remove an offline blade server from its chassis. To manually send a blade server to maintenance, identify it first in the list of blade servers available for a specific blade chassis admin, click its **More** button, and then select the **To maintenance** menu option.

To remove the offline blade server from its chassis:

1. Click **Data center assets and resources > Inventory > Blade Servers > <Blade Server Chassis name>**. The **Info** tab lists all the blade servers available in the current blade chassis.
2. Identify the offline blade server you want to remove, click its **More** button, and then select the **Delete** option.
3. When prompted, click **OK**.

Removing the blade server chassis

To remove the blade server chassis:

1. Click **Data center assets and resources > Inventory > Blade Servers**. The **Blade Server Chassis** tab lists all the currently available blade chassis.
2. In the list, identify the blade chassis you want to remove, and then click its **Delete Blade Server Chassis** button.
3. When prompted, click **OK**.

Configuring terminal servers

Terminal servers provide network access to servers with network connections that have been lost or to servers that have no built-in network support. The terminal servers can also provide direct communication to hosts that support TCP/IP protocols in their system operating software. Terminal servers handle multiple device interfaces. They multiplex asynchronous RS-232 serial lines onto a high-speed network interface. The network interface is typically to Ethernet.

This section describes the following procedures:

- Adding a new terminal server
- Editing the properties of the terminal server
- Editing the user-defined variables for the terminal server
- Assigning workflows to the terminal server
- Setting up a NIC for the terminal server
- Adding an interface card to the terminal server
- Setting up a network interface for the terminal server
- Setting up the routing configuration for the network interface
- Configuring credentials for the terminal server
- Removing the terminal server

Adding a new terminal server

To add a new terminal server:

1. Click **Data center assets and resources > Inventory > Terminal Servers**. The Terminal Server Inventory screen lists all the currently available terminal servers.
2. Click the **Add Terminal Server**  button. The Terminal Server Configuration dialog box is displayed.
3. Type a name for the new terminal server, and then click **Save**.

Editing the properties of the terminal server

To edit the terminal server properties:

1. Click **Data center assets and resources > Inventory > Terminal Servers**. The Terminal Server Inventory screen lists all the currently available terminal servers.
2. In the list, identify the terminal server whose properties you want to edit, and then click its **Edit terminal server**  button. The **Properties** tab is displayed.
3. Edit the terminal server's name as required, and then click **Save**.

Editing the user-defined variables for the terminal server

To edit the user-defined variables of a terminal server:

1. Click **Data center assets and resources > Inventory > Terminal Servers > <Terminal server name>**.
2. Click the **Variables** tab.
3. Complete the fields as follows:
 - a. In the **Key** box, type the name of the SNMP community that specifies the access rights for the terminal server:
 - snmp.community.read** to query the server
 - snmp.community.write** to change it.
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
4. Click **Add**.

Assigning workflows to the terminal server

To learn how to assign workflows to a terminal server, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Setting up a NIC for the terminal server

To set up a NIC for the terminal server:

1. Click **Data center assets and resources > Inventory > Terminal Servers > <Terminal server name>**.
2. On the **Info** tab, click **Edit > Add NIC** . The New Network Interface Card dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **MAC** text box, type the MAC address of the NIC.
 - b. In the **Switch** list, click a switch. In the **Port** box, type the number of the physical port connected to the NIC.
 - c. If required, specify the module for the port as well. In the **VLAN** list, click a VLAN.
 - d. Select the **Managed** check box if the NIC is intended for management purposes.
 - e. Clear the **Failed** check box.
 - f. Select the **Netboot Enabled** check box if you want the NIC to be boot-enabled.
4. Click **Save**.

Adding an interface card to the terminal server

To add a new interface card to the terminal server:

1. Click **Data center assets and resources > Inventory > Terminal Servers > <Terminal server name>**.
2. On the **Info** tab, click **Edit > Add serial IC** . The New Interface Card dialog box is displayed.
3. Type the name of the new interface card.
4. Click **Save**.

Setting up a network interface for the terminal server

To set up a network interface for the terminal server:

1. Click **Data center assets and resources > Inventory > Terminal Servers > <Terminal server name>**.
2. On the **Info** tab, click **Edit > Add Interface** . The New Network Interface dialog box is displayed.
3. Complete the fields as follows:
 - a. Type a name for the new network interface.
 - b. In the **IP Address** text box, type the IP address for the terminal server.
 - c. In the corresponding lists, click a subnet and an ACL for the terminal server.
 - d. Select the **Managed** check box if the network interface is intended for management purposes.
 - e. Clear the **Failed** check box.
4. Click **Add**.

Setting up the routing configuration for the network interface

A terminal server with multiple NICs can act as a router. To set up the routing configuration for its network interfaces:

1. Click **Data center assets and resources > Inventory > Terminal Servers > <Terminal server name>**. The **Info** tab is displayed.
2. In the list, identify the network interface for which you want to set up the routing configuration, and then click its **More** button.
3. Select the **Add Route** menu option. The New Route dialog box is displayed.
4. Under **Destination**, specify the IP address of the destination subnet, and then the gateway for the enabled router. Click **Save**.
5. Repeat steps 2 through 4 to define more routes for the router.
6. If applicable, instead of defining one route after another through steps 2 through 4, you can select the **Apply Routing Table** menu option. The Apply Routing Table dialog box is displayed.
 - a. Click the routing table you want to apply to the router.
 - b. Select the **Remove existing routes** check box if you want to remove all the individual routes that have been previously set up for this device.
 - c. Click **Save**.

Configuring credentials for the terminal server

For information on how to define the service access points and configure credentials for the terminal server, refer to the “Configuring credentials for the server” section, earlier in this chapter.

Removing the terminal server

To remove the terminal server:

1. Click **Data center assets and resources > Inventory > Terminal Servers**. The Terminal Server Inventory screen lists all the currently available terminal servers.
2. In the list, identify the terminal server you want to remove, and then click its **Delete terminal server** button.
3. When prompted, click **OK**.

Adding a console server network interface

To add a console server network interface:

1. Click **Data center assets and resources > Inventory > Terminal Servers**.
2. Click **Console Server**.
3. Type the port number for the new console server.
4. Select the port type from the drop-down list.
5. Type the speed of the interface and port.
6. Select the service access point for the console server. A service access point is an abstract definition of a channel that is used by a client server to access a service.
7. Select a connection.
8. Click **Add**.

To edit the properties of a console server network interface, click Edit

Configuring ACLs

This section describes the following procedures:

- Adding a new ACL
- Editing the properties of the ACL
- Setting up access rules for the ACL
- Editing the user-defined variables for the ACL
- Assigning workflows to the ACL

Adding a new ACL

To add a new ACL:

1. Click **Data center assets and resources > Inventory > ACLs**. The Inventory ACLs screen lists all the currently available ACLs.
2. Click **Add ACL** . The ACL Configuration dialog box is displayed.
3. Type the name of the new ACL, and then click **Save**.

Editing the properties of the ACL

To edit the ACL's properties:

1. Click **Data center assets and resources > Inventory > ACLs**. The Inventory ACLs screen lists all the currently available ACLs.
2. In the list, identify the ACL whose properties you want to edit, and then click its **Modify ACL**  button. The ACL Configuration dialog box is displayed.
3. Edit the ACL name as required, and then click **Save**.

Setting up access rules for the ACL

This section describes the following procedures:

- Adding a new access rule
- Editing an existing access rule
- Removing the access rule

Adding a new access rule

To add a new access rule:

1. Click **Data center assets and resources > Inventory > ACLs > <ACL name>**. The **Access Rules** tab is displayed.
2. Type the required configuration data for the new rule, as follows:
 - a. Under **Target**, enter the type of access, **PERMIT** or **DENY**.
 - b. Specify the protocol for the access rule.
 - c. Select the **Source** subnet and specify the corresponding port range, and then select the **Destination** subnet and specify the corresponding port range for the new access rule.
 - d. If required, type additional options, such as port number, log, and so on.
 - e. Click **Add**.

Editing an existing access rule

To edit an existing access rule:

1. Click **Data center assets and resources > Inventory > ACLs > <ACL name>**. The **Access Rules** tab displays all the access rules available for the selected ACL.
2. In the list, identify the access rule you want to modify, and then click its **Edit rule**  button.
3. Edit the properties of the rule as required.
4. Click **Save**.

Removing the access rule

To remove the access rule:

1. Click **Data center assets and resources > Inventory > ACLs > <ACL name>**. The **Access Rules** tab displays all the access rules available for the selected ACL.
2. In the list, identify the access rule you want to remove, and then click its **Delete rule**  button.
3. When prompted, click **OK**.

Editing the user-defined variables for the ACL

To edit the user-defined variables of the ACL:

1. Click **Data center assets and resources > Inventory > ACLs > <ACL name>**.
2. Click the **Variables** tab.
3. Complete the fields as follows:
 - a. In the **Key** text box, type the name of the SNMP community that specifies the access rights for the ACL:
snmp.community.read
snmp.community.write
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
4. Click **Add**.

Assigning workflows to the ACL

To learn how to assign workflows to an ACL, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Configuring power units

This section describes the following procedures:

- Adding a new power unit
- Editing the properties of the power unit
- Editing the user-defined variables for the power unit
- Assigning workflows to the power unit
- Setting up a NIC for the power unit
- Adding an interface card to the power unit
- Setting up a network interface for the power unit
- Setting up the routing configuration for the network interface
- Configuring credentials for the power unit
- Setting up the power unit configuration

- Removing the power unit

Adding a new power unit

To add a new power unit:

1. Click **Data center assets and resources > Inventory > Power Units**. The Power Unit Inventory screen lists all the currently available power units.
2. Click the **Add power unit**  button. The Power Unit Data dialog box is displayed.
3. Type the name of the new power unit.
4. Click **Save**.

You have finished adding a new power unit. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the properties of the power unit

To edit the properties of the power unit:

1. Click **Data center assets and resources > Inventory > Power Units**. The Power Unit Inventory screen lists all the currently available power units.
2. In the list, identify the power unit whose properties you want to edit, and then click its **Edit power unit**  button. The **Properties** tab is displayed.
3. Edit the name of the power unit as required.
4. Click **Save**.

You have finished editing the power unit properties. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the user-defined variables for the power unit

To edit the user-defined variables for the power unit:

1. Click **Data center assets and resources > Inventory > Power Units > <Power unit name>**.
2. Click the **Variables** tab.
3. Complete the fields as follows:
 - a. In the **Key** text box, type the name of the SNMP community that specifies the access rights for the power unit:
snmp.community.read to query it
snmp.community.write to change it.
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
4. Click **Add**.

Assigning workflows to the power unit

To learn how to assign workflows to the power unit, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Setting up a NIC for the power unit

To set up a NIC for the power unit:

1. Click **Data center assets and resources > Inventory > Power Units -> <Power unit name>**.
2. On the **Info** tab, click **Edit > Add NIC** . The New Network Interface Card dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **MAC** text box, type the MAC address of the NIC.
 - b. In the **Switch** list, click a switch.
 - c. In the **Port** box, type the number of the physical port connected to the NIC.
 - d. If required, specify the module for the port as well.
 - e. In the **VLAN** list, select a VLAN.
 - f. Select the **Managed** check box if the NIC is intended for management purposes.
 - g. Clear the **Failed** check box.
 - h. Select the **Netboot Enabled** check box if you want the NIC to be boot-enabled.
4. Click **Save**.

Adding an interface card to the power unit

To add a new interface card to the power unit:

1. Click **Data center assets and resources > Inventory > Power Units > <Power unit name>**. On the **Info** tab, click **Edit > Add serial IC** . The New Interface Card dialog box is displayed.
2. Type the name of the new interface card.
3. Click **Save**.

Setting up a network interface for the power unit

To set up a network interface for the power unit:

1. Click **Data center assets and resources > Inventory > Power Units > <Power unit name>**.
2. On the **Info** tab, click **Edit > Add Interface** . The New Network Interface dialog box is displayed.
3. Complete the fields as follows:
 - a. Type a name for the new network interface.
 - b. In the **IP Address** text box, type the IP address for the power unit.
 - c. In the corresponding lists, click a subnet and an ACL for the power unit.
 - d. Select the **Managed** check box if the network interface is intended for management purposes.
 - e. Clear the **Failed** check box.
4. Click **Add**.

Setting up the routing configuration for the network interface

To set up the routing configuration for the network interfaces of the power unit:

1. Click **Data center assets and resources > Inventory > Power Units > <Power unit name>**. The **Info** tab is displayed.
2. In the list, identify the network interface for which you want to set up the routing configuration, and then click its **More** button.
3. Select the **Add Route** menu option. The New Route dialog box is displayed.
4. Under **Destination**, specify the IP address of the destination subnet, and then the gateway for the enabled router. Click **Save**.
5. Repeat steps 2 through 4 to define more routes for the router.
6. If applicable, instead of defining one route after another through steps 2 through 4, you can select the **Apply Routing Table** menu option. The Apply Routing Table dialog box is displayed.
7. Complete the fields as follows:
 - a. In the **Routing Table** list, click the routing table you want to apply to the router.
 - a. Select the **Remove existing routes** check box if you want to remove all the individual routes that have been previously set up for this device.
 - b. Click **Save**.

Configuring credentials for the power unit

For information on how to define the service access points and configure credentials for the power unit, refer to the “Configuring credentials for the server” section, earlier in this chapter.

Setting up the power unit configuration

Setting up the power unit configuration means specifying the power unit outlets and, for each outlet, the data center device that is connected to it. The following procedures are available:

- Adding a new outlet to the power unit
- Editing the outlet properties
- Removing the outlet

Adding a new outlet to the power unit

To add a new outlet to the configuration data of a power unit:

1. Click **Data center assets and resources > Inventory > Power Units > <Power unit name>**. The **Outlets** tab lists all the outlets available in the power unit. For each outlet, the device that is connected to it is specified.
2. In the **Outlet** box, type the number for the new outlet.
3. In the list of devices, click the device that is connected to the new outlet.
4. Click **Add**.

Editing the outlet properties

To edit the outlet properties:

1. Click **Data center assets and resources > Inventory > Power Units > <Power unit name>**.
2. On the **Outlets** tab, identify the outlet whose properties you want to edit, and then click its **Edit power outlet**  button. The fields corresponding to that outlet become editable.
3. Edit the outlet's properties as required.
4. Click **Save**.

Removing the outlet

To remove an outlet from the configuration data of a power unit:

1. Click **Data center assets and resources > Inventory > Power Units > <Power unit name>**.
2. On the **Outlets** tab, identify the outlet you want to remove, and then click its **Delete power outlet**  button.
3. When prompted, click **OK**.

Removing the power unit

To remove the power unit:

1. Click **Data center assets and resources > Inventory > Power Units**. The Power Unit Inventory screen lists all the currently available power units.
2. In the list, identify the power unit you want to remove, and then click its **Delete power unit**  button.
3. When prompted, click **OK**.

You have finished removing the specified power unit. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Configuring license pools

The license pool concept is used to model the license reusability: shared licenses are pooled to minimize the licensing costs. License pools can hold homogeneous licenses for reuse purposes. A license pool has an owner and also a number of configured clients. The clients are software installations for which the license pool was specified at software configuration design time.

This section describes the following procedures:

- Adding a new license pool
- Editing the properties of the license pool
- Setting up the license pool's configuration
- Editing the user-defined variables for the license pool
- Removing the license pool

Adding a new license pool

To add a new license pool:

1. Click **Data center assets and resources > Inventory > License Pools**. The **License Pools** tab lists all the available license pools.
2. Click  **Add License Pool**. The License Pool Data dialog box is displayed.
3. Complete the fields as follows:
 - a. Type the name of the new license pool, and then select a software product for it.
 - b. Specify the owner of the license pool owner.
 - c. Specify a valid from date and expiration date for the license pool (YYYY-MM-DD hh:mm:ss).
 - d. Select **No expiration date** if the new license pool will not expire.
 - e. Type the number of license keys for the new license pool.
4. Click **Save**.

Editing the properties of the license pool

To edit the properties of a license pool:

1. Click **Data center assets and resources > Inventory > License Pools**. The **License Pools** tab lists all the available license pools.
2. In the list, find the license pool whose properties you want to edit, and then click its **Edit License Pool** button . The Properties tab is displayed.
3. Edit the license pool's properties as required.
4. Click **Save**.

Setting up the license pool's configuration

This section describes the following procedures:

- Adding a new license key
- Modifying the license key's parameters
- Renaming the license key
- Removing the license key

Adding a new license key

To add a new key to the license pool:

1. Click **Data center assets and resources > Inventory > License Pools**. The **License Pools** tab lists all the available license pools.
2. In the list, click the name of the license pool to which you want to add a new key.
3. Click **Add License Key**. The License Key Data dialog box is displayed.
4. Type the name of the new key.
5. Click **Save**.

Modifying the license key's parameters

To modify the parameters of the license key:

1. Click **Data center assets and resources > Inventory > License Pools > <License pool name>**. The **Info** tab lists all the keys available in the license pool.
2. In the list, find the key whose parameters you want to modify, and then click its **Modify license key parameters**  button. The **Variables** tab is displayed.
3. In the parameter list, find the parameter that you want to modify, and then click its **Edit variable**  button.
4. Edit the selected parameter value, and then click **Save**.
5. If you want to add a new parameter to the existing list, do the following:
 - a. Type a name for the new parameter.
 - b. Select the system component that will use it.
 - c. Type a value for the new parameter.
 - d. Click **Add**.
6. If you want to remove an existing parameter from the list, click its **Delete variable**  button in the parameter list. When prompted, click **OK**.

Renaming the license key

To rename the license key:

1. Click **Data center assets and resources > Inventory > License Pools > <License pool name>**. The **Info** tab lists all the keys available in the pool.
2. In the list, click the name of the key whose name you want to change.
3. Edit the name of the license key name as required.
4. Click **Save**.

Removing the license key

To remove the license key:

1. Click **Data center assets and resources > Inventory > License Pools > <License pool name>**. The **Info** tab lists all the keys available in the pool.
2. In the list, find the key you want to remove, and then click its **Delete License Key**  button.
3. When prompted, click **OK**.

Editing the user-defined variables for the license pool

To edit the user-defined variables for the license pool:

1. Click **Data center assets and resources > Inventory > License Pools > <License pool name>**. Click the **Variables** tab.
2. Complete the fields as follows:
 - a. In the **Key** text box, type the name of the SNMP community that specifies the access rights for the license pool:
snmp.community.read
snmp.community.write
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
3. Click **Add**.

Removing the license pool

To remove the license pool from the list of available pools:

1. Click **Data center assets and resources > Inventory > License Pools**. The **License Pools** tab lists all the available license pools.
2. In the list, find the license pool you want to remove, and then click its **Delete Software Product** button .
3. When prompted, click **OK**.

Configuring software products

You can easily access an overall list of software products that are available to be installed on any server or data center device managed by the system. This section describes the following procedures:

- Adding a new software product
- Editing the properties of the software product
- Editing the user-defined variables for the software product
- Assigning workflows to the software product
- Installing a software product on a server
- Uninstalling a software product from a server
- Removing the software product

Adding a new software product

To add a new software product to the list of available software products:

1. Click **Data center assets and resources > Software Products**. The **Software** tab lists all the software products that are available to install.
2. Type the product category and a brief description for the new software product.
3. Click **Add**.

Editing the properties of the software product

To edit the properties of an existing software product:

1. Click **Data center assets and resources > Software Products**. The **Software** tab lists all the software products that are available to install.
2. In the list, locate the software product whose properties you want to edit, and then click its **Edit Software Product** button . The **Properties** tab is displayed.
3. In the Software Product Configuration dialog box, edit the software product's properties as required.
4. Click **Save**.
5. You can also add new parameters to the software product or edit the existing ones. You can define parameters that are specific to this product, such as host and product path, as well as configuration parameters that are required by workflows, such as installation path. Click **Add**.

Editing the user-defined variables for the software product

To edit the user-defined variables for the software product:

1. Click **Data center assets and resources > Inventory > Software Products > <Software product name>**. Click the **Variables** tab.
2. Complete the fields as follows:
 - a. In the **Key** text box, type the name of the SNMP community that specifies the access rights for the product:
snmp.community.read
snmp.community.write.
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
3. Click **Add**.

Assigning workflows to the software product

To learn how to assign workflows to the software product, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, later in this guide.

Installing a software product on a server

To install a software product on a server:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**.
2. Click the **Software** tab. The “Server: <Server name>” screen displays all the software products installed or available to be installed on the server.

Note: You can also click **Customer applications** and select the server that interests you. Click the **Software** tab.

3. In the **Install Software** list, click the product that you want to install on the server.
4. Click **Install**. The workflow triggered by the *Software.Install* logical device operation is executed. If the execution is successful, the software product is installed on the server.

Uninstalling a software product from a server

To uninstall a software product on the server:

1. Click **Data center assets and resources > Inventory > Servers > <Server name>**. Click the **Software** tab.

Note: You can also click **Customer applications** and select the server of interest. Click the **Software** tab.

2. In the **Software** list, identify the software product you want to uninstall from the server, and then click its **Delete Software**  button. The workflow triggered by the *Software.Uninstall* logical device operation is executed. If the execution is successful, the software product is uninstalled from the server.

Removing the software product

To remove a software product from the list of available products, follow the steps:

1. Click **Data center assets and resources > Software Products**. The **Software** tab lists all the software products available to be installed.
2. In the list, find the software product you want to remove, and then click its **Delete Software Product** button .
3. When prompted, click **OK**.

Configuring software patches

This section describes the following procedures:

- Adding a new software patch
- Querying for servers that require patching
- Applying a software patch to a server
- Editing the properties of the software patch
- Editing the user-defined variables for the software patch
- Assigning workflows to the software patch
- Removing the software patch

Adding a new software patch

To add a new software patch:

1. Click **Data center assets and resources > Inventory > Software Patches**. The Inventory Software Patches screen lists all the available patches.
2. Click the **Add Software Patch**  button. The Software Patch Configuration dialog box is displayed.
3. Type the name and a brief description for the new software patch.
4. In the **Package Path** text box, type the full patch path.
5. Select the patch type. Select the following check boxes as required:
 - **Restart Application** if the new patch needs to restart the customer application.
 - **Reboot Device** if the new patch needs to reboot the OS for that device.
 - **Co-Exist** specifies whether the patched servers can coexist in the same cluster with other servers that have not been patched.
 - **Approved** specifies whether the new patch is approved or not.
6. Click **Save**.
7. Based on the patch type you've selected, a second dialog box is displayed, prompting you for further information about the new patch.

For example, if you selected the **Copy, UnTar and Install** patch type, you would be prompted for the following additional information:

- **Filename**, specifying the parameters that are needed by the workflows for this patch type, and
- **Install Directory**, specifying the location where you want the new patch to be installed on the target server.

Querying for servers that require patching

The querying capability offered by Tivoli Intelligent Orchestrator enables you to perform various ad hoc queries in order to identify vulnerable systems in need of a particular patch. As soon as these servers are identified, you can optionally apply the required software patch. The following criteria are used for querying the data center:

Based on server location	Based on server contents
Resource pool	Products
Cluster	Patches
Application	
Customer	

To query for servers requiring patching:

1. Click **Data center assets and resources > Inventory > Software Patches**, and then click the **Query** tab.
2. Build your query by selecting all the appropriate criteria under **Server Location** and **Server Contents** until the query statement is as close to your needs as possible.
3. For each criterion you're using, save your selection by clicking the **Save Selection** button. For example, to select the resource pool containing the servers whose status you are interested in, do the following:
 - a. Expand the **Resource Pool** line.
 - b. Select the pool.
 - c. Click **Save Selection**.
4. If necessary, you also have the option to edit the SQL query you've just built. To edit the query, expand the **Modify Query** line, and edit the SQL statement as required.
5. When you have finished building your query, click the **Execute Query** button. Your query is executed, and the status of all selected servers is displayed on the screen.
6. To apply a certain patch to some or all of the displayed servers, check the names of all of the servers that require the patch, select the patch in the **Patch** list, and then click **Apply Patch to Chosen Servers**.

Applying a software patch to a server

There are several ways to apply a software patch to a server:

- Applying a software patch to a server from the Info tab of the patch
- Applying a software patch to a server from the Software Stack tab of the server
- Applying a software patch to a server within a cluster
- Editing the properties of the software patch

These procedures are described in the following subsections.

Applying a software patch to a server from the Info tab of the patch

To apply the software patch to the server:

1. Click **Data center assets and resources > Software Patches > <Software patch name>**. The **Info** tab is displayed.
2. In the **Install this patch on** list, click the server you want to install the patch on.
3. Click **Apply Patch**.

Applying a software patch to a server from the Software Stack tab of the server

To apply the software patch to the server:

1. Click **Data center assets and resources > Servers > <Server name>**, and then click the **Software** tab.
2. Select the **Install Patch** button, and then choose the software patch to be installed on the server from the list.
3. Click **Install**.

Applying a software patch to a server within a cluster

You can apply a software patch to a server that is functioning in a specific cluster by applying the patch to the entire cluster, as follows:

1. In the **Customer applications** navigation tree, select the cluster you are interested in, and then click the **Software Stack** tab.
2. In the **Install Patch** list, locate the patch you want to apply to the cluster. Select either **Live Upgrade** or **Rolling Upgrade**. The following two patching methods are available:
 - **Live Upgrade**, to successively apply the patch to all servers in the cluster. The cluster is put in **manual** mode. The workflow triggered by the *Cluster.LiveUpgrade* logical operation is executed, and all servers in the cluster are patched. The cluster is put back in its original mode.
 - **Rolling Upgrade**, to apply the patch to a server within the cluster without any service interruption, by first removing the server from the cluster before patching it. The cluster is put in **manual** mode. The workflow triggered by the *Cluster.RollingUpgrade* logical operation is executed, and one by one, unpatched servers are replaced with patched servers. This method requires that patched servers be able to coexist in the same cluster with unpatched servers. When all servers are patched and returned to their cluster, the cluster is put back in its original mode.
3. Click **Apply Patch**

Editing the properties of the software patch

To edit the properties of the software patch:

1. Click **Data center assets and resources > Software Patches**. The **Software Patches** tab lists all the software patches available to be installed.
2. In the list, find the patch whose properties you want to edit, and then click its **Edit Software Patch** button . The **Properties** tab is displayed.
3. Edit the properties of the software patch as required.
4. Click **Save**.

Editing the user-defined variables for the software patch

To edit the user-defined variables of the software patch:

1. Click **Data center assets and resources > Inventory > Software Patches > <Software patch name>**.
2. Click the **Variables** tab.
3. Complete the fields as follows:
 - a. In the **Key** text box, type the name of the SNMP community that specifies the access rights for the patch:
snmp.community.read
snmp.community.write.
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
4. Click **Add**.

Assigning workflows to the software patch

To learn how to assign workflows to the software patch, refer to the “Assigning workflows to a data center device” section, in the “Configuring and using workflows” chapter, described later in this guide.

Removing the software patch

To remove a software patch from the list of available patches:

1. Click **Data center assets and resources > Software Patches**. The **Software Patches** tab lists all the software patches available to be installed.
2. In the list, locate the software patch you want to remove, and then click its **Delete Software Patch** button .
3. When prompted, click **OK**.

Configuring software stacks

The *software stack* concept is introduced to simplify the relationship between workflows and software products, by modeling all the software dependencies and identifying software products. You can interactively add and remove software by automatically invoking the appropriate workflows as defined by the *device driver*.

This section describes the following procedures:

- Adding a new software stack
- Editing the properties of the software stack
- Editing the user-defined variables for the software stack
- Configuring the software stack entries
- Assigning workflows to the software stack
- Cloning the software stack
- Removing the software stack

Adding a new software stack

To add a new stack to the list of available software stacks:

1. Click **Data center assets and resources > Inventory > Software Stacks**. The Software Stack Configuration screen lists all the software stacks currently available in the system.
2. Click the **Add Software Stack**  button. The Software Stack Configuration dialog box is displayed.
3. Type the name and a brief description for the new software stack.
4. Click **Save**.

You have finished adding a new software stack. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the properties of the software stack

To edit the properties of an existing software stack:

1. Click **Data center assets and resources > Inventory > Software Stacks**. The Software Stack Configuration screen lists all the available software stacks.
2. In the list, identify the software stack whose properties you want to edit, and then click the **Edit Software Stack**  button. The Software Stack Configuration dialog box is displayed.
3. Edit the stack's properties as required.
4. Click **Save**.

You have finished editing the properties of the software stack. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the user-defined variables for the software stack

To edit the user-defined variables for the software stack:

1. Click **Data center assets and resources > Inventory > Software Stacks > <Software stack name>**.
2. Click the **Variables** tab.
3. Complete the fields as follows:
 - a. In the **Key** text box, type the name of the SNMP community that specifies the access rights for the stack:
snmp.community.read
snmp.community.write
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
4. Click **Add**.

Configuring the software stack entries

This section describes the following procedures:

- Adding a new entry to the stack
- Editing the configuration of the stack entry
- Removing the stack entry

Adding a new entry to the stack

A new entry added to the software stack can be either a software product or another nested stack. To add a new entry to the software stack:

1. Click **Data center assets and resources > Inventory > Software Stacks**. The Software Stack Configuration screen lists all the available software stacks.
2. Select the software stack to which you want to add the new entry. The **Info** tab is displayed.
3. If the new entry is a software product, do the following:
 - a. Select the **Install Software** button under **Software Stack Entries**.
 - b. Select the product you want to add in the list.
 - c. Click **Install**.
4. If the new entry is another stack, select its type in the **Software Stack Type** list first. Two stack types are available: **Regular Software Stack** and **Image Software Stack**, and different steps must be taken to add each of them, as described below:
 - For a regular software stack being added, do the following:
 1. In the **Software Stack Type** list, select **Regular Software Stack**, and then click **Change**.
 2. Under **Software Stack Entries**, select the **Install Stack** button first, and then select the stack you want to add from the corresponding list.
 3. Click **Install**. The new stack is added to the list of entries. If necessary, use the **Move up** and **Move down** buttons to rearrange the new entry in the list.
 - For an image software stack being added, do the following:
 1. In the **Software Stack Type** list, select **Image Software Stack**, and then click **Change**.
 2. Select the boot server where the image stack configuration should be taken from.
 3. In the corresponding text boxes, type the host where the image file is located, and the full path for the image file.
 4. Click **Save**. The specified image stack is added to the list of entries. If necessary, use the **Move up** and **Move down** buttons to rearrange the new entry in the list.
5. Repeat steps 3 and 4 to add more software products or nested stacks to the current software stack.

All other software stacks and objects that are using the current stack are displayed on the lower half of the screen.

Editing the configuration of the stack entry

To edit the configuration of the stack entry:

1. Click **Data center assets and resources > Inventory > Software Stacks > <Software stack name>**. The **Info** tab lists all products and nested stacks composing the current stack.
2. Click the product whose configuration you want to edit. The **Properties** tab is displayed.
3. You can add new parameters or edit the existing ones. You can also define parameters that are specific to this product, such as host and product path, as well as configuration parameters that are required by workflows, such as installation path. Click **Add**.

Note: Software products can only be configured in their original software stacks.

Removing the stack entry

To remove a stack entry:

1. Click **Data center assets and resources > Inventory > Software Stacks > <Software stack name>**. The **Info** tab lists all the stack entries composing the current stack.
2. In the list, identify the stack entry that you want to remove, and then click its **Remove Software Product**  button.
3. When prompted, click **OK**.

Assigning workflows to the software stack

You can assign workflows to software stacks, as it is the case with all other data center devices. In order for the workflow to be executed in the specified software stack, you need to perform the following procedures:

- Tagging the workflow with a logical device operation
- Assigning workflows to the software stack

These procedures are described in the following subsections.

Tagging the workflow with a logical device operation

To learn how to tag a workflow with a logical operation, and how to define a new logical operation if required, refer to the “Tagging the workflow with a logical device operation” section, in the “Configuring and using workflows” chapter.

Assigning workflows to the software stack

To assign workflows to a software stack:

1. Click **Data center assets and resources > Inventory > Software Stacks > <Software stack name>**.
2. Click the **Workflows** tab.
3. Under **Device Driver for Software Stack**, select the device driver whose workflows you want to assign to the current software stack.
4. Click **Change Device Driver**. The software stack will inherit all of its workflows from the selected device driver.

Cloning the software stack

To clone the software stack:

1. Click **Data center assets and resources > Inventory > Software Stacks**.
2. In the stack list, identify the software stack you want to clone, click its **More**  button, and then select the  **Clone** option.
3. When prompted, click **OK**. The Software Stack Configuration dialog box is displayed.

Note: Only the top-level stack is cloned: nested stacks are not cloned, but displayed as simple references to the original nested stack. If you want to change the nested stack, you have to clone it first, and then change the reference to the first cloned

stack so that it points to the nested stack's clone, which is now editable.

4. Edit the properties of the cloned stack as required.
5. Click **Save**.

Removing the software stack

To remove a stack entry:

1. Click **Data center assets and resources > Inventory > Software Stacks**. The **Software Stacks** tab lists all the available software stacks.
2. In the stack list, identify the software stack you want to remove, click its **More**  button, and then select the **Delete**  option.
3. When prompted, click **OK**.

Configuring resource pools

Note: This procedure requires that you have already created and configured a subnet and a VLAN. If you have not configured a subnet or VLAN, refer to the previous sections before you start creating a new resource pool.

A new resource pool is added to the list of currently available resource pools. You can edit the resource pool's properties, edit its user variables, assign a software stack to it, or remove it.

This section describes the following procedures:

- Adding a new resource pool
- Editing the resource pool properties
- Editing the user-defined variables for the resource pool
- Assigning workflows to a resource pool
- Assigning a software stack to the resource pool
- Removing a resource pool

Adding a new resource pool

To add a new resource pool:

1. Click **Data center assets and resources > Resource Pools**. The Resource Pools screen lists all the resource pools that are currently available in the system.
2. Click **Add pool** . The Resource Pool Data dialog box is displayed.
3. In the **Name** text box, type a name for the new resource pool.
4. In the **VLAN** list, click a VLAN for the new resource pool.
5. Click **Save**.

The new resource pool is added to the **Resource Pools** list, under the **Data center assets and resources > Resource Pools** menu.

You have finished adding a new resource pool. Click the **Data center assets and resources** navigation tree to refresh the displayed data.

Editing the resource pool properties

To edit the properties of a resource pool:

1. Click **Data center assets and resources > Resource Pools** to list all the available resource pools.
2. In the list, find the resource pool whose properties you want to edit, and then click its **Edit pool** button . The Resource Pool Data dialog box is displayed.
3. Edit the pool properties as required.
4. Click **Save**.

Editing the user-defined variables for the resource pool

To edit the user-defined variables of the resource pool:

1. Click **Data center assets and resources > Resource Pools**, and click the pool whose variables you want to edit.
2. Click the **Variables** tab.
3. Complete the fields as follows:
 - a. In the **Key** text box, type the name of the SNMP community that specifies the access rights for the servers in the resource pool:
snmp.community.read to query the servers
snmp.community.write to change the servers.
 - b. In the **Component** list, click the system component that will use this variable.
 - c. In the **Value** box, type a value for the chosen SNMP community.
4. Click **Add**.

Assigning workflows to a resource pool

To learn how to assign workflows to a resource pool, refer to the “Assigning workflows to a resource pool” section, in the “Configuring and using workflows” chapter, described later in this guide.

Assigning a software stack to the resource pool

To assign a software stack to the resource pool:

1. Click **Data center assets and resources > Resource Pools** and click the pool to which you want to assign a software stack.
2. Click the **Software Stack** tab.
3. In the **Software Stack Configuration** list, click the stack you want to assign to the pool.
4. Click **Save**. The selected software stack is assigned to the pool, and the content of the stack is displayed on the screen.
5. From this screen, you can edit the stack’s configuration by clicking either the pool’s **Current Software Stack** link or its corresponding **Edit stack configuration**  button.
6. You can also edit the properties of any software product in the stack, by clicking the **Edit product configuration**  button corresponding to that product.

Removing a resource pool

To remove a resource pool:

1. Click **Data center assets and resources > Resource Pools** to list all the available resource pools.
2. Select the resource pool you want to remove, and then click its **Delete pool** button .
3. When prompted, click **OK**.

Chapter 3. Managing Tivoli Intelligent Orchestrator

This chapter provides you with the necessary information to operate and control the behavior of Tivoli Intelligent Orchestrator. It describes the procedures to add and remove servers to and from customer application environments and resource pools, manage customer accounts, applications, clusters, and individual servers, as well as switch operating modes at the application level.

Also provided are the procedures to monitor the customers, applications, and resource pools available in your system, and evaluate the performance of individual servers.

This chapter includes information on the following:

- Changing the operating mode of an application
- Changing the operating mode of a cluster
- Managing customer accounts
- Managing applications and clusters
- Managing servers
- Performing deployment tasks from the interactive views
- Monitoring the system
- Changing credential data in database tables using a new encryption key

Changing the operating mode of an application

Changing the operating mode of an application can be done independently from the global operating mode. The available operating modes for an application are: **automatic**, **semiautomatic**, **manual**, and **default**.

Note: A description of the **automatic**, **semiautomatic**, and **manual** modes is provided in the “The Operating Modes” section, in the “Introduction to the IBM Tivoli Intelligent Orchestrator” chapter.

The global operating mode can be inherited by all the applications managed by the system, provided the operating mode for each application was initially set to **default**. The global operating mode does not override any application mode other than **default**.

You can change the operating mode of an application by using either the system’s Web-based interface, or the command-line interface. These procedures are described in the following sections.

Using the Web-based interface

The Web-based interface provides two different ways to change the current operating mode of an application:

- Using the Customer applications menu
- Using the Command box

Using the Customer applications menu

To change the application operating mode using the **Customer applications** menu:

1. Click the **Customer applications** navigation tree. Select the application whose operating mode you want to change. The **Clusters** tab is displayed.
2. In the **Operating Mode** list, click the application operating mode you want to switch to. The default mode is **default**.
3. Click **Change**.

Using the Command box

To change the current operating mode of an application using the **Command** box:

1. In the **Command** box, type the appropriate SOAP command to switch to the operating mode you want. For example, to change the current application operating mode to **manual**, the execution line that you have to type has the following format:

```
soapcli.sh <username> <password>  
http://localhost:9080/tcSoap/wsdl/com/ibm/tc/soap/OperationsModeService.w  
sdl setApplicationMode <appl-id> <new-mode>
```

where the variables are as follows:

<app-id> is the numeric identifier of the application. To determine the numeric identifier, hold the mouse cursor over the name of the application.

<new-mode> is the operating mode to which you want to switch

2. Press **Enter**.

Using the command-line interface

To change the current operating mode of an application using the command-line interface, open a telnet session to Tivoli Intelligent Orchestrator in a Command prompt window, and type the appropriate execution line. For example, for changing the application's current operating mode to **manual**, the execution line is as follows:

```
soapcli.sh <username> <password>  
http://localhost:9080/tcSoap/wsdl/com/ibm/tc/soap/OperationsModeService.w  
sdl setApplicationMode <appl-id> <new-mode>
```

where the variables are as follows:

<app-id> is the numeric identifier of the application. To determine the numeric identifier, hold the mouse cursor over the name of the application.

<new-mode> is the operating mode to which you want to switch

Changing the operating mode of a cluster

Changing the operating mode of a cluster can be done independently from changing the operating mode of the parent application. The available operating modes for a cluster are: **automatic**, **semiautomatic**, **manual**, and **default**.

The application operating mode can be inherited by all the clusters within that application, provided the operating mode for each cluster was initially set to **default**. The application operating mode does not override any cluster mode other than **default**.

You can change the operating mode of the cluster using either the Web-based interface, or the command-line interface. These procedures are described in the following sections.

Using the Web-based interface

There are two different ways to change the current operating mode of a cluster using the Web-based interface, as follows:

- Using the Customer applications menu
- Using the Command box

Using the Customer applications menu

To change the operating mode using the **Customer applications** menu:

1. Click the **Customer applications** navigation tree. Select the cluster whose operating mode you want to change. The **Info** tab is displayed.
2. In the **Operating Mode** list, click the cluster operating mode you want to switch to. The default mode is **default**.
3. Click **Change**.

Using the Command box

To change the current operating mode of a cluster using the **Command** box:

1. In the **Command** box, type the appropriate command to switch to the operating mode you want. For example, to change the current cluster operating mode to **automatic**, the command line that you have to type has the following format:

```
soapcli.sh <username> <password>  
http://localhost:9080/tcSoap/wsdl/com/ibm/tc/soap/OperationsModeService.w  
sdl setClusterMode <cluster-id> <new-mode>
```

where the variables are as follows:

<cluster-id> is the numeric identifier of the cluster. To determine the numeric identifier, hold the mouse cursor over the name of the cluster.

<new-mode> is the operating mode to which you want to switch

2. Press Enter.

Note: For additional information on all the SOAP scripts that are available through Tivoli Intelligent Orchestrator's command-line interface, refer to the "SOAP commands available through command line" chapter, described later in this guide. For more details, the "SOAP command syntax reference" appendix describes the specific syntax of some of these commands and provides relevant examples.

Using the command-line interface

To change the current operating mode of a cluster using the command-line interface, open a telnet session to Tivoli Intelligent Orchestrator in a command prompt window, and type the appropriate command.

For example, for changing the cluster's current operating mode to **manual**, type the following command line:

```
soapcli.sh <username> <password>  
http://localhost:9080/tcSoap/wsdl/com/ibm/tc/soap/OperationsModeService.w  
sdl setClusterMode <cluster-id> <new-mode>
```

where the parameters are as follows:

<cluster-id> is the numeric identifier of the cluster. To determine the numeric identifier, hold the mouse cursor over the name of the cluster in the Web-based interface.

<new-mode> is the operating mode to which you want to switch

Managing customer accounts

Tivoli Intelligent Orchestrator customer accounts have a hierarchical structure. Each customer account contains one or more customer applications. Each customer application contains one or more application clusters, and each cluster contains one or more dedicated servers.

To set up a new customer account, you need to use the **Data center assets and resources** navigation tree to specify all the information that is necessary to configure the customer and its applications, as well as its infrastructure. These configuration procedures are provided in the “Configuring Tivoli Intelligent Orchestrator” chapter, described earlier in this guide.

This section describes the following procedures:

- Setting up a new customer account
- Editing the customer account properties
- Removing the customer account

To turn a newly created customer account into a fully operational customer account, you need to set up and operate all the components of the customer account, that is, a new application, a new cluster, and a new dedicated server. These procedures are described in the “Managing applications and clusters” section, later in this chapter.

Setting up a new customer account

Note: The following procedure requires that the dedicated servers and all the network hardware and software are already installed and configured.

To set up a new customer account:

1. Click **Customer applications > Customers**. The “Customer accounts” screen lists all the customer accounts that are currently available in the system.
2. Click the **Add customer**  button. The **Properties** tab is displayed.
3. Type the name of the new customer account, and then click **Save**.

To be able to see the new customer account added to the **Customers** list, click the **Customer applications** navigation tree to refresh the displayed data.

Editing the customer account properties

To edit the properties of the customer account:

1. Click **Customer applications > Customers**. The **Customer accounts** screen lists all the customer accounts that are currently available.
2. In the list, identify the customer account whose properties you want to edit, and then click its **Edit customer**  button. The **Customer Data** dialog box is displayed.
3. Edit the customer properties as required.
4. Click **Save**.

Removing the customer account

To remove an empty customer account:

1. Click **Customer applications > Customers**. The Customer accounts screen is displayed.
2. In the list, identify the customer account you want to remove, and then click its **Delete customer** button .
3. When prompted, click **OK**.

Managing applications and clusters

After you have set up a new customer account, you need to set up a new application for it. To do this, you need to use the **Customer applications** menu to configure the application, as well as the **Data center assets and resources** menu to configure its infrastructure.

This section describes the following procedures:

- Setting up a new application
- Editing the properties of the application
- Adding a new cluster to the application
- Setting up the routing properties for the cluster
- Assigning workflows to the application cluster
- Assigning a software stack to the application cluster
- Adding a new dedicated server to the application cluster
- Bringing an application online
- Manually moving the application to maintenance
- Removing an application from a customer account
- Removing a cluster from an application

Setting up a new application

To set up a new application:

1. Click **Customer applications > Customers** and select the customer to which you want to add a new application. The **Applications** tab lists all the applications that are currently available for the selected customer.
2. Click **Add application** . The Application Data dialog box is displayed.
3. Complete the fields as follows:
 - a. Type a unique name for the new application.
 - b. In the **SLA Service Plan** list, click an SLA service plan. Available are: *Platinum*, *Gold*, and *Silver*. The value in this field determines the resource priority assigned to the application. The *Platinum* service plan offers the application the highest priority to be provided with resources when needed.
 - c. Type the appropriate value in the **Minimum percentage of time available** text box.
 - d. Type the appropriate number of seconds in the **Maximum response time available** text box.

4. Click **Save**. The new application is added to the list of applications available for the selected customer.
5. If necessary, you can also select the patch you want to install on the new application, and specify whether you want it installed as a live or a rolling upgrade, and then click **Apply Patch**.

You have finished adding a new application. To see the new application added to the **Applications** list, click the **Customer applications** navigation tree to refresh the data displayed on it.

The newly added application is kept in maintenance mode until you set up and properly configure its components.

Editing the properties of the application

To edit the properties of the application in maintenance:

1. Click **Data center assets and resources > Maintenance > Applications** and select the application whose properties you want to edit.
2. Click the **Properties** tab. The Application Data dialog box is displayed.
3. Edit the application properties as required.
4. Click **Save**.

Adding a new cluster to the application

Before you set up a new application cluster, you must have all the configuration data for the infrastructure that the cluster will rely on. Ensure that the following data is available:

- The number of the VLAN in which the application cluster will be deployed
- The subnet in which cluster server IP addresses will be defined
- The virtual IP address for each load balanced application cluster
- The identity of the associated load balancer
- The TCP port for each cluster, if required by the cluster technology
- The operating system and platform parameters for each application cluster
- The default gateway IP address for each application cluster
- Details of all the static routes defined on the servers of each cluster

Note: The application to which you want to add a new cluster must be in maintenance mode.

To add a new cluster to an application in maintenance:

1. Click **Data center assets and resources > Maintenance > Applications** and click the application to which you want to add the new cluster.
2. On the **Clusters** tab, click **Edit > Add Cluster**. The Cluster Data dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **Name** text box, type a unique name for the new cluster.
 - b. In the **Tier #** box, type the number of tiers for the cluster's application (the default value is 0).
 - c. Ensure the **Managed by the Tivoli Intelligent Orchestrator** check box is selected.
 - d. In the **Virtual IP** list, click the VIP address for the new cluster.
 - e. In the **Resource Pool** list, click a resource pool.
 - f. Type the appropriate numbers in the **Overflow Servers** text boxes.
 - g. In the **VLAN** list, click the VLAN for the cluster.
 - h. If necessary, click **Add VLAN** to add a new VLAN.
4. Click **Save**.

You have finished adding a new cluster to the application in maintenance. To see the new cluster added to the **Clusters** list for that application, click the **Customer applications** navigation tree to refresh the displayed data.

Setting up the routing properties for the cluster

To be able to set up the routing properties for a cluster, the cluster must be in maintenance. In order to put a cluster in maintenance, the parent application must be in maintenance.

To set up the routing properties of a cluster in maintenance:

1. In the **Customer applications** navigation tree, select the cluster you are interested in.
2. Click the **Routing** tab.
3. Type the subnet IP address for the cluster, the subnet mask, and the default gateway for the cluster. Select the **Default** check box to set up this gateway IP as the default for the cluster.
4. Click **Add**.
5. Repeat the steps 3 and 4 for all the subnet IP addresses, so that these IP addresses will not be used for other devices.

Assigning workflows to the application cluster

To learn how to assign workflows to an application cluster in maintenance, refer to the “Assigning workflows to an application cluster” section, in the “Configuring and using workflows” chapter.

Assigning a software stack to the application cluster

To assign a software stack to the cluster:

1. In the **Customer applications** navigation tree, select the cluster you are interested in.
2. Click the **Software Stack** tab.
3. In the **Software Stack Configuration** list, click the stack you want to assign to the cluster.

4. Click **Save**. The selected software stack is assigned to the cluster, and the content of the stack is displayed on the screen.
5. From this screen, you can edit the stack configuration by clicking either the cluster's **Current Software Stack** link or its corresponding **Edit stack configuration**  button.
6. You can also edit the properties of any software product in the stack, by clicking the **Edit product configuration**  button corresponding to that product.

Adding a new dedicated server to the application cluster

To add a new dedicated server to the application cluster in maintenance:

1. In the **Customer applications** navigation tree, select the cluster you are interested in. The **Info** tab is displayed.
2. Click **Add dedicated server** . The Server Data dialog box is displayed.
3. Complete the fields as follows:
 - a. In the **Name** box, type a unique name for the new dedicated server.
 - b. In the **Belongs to** list, click the application cluster the new dedicated server belongs to.
 - c. In the **OS** list, click the operating system for the new dedicated server.
 - d. Under **SNMP community names**, specify the access rights for the new dedicated server:
 - Read** to query the server
 - Write** to change the server.
4. Click **Save**.

You have finished adding a new dedicated server to the application cluster. To see the new dedicated server added to the **Servers** list for that cluster, click the **Customer applications** navigation tree to refresh the displayed data.

The newly added server is kept in maintenance to allow you to configure the remainder of its parameters (NICs and network interfaces).

Note: For more information on how to configure a NIC and a network interface for a server in maintenance, refer to the “Setting up a NIC for the server” and “Setting up a network interface for the server” sections in the “Configuring Tivoli Intelligent Orchestrator” chapter, described earlier in this guide.

Bringing an application online

When you have finished configuring all the application's components, you are ready to bring the application online, as follows:

1. Click **Data center assets and resources > Maintenance > Applications** and click the application you want to remove from maintenance and bring online. The **Clusters** tab is displayed.

Note: You can also perform these steps from the **Customer applications** navigation tree, by selecting the offline application you are interested in. The **Clusters** tab is displayed.

2. Click **Management >  Out of maintenance**.

You have finished bringing the application online. Its name is removed from the **Applications** list under **Data center assets and resources > Maintenance**.

Manually moving the application to maintenance

You can manually move an entire customer application to maintenance, when you need to reconfigure its components. To do this:

1. In the **Customer applications** navigation tree, select the application you are interested in. The **Clusters** tab is displayed.
2. Click **Management** >  **To maintenance**.

You have finished moving the customer application to maintenance. Its name is added to the **Applications** list, under **Data center assets and resources** > **Maintenance**. While in maintenance, the application and its components can be configured.

Removing an application from a customer account

To remove an application from a customer account, you need to put the application in maintenance first, as described in the previous section. To remove the application in maintenance from the customer account:

1. Click **Data center assets and resources** > **Maintenance** > **Applications** to list all the customer applications that are currently in maintenance.
2. In the list, identify the application you want to remove, and then click its **Delete application**  button.
3. When prompted, click **OK**.

Note: You can also remove an application by clicking **Customer applications** > **Customers** and select the customer from which you want to remove the application in maintenance. The **Applications** tab lists all the applications available in that customer account. In the list, identify the application in maintenance that you want to remove, click its **More**  button, and then select the  **Delete** option.

Removing a cluster from an application

To remove a cluster from an application, you need to put the application in maintenance first. To remove the cluster from the application in maintenance:

1. Click **Data center assets and resources** > **Maintenance** > **Applications** and click the application from which you want to remove the cluster. The **Clusters** tab is displayed.

Note: You can also access the Cluster tab by clicking **Customer applications** > **Customers** > <Customer name>, and click the offline application from which you want to remove the cluster. The **Clusters** tab is displayed.

2. Click the  **Delete cluster** button.
3. When prompted, click **OK**.

Managing servers

This section describes the following procedures:

- Manually adding overflow servers to an application cluster
- Manually freeing overflow servers
- Removing a dedicated server from an application cluster
- Moving available servers between resource pools

- Initializing a server within a resource pool
- Cleaning up a server within a resource pool
- Performing management tasks on a server
- Accessing the software stack for the server

Manually adding overflow servers to an application cluster

To manually add overflow servers to a cluster:

1. Click **Customer applications > Customers > <Customer name> > <Application name>**, and click the cluster to which you want to add the overflow servers. The **Info** tab is displayed.
2. In the **Operating mode** list, click **manual**, and then click **Apply**. You have changed the cluster's operating mode to **manual**.
3. In the **Overflow service** field, type the number of overflow servers you want to add, and then click **Add**.
4. Click **OK**. Your request to add the overflow servers to the application cluster is processed. The workflow triggered by the *Cluster.Add Server* logical device operation is executed. If the execution is successful, the requested servers are added to the server list in the Clusters screen.

Note: Before issuing the request to add servers, you can check if there are available servers in the resource pool. Depending on the number of available deployment workers and the number of concurrent workflows, it may take some time for your request to be processed. During this time, it is possible for another request to pass the same check and submit another request. The first workflow to obtain the servers will succeed and the second will fail.

Manually freeing overflow servers

To manually free overflow servers from a cluster:

1. Click **Customer applications > Customers > <Customer name> > <Application name>**, and click the cluster from which you want to remove the overflow servers. The **Info** tab is displayed.
2. In the **Operating mode** list, click **manual**, and then click **Apply**. You have changed the cluster's operating mode to manual.
3. Under **Overflow service**, and type the number of overflow servers you want to remove, then click **Remove**. The overflow servers are removed from the application cluster and sent back to their resource pool.

Removing a dedicated server from an application cluster

To remove a dedicated server from its application cluster, you need to put the server into maintenance first. perform the following steps:

1. Click **Customer applications > Customers > <Customer name> > <Application name>**, and click the cluster from which you want to remove the dedicated server. The **Info** tab is displayed.
2. In the **Operating mode** list, click **manual**, and then click **Apply**. You have changed the cluster's operating mode to manual.
3. In the server list, find the dedicated server you want to remove. Click the **Stop** button. The dedicated server is removed from its application cluster and sent to maintenance.
4. Click its **More**  button, and then select the  **Delete** option.
5. When prompted, click **OK**.

Moving available servers between resource pools

If a resource pool runs out of available servers, you can “borrow” available servers from another resource pool. To move the servers from the source resource pool to the destination resource pool:

1. Click **Data center assets and resources > Resource Pools > <Resource pool name>**. The **Servers** tab is displayed.
2. In the server list, select the available servers that you want to move from the current resource pool to the destination pool. In the **Move selected servers to** list, click the destination resource pool. Click **Proceed**.

Initializing a server within a resource pool

To initialize a server within a resource pool:

1. Click **Data center assets and resources > Resource Pools > <Resource pool name>**. The **Servers** tab is displayed.
2. In the list, identify the server you want to initialize, click its **More**  button, and then select the  **Initialize** option. The workflow triggered by the **SparePool.InitializeServer** logical operation is executed. If the execution is successful, the selected server is initialized.
If no workflow triggered by the **SparePool.InitializeServer** logical operation is available, and then the default workflow triggered by the **SoftwareStack.Install** logical operation is executed, installing the resource pool's software stack on that server.

Cleaning up a server within a resource pool

To clean up a server within a resource pool:

1. Click **Data center assets and resources > Resource Pools > <Resource pool name>**. The **Servers** tab is displayed.
2. In the list, identify the server you want to clean up, click its **More**  button, and then select the  **Cleanup** option. The workflow triggered by the **SparePool.CleanupServer** logical operation is executed. If the execution is successful, the selected server is cleaned up.
If no workflow triggered by the **SparePool.CleanupServer** logical operation is available, and then the default workflow triggered by the **SoftwareStack.Uninstall** logical operation is executed, removing the existing software stack from that server.

Performing management tasks on a server

A number of management tasks can be performed on a server, by accessing logical commands that enable you to execute these tasks. The management tasks that can be initiated and executed on the server are:

- Powering the server on 
- Powering the server off 
- Rebooting the server 
- Initializing the server 

These procedures are described in the following subsections.

In a similar way, the device rebooting and initialization management tasks can be performed for such data center devices as switches and load balancers.

Powering the server on

To power on a server:

1. Click **Customer applications** and select the server you want to power on. The **Info** tab is displayed.

Note: You can also click **Data center assets and resources > Inventory > Servers**, and select the server that interests you. The **Info** tab is displayed.

2. Click **Management >  Power On**.
3. When prompted whether you want to power on the server, click **OK**. The workflow triggered by the *Device.PowerOn* logical operation is executed. If the execution is successful, the server's power supply/supplies are turned on.

Powering the server off

To power off the specified server:

1. Click **Customer applications** and select the server you want to power off. The **Info** tab is displayed.

Note: You can also click **Data center assets and resources -> Inventory -> Servers**, and select the server that interests you. The **Info** tab is displayed.

2. Click **Management >  Power Off**.
3. When prompted whether you want to power off the server, click **OK**. The workflow triggered by the *Device.PowerOff* logical operation is executed. If the execution is successful, the server's power supply/supplies are turned off.

Rebooting the server

Different reboot procedures can be applied on a server or on any other data center device, as follows:

- *Software Asynchronous Reboot*,
- *Software Synchronous Reboot*,
- *Hardware Reboot*, or
- *Generic Reboot*, the default option.

The type of reboot to be performed on a device during a *Generic Reboot* will consider first the reboot preference specified for that device through user-defined variables. If no reboot preference is specified, the following sequence will be performed: an *asynchronous software reboot* is attempted first; if this fails, a *synchronous software reboot* will be attempted, and if this one fails too, a *hardware reboot* will be performed on that device.

The following reboot procedures are described:

- Performing an asynchronous software reboot
- Performing a synchronous software reboot
- Performing a hardware reboot
- Performing a generic reboot

Performing an asynchronous software reboot

To perform an asynchronous software reboot on the server:

1. Click **Customer applications** and select the server you want to reboot. The **Info** tab is displayed.

Note: You can also click **Data center assets and resources > Inventory > Servers**, and select the server to be rebooted. The **Info** tab is displayed.

2. Click **Management >  Software Reboot Async**.
3. When prompted whether you want to reboot this server, click **OK**. The workflow triggered by the *Device.Software Reboot Async* logical operation is executed. If the execution is successful, the asynchronous software reboot is performed on the server.

Performing a synchronous software reboot

To perform a synchronous software reboot on the server:

1. Click **Customer applications** and select the server you want to reboot. The **Info** tab is displayed.

Note: You can also click **Data center assets and resources > Inventory > Servers**, and select the server to be rebooted. The **Info** tab is displayed.

2. Click **Management ->  Software Reboot Sync**.
3. When prompted whether you want to reboot this server, click **OK**. The workflow triggered by the *Device.Software Reboot Sync* logical operation is executed. If the execution is successful, the synchronous software reboot is performed on the server.

Performing a hardware reboot

To perform a hardware reboot on the server:

1. Click **Customer applications** and select the server you want to reboot. The **Info** tab is displayed.

Note: You can also click **Data center assets and resources > Inventory > Servers**, and

select the server to be rebooted. The **Info** tab is displayed.

2. Click **Management** >  **Hardware Reboot**.
3. When prompted whether you want to reboot this server, click **OK**. The workflow triggered by the *Device.Hardware Reboot* logical operation is executed. If the execution is successful, the hardware reboot is performed on the server.

Performing a generic reboot

To perform a generic reboot on the server:

1. Click **Customer applications** and select the server you want to reboot. Click the **Variables** tab.

Note: You can also click **Data center assets and resources** > **Inventory** > **Servers**, and select the server to be rebooted. Click the **Variables** tab.

2. Complete the fields as follows:
 - a. In the **Key** text box, type *reboot_preference*.
 - b. In the **Component** list, click **Deployment engine**.
 - c. In the **Value** text box, type the specific type of reboot you want to be performed on this server, for example, *software async*.
 - d. Click **Add**. Once this user-defined variable is specified, an asynchronous software reboot will be performed by default on this server.
3. Click the **Info** tab, and then click **Management** >  **Generic Reboot**. The workflow triggered by the *Device.Software Reboot Async* logical operation is executed. If the execution is successful, the software asynchronous reboot is performed on the server. If this execution is unsuccessful, the workflow triggered by the *Device.Software Reboot Sync* logical operation gets executed. If this one fails too, the workflow triggered by the *Device.Hardware Reboot* logical operation is executed, and a hardware reboot is performed on the server.

Initializing the server

To initialize the specified server:

1. Click **Customer applications** and select the server you want to initialize. The **Info** tab is displayed.

Note: You can also click **Data center assets and resources** > **Inventory** > **Servers**, and select the server to be initialized. The **Info** tab is displayed.

2. Click **Management** > **Initialize** .
3. When prompted whether you want to initialize the server, click **OK**. The workflow triggered by the *Device.Initialize* logical operation is executed. If the execution is successful, the server is initialized.

Accessing the software stack for the server

For information on how to access the software stack for a server, refer to the “Configuring the software stack” section in the “Configuring Tivoli Intelligent Orchestrator” chapter, earlier in this guide.

Performing deployment tasks from the interactive views

Note: The procedures described in this section require that you have already accessed

the interactive view that interests you, as described in the “Configuring Tivoli Intelligent Orchestrator” chapter.

A number of deployment tasks can be initiated from each interactive view, by right-clicking the icon of the device on which you want to execute the deployment, and then choosing the appropriate menu option from the displayed pop-up menu.

In the background, workflows are initiated for each deployment task. A successful workflow execution will perform the requested deployment task. In case of an unsuccessful workflow execution, you may want to go to **System configuration and workflow management > Workflow Executions**, and verify why and at what step the workflow execution failed.

The following deployment tasks are available from each interactive view:

From the switch view

Turn port ON	Right-click a port icon
Turn port OFF	
Hardware Reboot	Right-click a server icon

From the switch fabric view

Hardware Reboot	Right-click a server icon
Initialize Device	Right-click a load balancer icon

From the VLAN view

Create VLAN	Right-click the VLAN icon
Remove VLAN from fabric	
Move port to VLAN	
Create a VLAN	Right-click the switch fabric icon
Remove a VLAN	
Create connection	Right-click a router icon
Remove connection	

From the router view

Create connection	Right-click the router icon
Remove connection	
Create VLAN	Right-click the VLAN icon
Remove VLAN from fabric	
Move port to VLAN	

From the firewall view

Turn port ON	Right-click a port icon
Turn port OFF	
Move port to VLAN	
Hardware Reboot	Right-click a server icon

From the NIC view on a server

Hardware Reboot	Right-click the server icon
Software Reboot	
Device Initialize	

Turn port ON
Turn port OFF
Move port to VLAN

Right-click the port icon

Create Virtual IP
Remove Virtual IP
Add Server
Remove Server
Hardware Reboot
Initialize Device

Right-click the load balancer icon

The step-by-step procedures to perform these deployment tasks are described in the following sections:

- Creating a VLAN
- Removing the VLAN from the switch fabric
- Turning a port on
- Turning a port off
- Moving a port to a VLAN
- Creating a router connection
- Removing a router connection
- Rebooting a device
- Initializing a device
- Creating a virtual IP
- Removing a virtual IP
- Adding a server
- Removing a server

Creating a VLAN

Creating a VLAN involves adding a new VLAN to the current switch fabric. To create a VLAN from an interactive view:

1. Access the interactive view that interests you.
2. Right-click the appropriate device icon (either switch, switch fabric, or VLAN), and then choose the **Create VLAN** option from the displayed pop-up menu. The Create VLAN dialog box is displayed.
3. Complete the fields as follows:
 - a. Type the VLAN name and number in the corresponding boxes.
 - b. From the **Subnetwork** list, select a subnetwork for the new VLAN.
4. Click **Add**. Depending on what device icon you've clicked, the workflow triggered by the *<Logical Device name>.Create VLAN* logical operation is executed. If the execution is successful, the new VLAN is added to the switch fabric.

Removing the VLAN from the switch fabric

To remove the VLAN from the switch fabric:

1. Access the interactive view that interests you.
2. Right-click the appropriate device icon (either switch, switch fabric, or VLAN), and then choose the **Remove VLAN** option from the displayed pop-up menu. The Remove VLAN dialog box is displayed.
3. From the list, select the VLAN you want to remove from the switch fabric. Click **Remove**. Depending on what device icon you've clicked, the workflow triggered by the *<Logical Device name>.Remove VLAN* logical operation is executed. If the execution is successful, the VLAN is removed from the switch fabric.

Turning a port on

To turn a port on from an interactive view:

1. Access the interactive view that offers access to the port you're interested in. The interactive view can be either the switch, firewall, or NIC view.
2. Right-click the icon of the port that you want to turn on, and then choose the **Turn Port ON** option from the displayed pop-up menu. Depending on what device icon you've clicked, the workflow triggered by the *<Logical Device name>.Turn Port ON* logical operation is executed. If the execution is successful, the chosen port is turned on.

Turning a port off

To turn a port off from an interactive view:

1. Access the interactive view that offers access to the port you're interested in. The interactive view can be either the switch, firewall, or NIC view.
2. Right-click the icon of the port that you want to turn off, and then choose the **Turn Port OFF** option from the displayed pop-up menu. Depending on what device icon you've clicked, the workflow triggered by the *<Logical Device name>.Turn Port OFF* logical device operation is executed. If the execution is successful, the chosen port is turned off.

Moving a port to a VLAN

To move a port to a VLAN from an interactive view:

1. Access the interactive view that interests you.
2. Right-click the appropriate device icon (either switch, VLAN, firewall, router, or NIC), go to the **Move Port to VLAN** option on the displayed pop-up menu, and then choose the number of the destination VLAN from the displayed VLAN list. Depending on what device icon you've clicked, the workflow triggered by the *<Logical Device name>.Move Port to VLAN* logical device operation is executed. If the execution is successful, the chosen port is moved to the specified VLAN.

Creating a router connection

To create a connection or add a route to a router:

1. Access the router view that interests you, or access the VLAN view that offers you access to the router you are interested in.
2. Right-click the icon of the router to which you want to add the new route, and then choose the **Create connection** option from the displayed pop-up menu. The Create Connection dialog box is displayed.
3. From the displayed lists, select the source subnet and the destination subnet for the new route, and then click **Add**. The workflow triggered by the *Router.Create Route* logical device operation is executed. If the execution is successful, the new route is added to the router.

Removing a router connection

To remove a connection or route from a router:

1. Access the router view that interests you, or access the VLAN view that offers you access to the router you are interested in.
2. Right-click the icon of the router from which you want to remove the route, and then choose the **Remove connection** option from the displayed pop-up menu. The Remove Connection dialog box is displayed.
3. From the displayed lists, select the source subnet and the destination subnet for the route to be removed, and then click **Remove**. The workflow triggered by the *Router.Remove Route* logical operation is executed. If the execution is successful, the specified route is removed from the router.

Rebooting a device

Data center devices such as servers, switches, firewalls, or load balancers can be hardware rebooted. To initiate a hardware reboot on a certain device from an interactive view:

1. Access the interactive view that offers access to the device you want to reboot.
2. Right-click the device icon, and then choose the **Hardware Reboot** option from the displayed pop-up menu. The workflow triggered by the *Device.Hardware Reboot* logical operation is executed. If the execution is successful, the specified device is hardware rebooted.

Initializing a device

Data center devices such as servers, switches, firewalls, or load balancers can be initialized. To initialize a device from an interactive view:

1. Access the interactive view that offers access to the device you want to initialize.
2. Right-click the device icon, and then choose the **Initialize Device** option from the displayed pop-up menu. The workflow triggered by the *Device.Initialize* logical operation is executed. If the execution is successful, the specified device is initialized.

Creating a virtual IP

You can add a new virtual IP to a load balancer by initiating this deployment task from the NIC interactive view, as follows:

1. Access the NIC view that offers you access to the load balancer to which you want to add the VIP.
2. Right-click the load balancer icon, and then choose the **Create Virtual IP** option from the displayed pop-up menu. The **createVirtualIP** dialog box is displayed.
3. Type the name and the VIP address in the corresponding boxes.
4. From the **Cluster** list, select the load balanced cluster for the new VIP. In the appropriate boxes, specify the output TCP port number and the range of ports for which the algorithm will be performed. In the **Balancing Algorithm** text box, type the balancing algorithm, based on the specific device driver.
5. Click **Add**. The workflow triggered by the *LoadBalancer.Create Virtual IP* logical operation is executed. If the execution is successful, the specified VIP is added.

Removing a virtual IP

You can remove a virtual IP from a load balancer by initiating this deployment task from the NIC interactive view, as follows:

1. Access the NIC view that offers you access to the load balancer you want to work with.
2. Right-click the load balancer icon, and then choose the **Remove Virtual IP** option from the displayed pop-up menu. The **removeVirtualIP** dialog box is displayed.
3. From the list, select the VIP you want to remove, and then click **Remove**. The workflow triggered by the *LoadBalancer.Remove Virtual IP* logical operation is executed. If the execution is successful, the specified VIP is removed from the selected load balancer.

Adding a server

You can add an existing server to a load balancer by initiating this deployment task from the NIC interactive view, as follows:

1. Access the NIC view that offers you access to the load balancer you want to work with.
2. Right-click the load balancer icon, and then choose the **Add server** option from the displayed pop-up menu. The **addServer** dialog box is displayed.
3. From the **Destination Virtual IP** list, select the destination VIP for the server.
4. From the **Server** list, select the server that you want to add to the selected load balancer.
5. Click **Add**. The workflow triggered by the *LoadBalancer.Add Server* logical operation is executed. If the execution is successful, the specified server is added to the selected load balancer.

Removing a server

You can remove a server from a load balancer by initiating this deployment task from the NIC interactive view, as follows:

1. Access the NIC view that offers you access to the load balancer you want to work with.

2. Right-click the load balancer icon, and then choose the **Remove server** option from the displayed pop-up menu. The removeServer dialog box is displayed.
3. From the **Destination Virtual IP** list, select the destination VIP for the server.
4. From the **Server** list, select the server that you want to remove from the selected load balancer.
5. Click **Remove**. The workflow triggered by the *LoadBalancer.Remove Server* logical operation is executed. If the execution is successful, the specified server is removed from the selected load balancer.

Monitoring the system

The Web-based interface provides the ability to monitor, in real time, the overall performance of the customer accounts, applications, and resource pools available in your system. Also provided is the ability to monitor the performance of individual servers.

For auditing and billing purposes, you can display deployment requests for applications, clusters, resource pools, or individual servers. All these procedures are available from the **Realtime performance monitoring** navigation tree.

The following tasks are described:

- Monitoring customers and resource pools
- Monitoring overall server performance for an application or resource pool
- Displaying cluster properties for an application
- Monitoring the server performance
- Displaying deployment requests

Monitoring customers and resource pools

Tivoli Intelligent Orchestrator provides the ability to monitor the overall performance of the customer applications or resource pools available in your system. The following performance parameters are monitored:

Servers overflow/dedicated

For an application, this is the total number of running servers, both dedicated and overflow. For a resource pool, this is the total number of overflow servers, both available and assigned.

Volatility (requests per hour)

Volatility refers to the deployment activity in a customer application or resource pool. For example, more than ten deployment requests per hour for a single resource pool are considered volatile, and may lead to breach. Or, if you have applications that are in breach, but you also have spare capacity, the problem is more likely due to volatility.

Capacity utilization

This is the CPU utilization.

Breach Probability

Three color-coded tabs, red, yellow, and green, indicate the likelihood of breach for pools and applications. For example, you can click the red tab to see all the applications or pools that are likely to breach. You can drill down into further levels of detail, right down to the server performance data, to find out what happened.

You can use these metrics to tune your customers or resource pools based on actual need.

Monitoring overall server performance for an application or resource pool

You can monitor the overall server performance for a customer application or resource pool available in the system. To monitor server performance:

For an application

1. In the **Customer applications** navigation tree, select the application for which you want to display the server information. The **Servers** tab lists all the servers currently available in that application.
2. For each server in the list, the following properties are displayed:

Name	The server's name
Cluster	The application cluster the server belongs to
Pool	In case of an overflow server assigned to this application, the resource pool the overflow server belongs to
Server utilization	The server's CPU utilization, in percentages

For a resource pool

1. Click **Data center assets and resources > Resource Pools** and select the resource pool for which you want to display server data. The **Servers** tab is displayed.
2. For each overflow server in the list, the following properties are displayed:

Name	The overflow server's name
Application	The application the overflow server was assigned to
Cluster	The cluster the overflow server was assigned to
Server utilization	The overflow server's CPU utilization, in percentages

Note: The server utilization is dynamically measured using an application protocol such as SNMP. If the system is not able to measure it (server failure, network connection severance, SNMP protocol not running, the server is not configured properly, and so on.), the displayed value for the **Server utilization** is **unknown**.

Displaying cluster properties for an application

To display cluster properties for any customer application, either managed by Tivoli Intelligent Orchestrator or in maintenance:

1. In the **Customer applications** navigation tree, select the application for which you want to display the cluster properties. The **Clusters** tab displays configuration information for all the clusters configured for that application.

Note: You can also display the cluster properties by clicking **Data center assets and resources > Maintenance > Applications**, and select the application in maintenance you're interested in. The **Clusters** tab displays configuration data for all the clusters available in that application.

2. Click the cluster you are interested in. The **Info** tab is displayed.

The tab displays configuration information for the current cluster, providing details on all the servers that are available in this cluster. Provided for every server are the name, the resource pool (for assigned servers), and the percentage of server utilization.

Monitoring the server performance

Tivoli Intelligent Orchestrator provides the ability to monitor the performance of individual servers. To display the real-time charts of the daily or weekly performance of the server:

1. In the **Customer applications** navigation tree, select the server you are interested in.
2. Click the **Performance** tab.

Parameters such as user CPU utilization, system CPU utilization, CPU idleness, and nice CPU utilization are charted. Depending on the configuration, other parameters can be charted, such as memory, disk I/O, or requests per second/minute.

Displaying deployment requests

For auditing and billing purposes, as well as for internal operations, you can display deployment requests for any application, cluster, resource pool, individual overflow server, or all resource pools available in your system.

Tivoli Intelligent Orchestrator provides different ways of initiating the resource deployment. The main component that is involved in this process is the Resource Broker, that receives and manages the resource requirements from all the application controllers. For each deployment request per resource, the automated resource management component inserts a new record into the specific database, which can then be used to keep track of all the deployments for that resource.

For example, to display the deployment requests for an entire application:

1. In the **Customer applications** navigation tree, select the application you're interested in.
2. Click the **Deployments** tab. By default, the table displays data describing the deployment requests in progress. The **See workflow execution history**  button enables you to switch to the Execution Overview screen. Also provided are buttons that can be used to populate the table with the needed deployment data.

The table includes the following data:

Cluster	The deployed server's cluster.
Request	The type of the deployment request, such as, adding (+) servers to or removing (-) servers from the specified cluster, followed by the number of deployed servers.
Status	The status of the deployed server.
Date	The date and time the server deployment has been issued.
Approval	The system's component (in <i>automatic</i> mode) or the operator (in <i>semiautomatic</i> or <i>manual</i> mode) who approved the deployment, and the date and time of the approval.

The buttons you can use to populate the table with deployment data are:

- ⌚ **In Progress**
Populates the table with data describing all the deployments that are being executed.
- ✓ **Deployed**
Populates the table with data describing all the successful deployments for all the servers available in that application.
- ? **Waiting for Approval**
Populates the table with data describing all the deployment requests pending approval.
- ✗ **Failed**
Populates the table with data describing all the failed deployments for all the servers available in that application. The term “failed” indicates that at least one failure occurred during the deployment process.
- 🚫 **Declined**
Populates the table with data describing all the deployments that have been declined before being executed. The deployments can be declined either manually when the cluster is in **semiautomatic** mode, or automatically by Tivoli Intelligent Orchestrator when the cluster is in **manual** mode.
- 🗑️ **Obsolete**
Populates the table with data describing all the deployment requests that no longer require execution, as more recent requests were issued overwriting the initial ones.

Once populated, the table can span multiple pages, based on the number of deployment requests submitted for that specific resource.

Understanding encryption in Tivoli Intelligent Orchestrator

Encryption is used for the following two purposes:

- To protect sensitive information in the database (such as credentials used to manage devices, servers, and so on). This encryption is mandatory. For this, the encryption key is defined outside of the database in the `%TC_CONFIG%/crypto.xml` configuration file in the `<database-key>` element. For example:

```
<database-key>rUY0VFtbGcSP2mIZH5s9H9CJFiySdfsx</database-key>
```
- To protect sensitive information in configuration files (for example, database and LDAP passwords). All encrypted settings are encrypted with the same key. The encrypted data is stored as Base64-encoded strings. The encryption key is specified in `%TC_CONFIG%/crypto.xml` file in the `<configuration-key>` element. This encryption is optional but highly recommended. If this key is not defined (the `<configuration-key>` element is not present or commented out altogether), the data is considered clear text. If this key is present, then the sensitive data is assumed encrypted and decryption is attempted. This is attempted for all known sensitive settings.

In both cases, the encryption algorithm is TripleDES. The implementation is provided by the JCE provider configured in `jre/lib/security/java.security` of the JVM being used. No configuration is needed here as the default settings of the JVM are appropriate.

All encrypted values and encryption keys are stored as Base64-encoded strings.

Performing cryptographic operations using a script

To perform cryptographic operations (key generation, encryption and decryption), the application uses `com.thinkdynamics.kanaha.util.CryptoUtils`. However, for the purpose of performing these operations from a standalone script, a separate but compatible implementation is available in `com.thinkdynamics.installation.InstallCryptoUtils`.

The following example shows how to perform these operations using a standalone script:

```
@echo off
SETLOCAL
rem this is a test batch file to invoke InstallCryptoUtils.
SET J_HOME=E:\IBM\WebSphere\AppServer\java
SET ICU=com.thinkdynamics.installation.InstallCryptoUtils
SET INVOCATION=%J_HOME%\bin\java -classpath
%J_HOME%\jre\lib\rt.jar;bin\StandaloneCrypto.jar %ICU%
SET TESTKEY=xIYqm3kWiqjsy8KM5cIHxAFwFz7xwQfV
SET TEST_CLEARTEXT_DATA=some_data_to_encrypt
SET TEST_ENCRYPTED_DATA=uZIpaT0lhuZ0pY0d7igU+V2PSppP33Mk
rem Generate a key
%INVOCATION% generate
rem Use the test key above to encrypt clear text
echo Encrypting %TEST_CLEARTEXT_DATA%
%INVOCATION% encrypt %TESTKEY% %TEST_CLEARTEXT_DATA%
rem Use the test key above to decrypt clear text
echo Decrypting %TEST_ENCRYPTED_DATA%, it should be %TEST_CLEARTEXT_DATA%
%INVOCATION% decrypt %TESTKEY% %TEST_ENCRYPTED_DATA%
```

Table 1: Return codes for InstallCryptoUtils

Return Code	Description	Details
0	SUCCESS	The action was successful.
1	NO_SUCH_ALGORITHM_EXCEPTION	Missing implementation of the crypto algorithm. The JCE provider may be missing.
2	NULL_KEY	Null encryption/decryption key provided.
3	INVALID_KEY_LENGTH	Invalid key length. 3DES keys are 24 byte-long (raw) or 32-byte long (Base64 encoded).
4	INVALID_KEY	The encryption/decryption key was invalid.

Table 1: Return codes for InstallCryptoUtils

Return Code	Description	Details
5	INVALID_BASE64_VALUE	The Base64 value provided is not a valid Base64 value.
6	BAD_PADDING	Low level crypto error.
7	FAILED_BASE64_OPERATION	There was an error performing a Base64 operation.
8	NO_DATA_TO_ENCRYPT	No data given to encrypt.
9	NO_DATA_TO_DECRYPT	No data given to decrypt.
10	ILLEGAL_BLOCKSIZE_EXCEPTION	Low level crypto error.
11	INVALID_KEY_SPECIFICATION	Low level crypto error.
12	NO_SUCH_PADDING	Low level crypto error.
13	INCORRECT_INVOCATION_ARGUMENTS	The invocation of the Java process did not have the correct arguments. See the output for details.

Note: If the database encryption key is changed, the data in the database must be re-encrypted. For more information, see “Changing credential data in database tables using a new encryption key” on page 101.

Performing cryptographic operations manually

To manually enable encryption for configuration settings:

1. Generate a key using `InstallCryptoUtils`.
2. Store the key in `crypto.xml` under `<configuration-key>` and make sure the `<configuration-key>` element is not commented out.
3. Encrypt passwords as needed, using `InstallCryptoUtils`.
4. Store the encrypted password in the appropriate configuration file.

Note: Database encryption is enabled by default and cannot be disabled.

Changing credential data in database tables using a new encryption key

You may need to migrate all existing credentials from your database tables based on a new encryption key. When you switch to a new encryption key, you must also change encrypted data inside the database table to re-encrypt them using the new encryption key.

The **credtables.xml** file lists the tables and fields that will be used to reencrypt the data:

Table 2: credtables.xml

Table	Field
credentials_password	password
credentials_rsa	public_key, private_key
credentials_snmp	community
request_context_variable	value, output_value
runtime_context_variable	new_value, old_value

To migrate all your existing credential data from the database:

1. Backup the five database tables listed above before you migrate all your existing credential data.
2. Stop the WAS server.
3. Run the following command: Usage: `applyNewEncryptKey.cmd new_encryption_key [JDBC Driver Zip]`.
4. Modify the **crypto.xml** file to set the new encryption key.
5. Start the WAS server.

Chapter 4. Configuring and using workflows

The Deployment Engine is responsible for the creation, storage, and execution of repeatable workflows that automate all data center processes, from configuring and allocating servers, to installing, configuring, and patching software. Workflows are a key component of the Deployment Engine design.

The relationship between workflows and software products is simplified by modeling all the software dependencies and identifying software device drivers. The Deployment Engine workflows become device drivers as part of the support for various hardware and software products.

This chapter includes information on the following:

- About Deployment Engine workflows
- Configuring and using workflows
- Setting up the workflow transitions
- Configuring device drivers
- Assigning workflows
- Managing the Deployment Engine

About Deployment Engine workflows

A workflow is a sequenced set of commands. It can be large and complex, or can be as simple as a single command. Both workflows and simple commands are considered commands. Since a workflow itself can be abstracted as a command, it can be included as a step in other workflows, therefore a workflow can contain simple commands, as well as other nested workflows. In this way, workflows make it possible to build powerful library of processes that can be assembled to meet any data center process requirement.

To be able to configure, customize, and successfully execute various workflows, you need to create, configure, and manage Java plug-ins, simple commands, and logical device operations, as described in detail in the “Configuring Java plug-ins, simple commands, and logical device operations” chapter, later in this guide. All the configuration procedures including workflows, Java plug-ins, simple commands, and logical device operations are available from the **System configuration and workflow management** navigation tree.

A Java plug-in is the Java class that contains the interface or the protocol code that will interact with the devices in the data center, including the Tivoli Intelligent Orchestrator Service itself. A simple command is a wrapper for a Java plug-in, describing the plug-in’s input and output requirements through its set of input and output variables. The simple command performs an action on the external environment, such as adding or removing a server, installing an application on a server, saving the configuration of a load balancer, and so on.

The following figure illustrates the generic structure of a nested workflow:

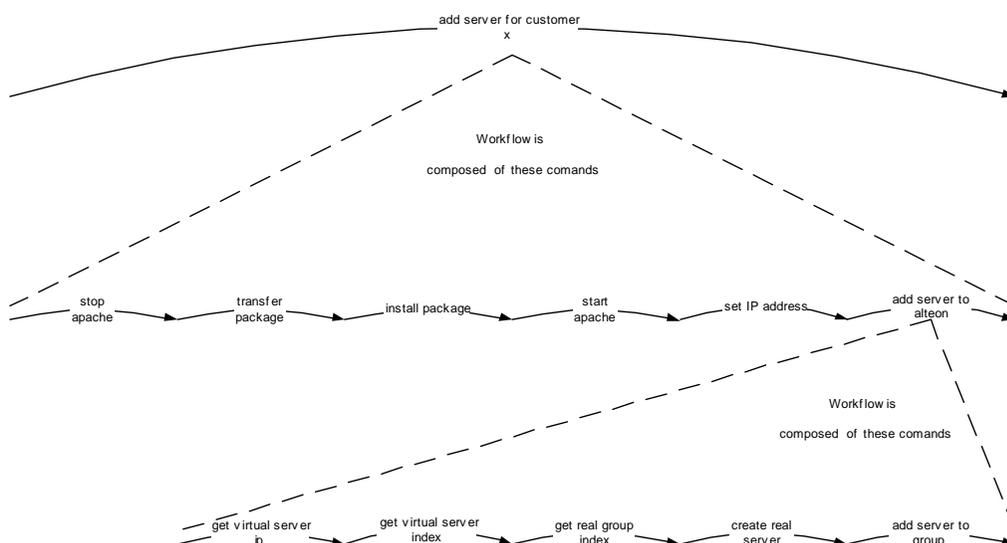


Figure 16. The generic structure of a nested workflow

The workflow framework and the plug-in design make it possible to manage any data center asset. Full editing capabilities for workflow design are provided and existing workflows can be copied, referenced, or extended. This allows you to begin with a standardized process and customize it.

The Deployment Engine supports a standard approach in creating system solutions. By being able to have multiple references to the enterprise's standard workflows, process standards can be established.

Changes to the data center environment made by the Deployment Engine are reversible in two ways:

- By design, Java plug-ins have a **do** and **undo** method. While executing normally, the **do** methods are consistently called. The **undo** methods are called in the event of a failure, reversing the work performed.
- By implementation, key workflows that build an environment can also have another workflow to break down that environment, and return the components to a resource pool.

Configuring and using workflows

This section describes the following tasks:

- Displaying workflow data
- Displaying workflow references
- Adding a new workflow
- Editing the workflow properties
- Setting up the workflow transitions
- Setting up the workflow variables
- Executing a workflow
- Cloning a workflow
- Exporting a workflow

- Importing a workflow
- Displaying the execution history for all workflows
- Displaying the execution history for an individual workflow
- Displaying execution details for a workflow
- Displaying execution details for an embedded workflow
- Deleting the execution history of a workflow

Displaying workflow data

To display the workflow transitions:

1. Click **System configuration and workflow management > Workflows**. The **Workflows** tab is displayed.
2. Select the workflow that interests you. The **Info** tab displays all the transitions composing the workflow.

There are three display areas:

Variable area

The first row, located underneath the workflow's header, can be expanded, as illustrated in the figure below. It displays the workflow variable stack, including **input**, **output**, and **local** variables.

- The *input* variables are used by the workflow as input parameters, and represent the initial set of data required by the workflow to perform its actions.
- The *output* variables represent the values resulted from the actions performed by the workflow. The output variables can be used by other workflows.
- The *local* variables are only used internally, within the workflow, to map variables from one transition to another.

Default values for some of these variables are also displayed. For every variable category, you can:

- Add a new variable, by clicking the  **Add variable** button.
- Edit an existing variable, by clicking its  **Edit variable** button.
- Remove a variable, by clicking its  **Delete variable** button.

The procedures to manage workflow variables are described in more detail in the following sections.

Transition area

The area where all transition setup and variable mapping for the current workflow takes place. All the transitions composing the workflow are listed, and every transition can be expanded to expose its input and output variables.

You can map these variables by dragging the output from one transition onto the input of another. You can drag and drop either mapped variables or unmapped variables (displayed as question marks). Dragging an unmapped variable will automatically create a new local variable in the workflow stack.

A variable that is used by multiple transitions within the same workflow is highlighted on the screen, making it easier for you to keep track of its usage. To obtain a printout of the workflow with all the variable mappings, you can expand the transitions and print the page.

Recovery Action area

The area where you can set up a recovery path for your workflow. Any workflow, command, Java plug-in, or logical operation in the nearby navigation tree can be set up as the default recovery path for your workflow, by dragging it from the navigation tree onto your workflow's **Recovery Action** area, in the right frame. You can also “borrow” a transition from another workflow and set it up as the default recovery path for your workflow.

Click the **Save Changes** button to save all the changes made to your workflow.

Displaying workflow references

Displaying references for a workflow includes all the other workflows that are currently using the selected one, as well as all the data center objects associated with the selected workflow.

To display the workflow references:

1. Click **System configuration and workflow management > Workflows**. The **Workflows** tab is displayed. In the list, select the workflow whose data you want to display. Click the **References** tab.

All workflows using the current workflow, as well as all data center objects that are associated with this workflow are displayed.

Adding a new workflow

To add a new workflow:

1. Click **System configuration and workflow management > Workflows**. The **Workflows** tab displays all the workflows that are currently defined in the system. Click  **Add workflow**.

Note: You can also click **System configuration and workflow management > Workflows > <Category group> > <Category name>**. Click **Add workflow** .

The Workflow Data dialog box is displayed.

2. In the corresponding text boxes, complete the fields as follows:
 - a. Type a unique name and a brief description for the new workflow.
 - b. Select the **Approved** check box if the new workflow is approved and can be executed.
 - c. In the **Implements Logical Operation** list, click a logical device operation for the new workflow.
 - d. In the **Category** list, click a category for the new workflow.
3. Click **Save**.

To complete the process of creating the new workflow, you also need to edit its properties, set up its transitions, and map its variables, as described in the following sections.

Note: You can also add a new workflow to the list of workflows that are specific to a certain device driver. To do this, click **System configuration and workflow management > Workflows > By Device Driver > <Device driver name>**, and then click the  **Create a new workflow** button. The remainder of the steps to create this new workflow are similar to the ones described above.

Editing the workflow properties

To edit the properties of an existing workflow:

1. Click **System configuration and workflow management > Workflows**. The **Workflows** tab lists all the currently available workflows.
2. In the workflow list, find the workflow whose properties you want to edit, click its **More**  button, and then select the **Edit**  option. Click the **Properties** tab.

Note: You can also click **System configuration > Workflows > <Group name> > <Workflow name>**. Click the **Properties** tab.

The Workflow Data dialog box is displayed.

3. Edit the workflow properties as required.
4. Click **Save**.

Setting up the workflow transitions

This section describes the following procedures:

- Adding a new transition to the workflow
- Editing transition properties
- Mapping the command variables
- Setting up the workflow recovery path
- Removing the transition

Adding a new transition to the workflow

Note: You cannot add new transitions to a workflow that has execution history. To be able to add new transitions to a workflow with execution history, you should delete its execution history first. For more information on how to delete the execution history of a workflow, refer to the “Deleting the execution history of a workflow” section, described later in this chapter.

To add a new transition to the workflow:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab lists all the transitions for the selected workflow.

Note: You can also click **System configuration > Workflows > <Category group> > <Category name> > <Workflow name>**. The **Info** tab is displayed.

2. In the navigation tree, identify the workflow, command, Java plug-in, or logical operation you want to add as a new transition to your workflow, and then drag it onto your workflow’s editing area, in the right frame.
You can also copy transitions from already existing workflows, by dragging them either from the navigation tree or from another browser window onto your workflow **Transition** area: In this case, a clone of the selected transition is created in your workflow.
3. If necessary, use the same drag and drop facility to rearrange the transitions within your workflow.

Editing transition properties

To edit the properties of a transition:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab displays all the transitions for the selected workflow.

Note: You can also click **System configuration > Workflows > <Category group> > <Category name> > <Workflow name>**.

The **Info** tab is displayed.

2. In the transition list, identify the transition whose properties you want to edit, and then click its  **Edit transition** button. A dialog box allowing you to edit the transition properties is displayed.
3. Edit the transition properties as required.
4. Click **OK**.

The next step is to map the command variables to workflow variables.

Mapping the command variables

To map the command variables to workflow variables:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab is displayed.

Note: You can also click **System configuration > Workflows > <Category group> > <Category name> > <Workflow name>**. The **Info** tab is displayed.

2. For every transition whose input variables require mapping, identify the transition, workflow, command, or logical operation whose variables can be used for mapping.
3. Drag the appropriate output variables onto the input variables that require mapping. You can drag both existing variables and unmapped variables displayed as question marks. If you drag a question mark, a new variable will be created automatically. You can also use local variables to map values between transitions.
4. Repeat step 2 and 3 for every transition that requires variable mapping.
5. Click **Save**.

Setting up the workflow recovery path

To set up the recovery path for a workflow:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab is displayed.
2. In the navigation tree, identify the workflow, command, Java plug-in, or logical operation you want to set up as the default recovery path for your workflow, and then drag it onto your workflow's **Recovery Action** area, in the right frame.

You can also copy a transition from another workflow, by dragging it either from the navigation tree or from another browser window onto your workflow **Recovery Action** area.

3. Click **Save Changes**.
If you want to check whether all variables are mapped correctly, make sure the **Compile after save** check box is also selected.

Removing the transition

To remove a transition from the current workflow:

Note: You cannot remove the last remaining transition in the workflow.

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab displays all the transitions for the workflow.
2. In the list, identify the transition you want to remove, and then click its  **Delete transition** button.
3. When prompted, click **OK**.

Setting up the workflow variables

For every workflow defined in the system, a set of input, output, and local variables can be set up and mapped.

- The *input* variables are used as input parameters by the workflow, and represent the initial set of data required by the workflow to perform its actions.
- The *output* variables represent the values resulted from the workflow's performed actions, and can be used by other workflows.
- The *local* variables are used internally within the workflow to map variables from one transition to another.

Setting up the workflow variables means:

- Adding a new variable
- Editing the variable properties
- Assigning a default value to a variable
- Removing the variable mapping
- Removing the variable

These procedures are described in the following subsections.

Adding a new variable

To add a new variable to the current workflow:

Note: To be able to add new variables to a workflow that has execution history, you need to delete its execution history first as described in the “Deleting the execution history of a workflow” section later in this chapter.

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab is displayed.
2. Expand the first top row, to display the workflow variable stack:
3. Click the  **Add a new variable** button corresponding to the type of new variable you want to add. A dialog box allowing you to enter the new variable's properties is displayed.
4. In the corresponding text boxes, type the name and a brief description for the new variable. In the **Type** list, click the variable type: **Input**, **Output**, or **Local**. If required, type a default value for the new variable in the **Default Value** text box.
5. Click **OK**. The newly defined variable is added to the workflow's variable stack.

Editing the variable properties

To edit the properties of an existing variable:

Note: You cannot edit the variable properties for a workflow with execution history unless you delete its execution history first.

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab is displayed.
2. Expand the first top row to display the workflow's variable stack.
3. In the stack, identify the variable whose properties you want to edit, and click its **Edit variable** button. A dialog box allowing you to edit the variable properties is displayed.
4. Edit the variable properties as required. You can change the variable name, its description, and its default value (if any). You can also change the variable type, in which case the variable will be transferred to the list of variables of the appropriate type.
5. Click **OK**.

Assigning a default value to a variable

To initialize a variable, you can either map it or assign a default value to it.

To assign a default value to the variable:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab is displayed.
2. In the transition list, identify the transition you are interested in, and expand it.
3. Identify the variable to which you want to assign a default value, and then click its **Assign a default value** button. A dialog box allowing you to set up the default value is displayed.
4. Type the default value for the variable, and then click **OK**. This will create a new local variable with a default value, which can be used for mapping between the workflow transitions.
5. To save the variable mapping changes, click the **Save Changes** button to the bottom of the page.
6. If you want to check whether all variables are mapped correctly, make sure the **Compile after save** check box is also selected.

Note: Warnings referring to variables that are not used or are incorrectly mapped are displayed or updated only after the workflow compilation.

Removing the variable mapping

To remove the variable mapping:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab is displayed.
2. In the transition list, identify the transition you are interested in, and expand it.
3. Identify the variable whose mapping you want to remove, and click its  **Remove variable mapping** button.
4. When prompted, click **OK**.
5. To save the variable mapping changes, click the **Save Changes** button. If you want to check whether all variables are mapped correctly, select the **Compile after save** check box.

Removing the variable

To remove a variable from the current workflow:

Note: You cannot remove a variable while another variable is mapped to it.

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. The **Info** tab is displayed.
2. Expand the first top row to display the workflow's variable stack:
3. In the variable list, identify the variable you want to remove, and then click its  **Delete variable** button.
4. When prompted, click **OK**.

Executing a workflow

To execute a workflow:

1. Click **System configuration and workflow management > Workflows**. The **Workflows** tab lists all the workflows currently defined in the system. In the list, identify the workflow that you want to execute, click its **More**  button, and then select the  **Execute** option.
2. The **Execute** tab is displayed. Specify the input variable values required for the workflow execution, then click the  **Execute** button. Your request for executing the workflow is processed, and the execution is initialized.

To learn how to display the execution details, refer to the “Displaying execution details for a workflow” section, later in this chapter. Once the workflow is executed, a new record is added to its execution history.

Note: Do not run deployment workflows manually when the destination cluster is in **automatic** mode.

Cloning a workflow

Cloning any of the existing workflows provides you with the possibility to customize it so as to meet any data center process requirement. To clone a workflow:

1. Click **System configuration and workflow management > Workflows**. The **Workflows** tab lists all the workflows that are currently defined in the system.
2. In the list, find the workflow you want to clone, click its **More**  button, and then select the  **Clone** option. The **Workflow: Copy of <Workflow name>** screen is displayed.

3. Type a unique name for the cloned workflow, and modify its properties as required.
4. Click **Save**.

To complete the process of cloning the workflow, you must also configure its transitions, as described earlier in the “Setting up the workflow transitions” section, and map its variables, as described earlier in the “Mapping the command variables” section.

Exporting a workflow

A workflow can be exported either to the screen for viewing purposes, or to a specified file:

- If exported to the screen, the workflow’s source code is displayed on your screen in the form of an unalterable XML file
- If exported to a file, the workflow is saved at a specified location as an XML file that can be used to build a powerful library of processes assembled to meet any data center process requirement. By default, the XML file is saved on the Tivoli Intelligent Orchestrator’s server.

Exporting a workflow to the screen

To export a workflow to the screen:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. Click the **Export** tab.

Note: You can also click **System configuration and workflow management > Workflows > <Category group> > <Category name> > <Workflow name>**. Click the **Export** tab.

2. Under **Destination**, select **Screen**. Select the **Include command data** and **Include Java plug-in data** check boxes if you want to include Java plug-in data as well.
3. Click **Export**. The source code for the current workflow is displayed on the screen.

Exporting a workflow to a file

To export a workflow to a file:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. Click the **Export** tab.

Note: Click **System configuration and workflow management > Workflows > <Category group> > <Category name> > <Workflow name>**. Click the **Export** tab.

2. Under **Destination**, select **File**. Complete the fields as follows:
 - a. In the corresponding text box, type the full path and the name of the xml file you want to export the workflow to.
 - b. Select the **Include command data** and **Include Java plug-in data** check boxes to include Java plug-in data as well.
3. Click **Export**. The workflow is exported to the specified xml file.

Importing a workflow

You can import a workflow from a specified XML source file. To import a workflow:

1. Click **System configuration and workflow management > Workflows**. The **Info** tab is displayed.
2. Click the **Import** tab.
3. Select **File** and specify the full path and the name of the xml source file to be imported.
4. You can import the workflow by doing one of the following:
 - Click **Import**. The workflow is imported from the specified xml file.
 - Select **Screen**. Copy the content of the xml source file into clipboard, and then paste it onto the **Screen** editable area.

Displaying the execution history for all workflows

To display the execution history for all the workflows in the system:

1. Click **System configuration and workflow management > Workflow Executions**. The **Execution Overview** tab is displayed.
2. In the **Find** list, click **All Workflows with executions**.
3. In the **For the period of** list, click the time interval you're interested in.
4. Click **Search**. The table below displays the execution history for all of the workflows with executions available in the system, for the specified period of time.

Displaying the execution history for an individual workflow

To display the execution history for an individual workflow:

1. Click **System configuration and workflow management > Workflow Executions**.
2. Do one of the following:
 - Use the request identifier of a certain workflow execution as a search criterion, and filter the displayed execution data by request ID. To do this, use the edit box located above the **Status** column and the  **Filter by request ID** button, to display only the workflow execution whose request ID you've entered.
 - Click **System configuration and workflow management > Workflow Executions > <Workflow name>**. The Workflow Execution screen is displayed.

A record is displayed for every execution of the current workflow, and includes data such as, the request number, the name of the workflow, the date and time of the execution, the name of the user who made the request, as well as the status of the execution.

In the case of a workflow with a longer execution history, you can “filter” the execution data to be displayed, using a combination of the execution's date and time, cluster, pool, and request ID as search criteria.

Search criterion	Description
Date Time	The date and time of the workflow execution. You can choose from: <i>last 2h, last 24h, last 48h, last week, last month, last year</i> .
Cluster	The cluster the current workflow was assigned to.
Pool	The resource pool that supported this workflow's execution.
Request ID	The request identifier of the workflow execution.

Displaying execution details for a workflow

To display the details of a workflow execution:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. In the list, identify the execution whose details you want to display. The **Execution Details** tab is displayed.

Displayed are all the transition executions for the workflow, including their date, time, and status.

2. Click any of the listed transition executions to view further levels of detail. The Transition Execution screen is displayed.

The displayed execution details include the date and time of the executed transition, the execution status, and the runtime context variables. Also displayed are a number of command properties, including the command name, the name of the parent workflow, the command's description and its status, whether or not it is reversible, its recovery path, the maximum number of times to retry it on *DO* and *UNDO* modes, as well as the number of seconds to wait before performing a new retry.

Displaying execution details for an embedded workflow

To display execution details for an embedded workflow:

1. Click **System configuration and workflow management > Workflow Executions > <Workflow name>**. On the **Execution History** tab, display the workflow execution history for a certain period of time.

2. Click the execution that you are interested in. The **Execution Details** tab is displayed.

Two buttons are available in the right upper corner of this screen, allowing you to toggle between two displaying modes:

Simple view mode

In the **Transition Execution** list, the name of the embedded workflow appears in bold and is preceded by the **Embedded workflow execution** icon, . This mode is illustrated in the figure above.

Detail view mode

In the **Transition Execution** list, the content of the embedded workflow is expanded, allowing you to view all of the workflows and commands comprising it. This mode is illustrated in the figure below.

3. Click the name of an embedded workflow to view further levels of detail.

The displayed details include the date and time of the executed transition, the execution status, and the runtime context variables. The details can span multiple pages.

Deleting the execution history of a workflow

You can delete the execution history of a workflow by either:

- Deleting one execution record at a time, or
- Deleting the entire execution history

These procedures are described in the following subsections.

Deleting one execution record at a time

To gradually delete the execution history of a workflow:

1. Click **System configuration and workflow management > Workflow Executions**. The **Execution Overview** tab is displayed.
2. In the **Find** list, click **All Workflows with executions**.
3. In the **For the period of** list, click the desired time interval.
4. Click **Search**. The table is populated with execution history data for all of the workflows, and for the specified period of time.

Each record in this table represents one execution of a certain workflow. For a workflow that has been executed more than once during the specified period of time, more execution records are displayed in this list, differentiated by execution IDs, date, and time.

5. Under **Workflow Name**, identify the workflow whose execution history you want to delete, and then click its **Delete execution history** button .
6. When prompted, click **OK**.
7. If more executions have been performed for the same workflow in the specified time interval, repeat the steps 2 through 6 for each execution record, until the entire execution history of the workflow is deleted.

Deleting the entire execution history

To delete the workflow's entire execution history:

1. Click **System configuration and workflow management > Workflow Executions > <Workflow name>**. The **Execution History** tab is displayed.
2. Click the  **Delete execution history** button, which is located at the top right-hand corner of the tab.

Configuring device drivers

The *device driver* abstraction represents different manufacturers' products and versions, and each device driver has a make and model number. The device driver abstraction ties the Deployment Engine workflows to a higher-level abstraction for all data center assets. For example, all Cisco CSS 11050 switches in a data center can perform the same action using the same workflow, which is dependent on the make and model and not the specific instance.

Each data center device can either inherit all of its workflows from the associated device driver or can override one or more with its own custom workflows.

This section describes the following procedures:

- Adding a new device driver
- Editing the properties of the device driver
- Assigning workflows to the device drivers
- Removing the device driver

Adding a new device driver

To add a new device driver to the device driver list:

1. Click **System configuration and workflow management > Device Drivers > <Device category>**.
2. If the new device driver requires that a new driver subcategory be defined, type the name of the new subcategory and a brief description for it, and then click **Add**.
3. In the **Drivers in** list, click the category for the new device driver. Type the name of the new driver, and then click **Add**.

Editing the properties of the device driver

To edit the properties of the device driver:

1. Click **System configuration and workflow management > Device Drivers > <Device category>**. The Device Drivers by Category screen lists all the drivers available for the selected category.
2. In the list, identify the device driver whose properties you want to modify, and then click its  **Edit device driver** button.
3. Edit the properties of the device driver as required, and then click **Save**.

Assigning workflows to the device drivers

To learn how to assign workflows to a device driver, refer to the “Assigning workflows to a device driver” section, described later in this guide.

Removing the device driver

To remove the driver from the device driver list:

1. Click **System configuration and workflow management > Device Drivers > <Device category>**. The Device Drivers by Category screen lists all of the drivers available in the current category.
2. In the list, identify the device driver you want to remove, and then click its **Delete device driver**  button.
3. When prompted, click **OK**.

Assigning workflows

Configurations and operations involving various resources managed by Tivoli Intelligent Orchestrator can be automated using workflows. Workflows can be assigned to application clusters, resource pools, device drivers, or such data center devices as servers, switches, load balancers, routers, firewalls, subnets, VLANs, and virtual IPs managed by the system.

Each application cluster, resource pool, and data center device can inherit all of its workflows from the associated device driver or can override one or more workflows with its own custom workflows.

The following procedures are described:

- Assigning workflows to an application cluster
- Assigning workflows to a resource pool
- Assigning workflows to a data center device

- Assigning workflows to a device driver

Assigning workflows to an application cluster

In order for the workflow to be executed in a certain cluster, you need to perform the following procedures:

- Tagging the workflow with a logical device operation
- Assigning workflows to the cluster

These procedures are described in the following subsections.

Tagging the workflow with a logical device operation

To tag the workflow with a logical device operation:

1. Click **System configuration and workflow management > Workflows > <Workflow name>**. Click the **Properties** tab.
2. In the **Implements Logical Operation** list, click the logical operation you want to tag the workflow with.
3. Click **Save**.

If the existing logical operations cannot be used for tagging your workflow, you may need to define a new logical operation.

Defining a new logical device operation

To define a new logical device operation:

1. Click **System configuration and workflow management > Logical Devices**. The Logical Operations screen lists all the logical operations currently defined in the system.
2. Click  **Add logical operation**. The Logical Operation dialog box is displayed.
3. Type a unique name and a brief description for the new logical operation, and then enter the EJB Proxy class name and method for the new logical operation.

Note: Note that the naming convention for logical operations is:

<Logical device name>.<Operation name>.

For example, the name of the logical operation for removing a failed server from a cluster is: *Cluster.Remove Failed Server*.

4. Click **Save**.

Assigning workflows to the cluster

To assign workflows to a cluster:

1. Click **Customer applications > Customers > <Application name>**, and select the cluster to which you want to assign the workflows. The **Info** tab is displayed.
2. Click the **Workflows** tab. The Workflows by Cluster screen is displayed.
3. In the **Logical Operation** list, click the logical operation that triggers the workflow you intend to add.
4. In the **Workflow** list, click the workflow that you want to add to the cluster.
5. Click **Add Selected Workflow**. The workflow is added to the list of available workflows for that cluster.
6. Repeat steps 3 through 5 if you want to add more workflows to the cluster.

Inheriting device driver workflows

Any application cluster can also inherit all of the workflows from an associated device driver. To have the application cluster inherit the device driver workflows:

1. Click **Customer applications > Customers > <Application name>**, and select the cluster which you want to assign the workflows to. Click the **Workflows** tab. The Workflows by Cluster screen is displayed.
2. In the **Device Driver** list under **Device Driver for cluster**, click the driver whose workflows you want the current cluster to inherit.
3. Click **Change Device Driver**. The workflows belonging to the newly selected driver are assigned to the cluster.

Assigning workflows to a resource pool

In order for the workflow to be executed in a certain resource pool, you need to perform the following procedures:

- Tagging the workflow with a logical device operation
- Assigning the workflow to the resource pool

These procedures are described in the following subsections.

Tagging the workflow with a logical device operation

To learn how to tag a workflow with a logical operation and how to define a new logical operation, refer to “Tagging the workflow with a logical device operation”.

Assigning the workflow to the resource pool

To assign a workflow to a resource pool:

1. Click **Data center assets and resources-> Resource Pools**, and select the resource pool which you want to assign the workflow to. Click the **Workflows** tab. The “Workflows by Resource pool” screen is displayed.
2. In the **Logical Operation** list, click the logical device operation that triggers the workflow you intend to add.
3. In the **Workflow** list, click the workflow that you want to add to the pool.
4. Click **Add Selected Workflow**. The workflow is added to the list of available workflows for that cluster.
5. Repeat steps 3 and 4 if you want to add more workflows to the pool.

Inheriting device driver workflows

Any resource pool can also inherit all of the workflows from an associated device driver. To have the pool inherit the device driver workflows:

1. Click **Data center assets and resources > Resource Pools**, and select the resource pool to assign to the workflows. Click the **Workflows** tab. The Workflows by Resource pool screen is displayed.
2. In the **Device Driver** list under **Device Driver for Spare Pool**, click the driver whose workflows you want the current pool to inherit.
3. Click **Change Device Driver**. The workflows belonging to the selected device driver are assigned to the resource pool.

Assigning workflows to a data center device

Workflows can be assigned to the following data center devices: servers, switches, load balancers, routers, firewalls, subnets, VLANs, virtual IPs, boot servers, and blade servers. In order for the workflow to be executed in the specified data center device, you need to perform the following procedures:

- Tagging the workflow with a logical device operation
- Assigning workflows to the data center device

These procedures are described in the following subsections.

Tagging the workflow with a logical device operation

To learn how to tag a workflow with a logical operation and how to define a new logical operation if required, refer to: “Tagging the workflow with a logical device operation”.

Assigning workflows to the data center device

To assign workflows to a data center device:

1. Click **Data center assets and resources > Inventory > <Device Group> > <Device name>**, and click the **Workflows** tab. The Workflows by <Device name> screen is displayed.
2. In the **Logical Operation** list, click the logical operation for the workflow you are interested in.
3. In the **Workflow** list underneath, click the workflow to be assigned to the device. Click **Add Selected Workflow**. The workflow is added to the list of all workflows that are currently assigned to that data center device.
4. If necessary, repeat steps 2 and 3 to assign other workflows to this device.

Inheriting device driver workflows

Any data center device can also inherit all of the workflows from an associated device driver. To have the data center device inherit the device driver workflows:

1. Click **Data center assets and resources > Inventory > <Device Group> > <Device name>**, and select the device which you want to assign the workflows to. Click the **Workflows** tab.
2. In the **Device Driver** list under **Device Driver for <Device name>**, click the driver whose workflows you want to add to the current data center device.
3. Click **Change Device Driver**. The workflows belonging to the selected device driver are assigned to the device.

Assigning workflows to a device driver

Workflows can be assigned to every supported device driver. A number of predefined workflows triggered by logical device operations are available.

In order for the workflow to be executed in the specified device driver, you need to perform the following procedures:

- Tagging the workflow with a logical device operation
- Assigning workflows to the device driver

These procedures are described in the following subsections.

Tagging the workflow with a logical device operation

To learn how to tag a workflow with a logical operation and how to define a new logical operation if required, refer to: “Tagging the workflow with a logical device operation”.

Assigning workflows to the device driver

To assign workflows to the device driver:

1. Click **System configuration and workflow management > Device Drivers > <Device driver name>**, and click the **Workflows by Device Driver** tab.
2. In the **Logical Operation** list, click the logical operation for the workflow you are interested in.
3. In the **Workflow** list underneath, click the workflow to be assigned to the specified device driver.
4. Click **Add Selected Workflow**. The workflow is added to the list of all workflows that are currently assigned to that driver.
5. If necessary, repeat steps 2 through 4 to assign other workflows to this driver.

Managing the Deployment Engine

You can manage Tivoli Intelligent Orchestrator’s Deployment Engine (DE) in either of the following ways:

- Through the Web-based user interface, accessing the **System configuration and workflow management > Configuration** menu, or
- Through a number of SOAP commands that can be used with the **Command** box in the Web-based interface

This section provides information on how to start or stop the Deployment Engine through the Web-based user interface, accessing the **System configuration and workflow management -> Configuration** menu. To do this:

1. Click **System configuration and workflow management > Configuration**. The Data Center Configuration screen is displayed.
2. The status of the Deployment Engine (**running** or **stopped**) is displayed on the lower half of the **Configuration** tab. If DE is running, as illustrated above, you can click the **Stop Deployment Engine**  button to stop it. If DE is stopped, click the **Start Deployment Engine**  button to start it.

Chapter 5. Configuring Java plug-ins, simple commands, and logical device operations

Creating, configuring, and managing Java plug-ins, simple commands, and logical device operations are required to configure, customize, and successfully execute various workflows, so as to meet any data center process requirement.

A Java plug-in is the Java class that contains the interface or the protocol code that will interact with the devices in the data center, including the Tivoli Intelligent Orchestrator. A simple command is a wrapper for a Java plug-in; describing the plug-in's input and output requirements through its variables.

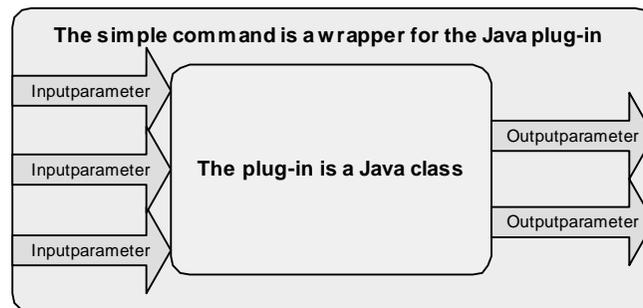


Figure 17. A simple command seen as a wrapper for a Java plug-in

The wrapper structure exists to create custom commands in which only the simple command's parameters have changed. You can have more than one simple command defining variables for the same Java plug-in. The simple command performs an action on the external environment, such as adding or removing a server, installing an application on a server, saving the configuration of a load balancer, and so on. A linked list of commands constitutes a workflow.

This chapter includes information the following information:

- Configuring Java plug-ins
- Configuring simple commands
- Configuring logical device operations

Configuring Java plug-ins

This section describes the following procedures:

- Displaying Java plug-in data in simple view and detail view modes
- Removing the Java plug-in
- Importing the Java plug-in

Displaying Java plug-in data in simple view and detail view modes

To display Java plug-in data in simple view and detail view modes:

1. Click **System configuration and workflow management -> Java Plug-ins**. The **Java Plug-ins** tab is displayed. In the list, select the plug-in you are interested in. The **Info** tab is displayed.

Or:

Click **System configuration and workflow management -> Java Plug-ins -><Java plug-in name>**. The **Info** tab is displayed.



Two displaying modes are available:

Simple view mode

Displays data for the selected Java plug-in, such as its name, description, and class name. This mode is illustrated in the preceding figure.

Detail view mode

In addition to the data displayed in the *simple view mode*, the *detail view mode* also lists all of the simple commands using this Java plug-in. This mode is illustrated in the following figure.



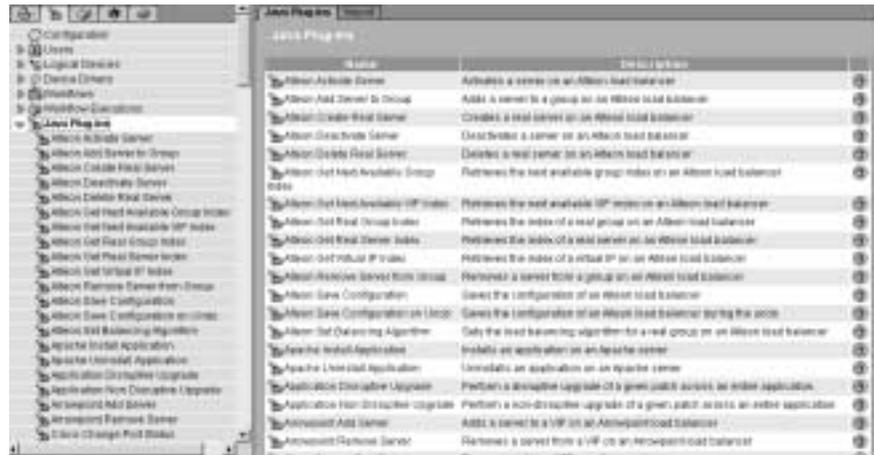
Removing the Java plug-in

To remove the Java plug-in from the list of all available Java plug-ins:

Note: You will not be allowed to remove a Java plug-in that is referenced by other

workflows, simple commands, or logical device operations.

1. Click **System configuration and workflow management -> Java Plug-ins**. The **Java Plug-ins** tab displays all the plug-ins currently available in the system.



2. In the list, identify the Java plug-in you want to remove, and then click its **Delete Java plug-in** button.
3. When prompted, click **OK**.

Note: A Java plug-in that is used by a workflow with execution history cannot be modified unless you delete the workflow's execution history. For more information on how to delete the execution history of a workflow, refer to the "Deleting the execution history of a workflow" section, described earlier in the "Configuring and using workflows" chapter of this guide.

Importing the Java plug-in

To import the Java plug-in:

1. Click **System configuration and workflow management -> Java Plug-ins**. The **Java Plug-ins** tab is displayed. Click the **Import** tab.



2. Under **Source**, select **File**. Specify the full path and the name of the xml source file to import and click **Import**. The plug-in is imported from the specified xml file.
- Or:

Under **Source**, select **Screen**. Copy the content of the xml source file into clipboard, and then paste it onto the **Screen** editable area.

Configuring simple commands

This section describes the following procedures:

- Displaying simple command data in simple view and detail view modes
- Adding a new simple command
- Editing the simple command properties
- Setting up the command's variables
- Removing the simple command
- Cloning the simple command
- Executing the simple command
- Exporting the simple command
- Importing the simple command

Displaying simple command data in simple view and detail view modes

To display the properties of the simple command in simple view and detail view modes:

1. Click **System configuration and workflow management -> Simple Commands**. The **Simple Commands** tab is displayed. Select the simple command you are interested in. The **Info** tab is displayed.

Or:

Click **System configuration and workflow management -> Simple Commands -> <Group name> -> <Simple command name>**. The **Info** tab is displayed.



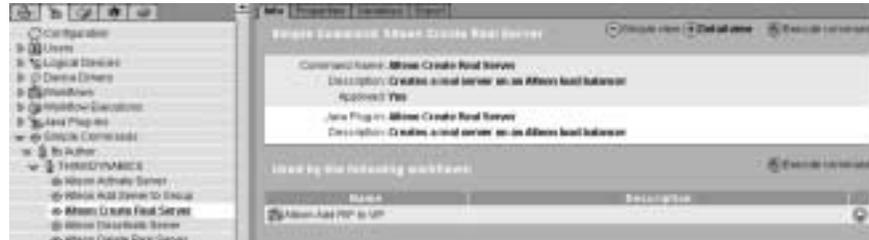
Two displaying modes are available:

Simple view mode

Displays data for the selected simple command, such as the command name and description, whether or not it is approved and reversible, the name and the description of its Java plug-in. Also provided is a button for executing the simple command. This mode is illustrated in the figure above.

Detail view mode

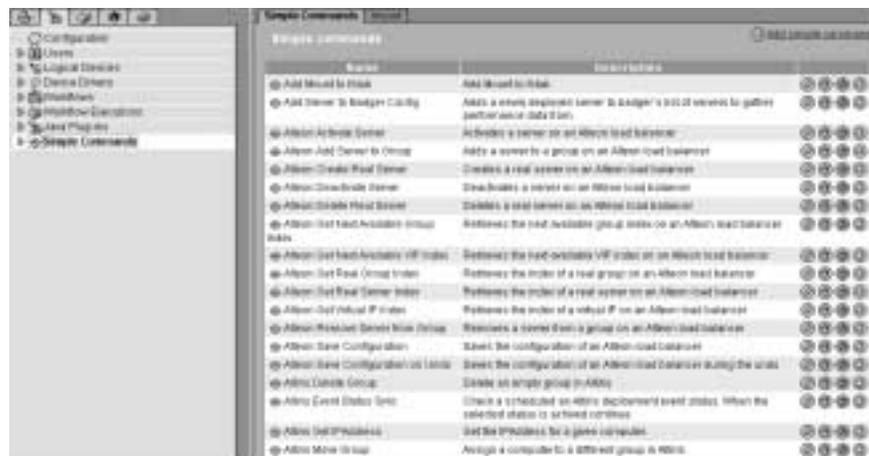
Lists the simple command data displayed in the *simple view* mode and lists all of the workflows using this command. This mode is illustrated in the following figure.



Adding a new simple command

To add a new simple command:

1. Click **System configuration and workflow management -> Simple Commands**. The **Simple Commands** tab displays all the commands available in the system.



2. Click the **Add simple command**  button. The Simple Command Property dialog box is displayed.



3. In the **Name** and **Description** text boxes, type a unique name and a brief description for the new command. In the **Java plug-in** list, click the plug-in for the new command. In the **Category** list, click a category for the new command. Select the **Approved** check box if the new command is approved and can be executed.
4. Click **Save**.

To complete the process of creating the new simple command, you also need to edit the command's properties and set up its variables, as described in the following sections.

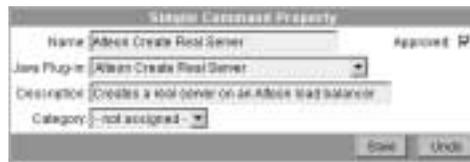
Editing the simple command properties

To edit the properties of the simple command:

1. Click **System configuration and workflow management -> Simple Commands**. The **Simple Commands** tab lists all the simple commands currently available.
2. In the list, identify the command whose properties you want to edit. Click the **Modify simple command** button . The Simple Command Property dialog box is displayed.

Or:

Click **System configuration and workflow management -> Simple Commands -> <Group name> -> <Simple command name>**. Click the **Properties** tab. The Simple Command Property dialog box is displayed.



3. Edit the command properties as required.
4. Click **Save**.

The next step in creating the new simple command is to set up the command's variables, as described in the following section.

Setting up the command's variables

For every simple command defined in the system, you must manage a set of variables. Setting up variables for a command means:

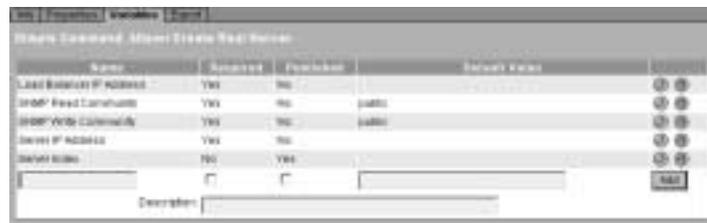
- Adding a new variable
- Editing the properties of an existing variable
- Removing the variable
- Mapping the command's variables

These procedures are described in the following subsections.

Adding a new variable

To add a new variable to the current command:

1. Click **System configuration and workflow management -> Simple Commands -> <Group name> -> <Simple command name>**. Click the **Variables** tab.



2. In the **Name** text box, type the name of the new variable you want to add. Select the **Required** check box if you want the new variable to be input for the command. Select the **Published** check box if you want the new variable to be output for the command. type the default value for the new variable in the corresponding text box. In the **Description** text box, type a brief description for the new variable.

3. Click **Add**. The newly defined variable is added to the variable list.

Editing the properties of an existing variable

To edit the properties of an existing variable:

1. Click **System configuration and workflow management -> Simple Commands -> <Group name> -> <Simple command name>**. Click the **Variables** tab. The “Simple Command: <Simple command name>” screen is displayed.
2. In the list, find the variable whose properties you want to edit. Click the **Edit variable** button . The fields corresponding to this variable become editable, as shown in the figure below:



Name	Enabled	Disabled	Default Value
Load Balancer IP Address	Yes	No	
OSMP Read Commands	Yes	No	public
OSMP Write Commands	Yes	No	public
Server IP Address	Yes	No	
Server Index	No	Yes	

Description: Server IP Address

3. Edit the properties of the variable as required.
4. Click **Save**. The properties of the variable are updated in the variable list.

Removing the variable

To remove a variable from the current command:

1. Click **System configuration and workflow management -> Simple Commands -> <Group name> -> <Simple command name>**. Click the **Variables** tab.
2. In the variable list, find the variable you want to remove, and then click its **Delete variable** button .
3. When prompted, click **OK**.

Mapping the command's variables

To learn how to map the command's variables to workflow variables, refer to the “Mapping the command variables” section, described earlier in the “Configuring and using workflows” chapter of this guide.

Removing the simple command

To remove the simple command from the list of all available commands:

Note: You will not be allowed to remove a simple command that is referenced by other workflows, Java plug-ins, or logical device operations.

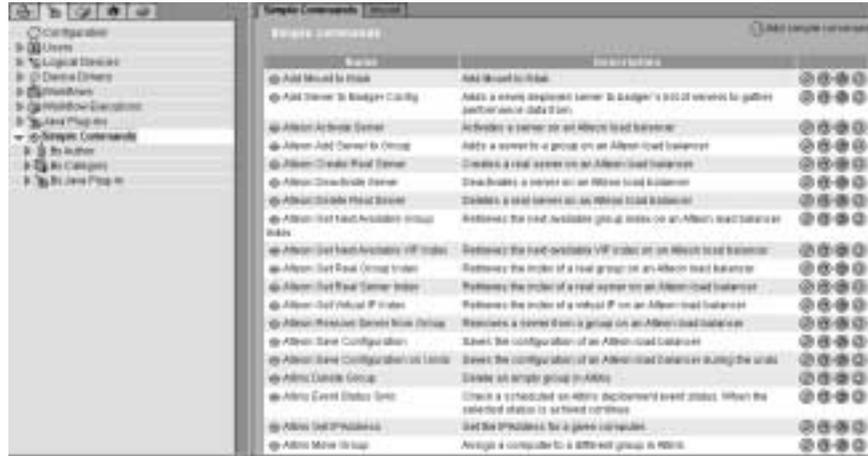
1. Click **System configuration and workflow management -> Simple Commands**. The **Simple Commands** tab lists all the simple commands currently available.
2. In the list, identify the command you want to remove, and then click its **Delete simple command** button .
3. When prompted, click **OK**.

Note: A simple command that is used by a workflow with execution history cannot be modified unless you delete the workflow's execution history. For more information on how to delete the execution history of a workflow, refer to the “Deleting the execution history of a workflow” section, described earlier in the “Configuring and using workflows” chapter of this guide.

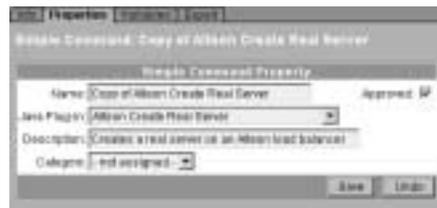
Cloning the simple command

To clone the simple command:

1. Click **System configuration and workflow management -> Simple Commands**. The **Simple Commands** tab displays all the commands available in the system.



2. In the list, identify the command you want to clone, and then click its **Clone simple command** button. The “Simple Command: Copy of <Simple command name>” screen is displayed.



3. Type a unique name for the cloned simple command, and edit its properties as required.
4. Click **Save**.

To finish the process of cloning the simple command, you must also set up the command's variables, as described earlier, in the Setting up the command's variables section.

Executing the simple command

To execute the simple command:

1. Click **System configuration and workflow management -> Simple Commands**. The **Simple Commands** tab displays all the commands available in the system.
2. In this list, identify the simple command that you want to execute. Click the **Execute simple command** button. The **Execute** tab is displayed.

Or:

Click **System configuration and workflow management -> Simple Commands -> <Group name> -> <Simple command name>**. The **Info** tab is displayed. Click the **Execute command** button. The **Execute** tab is displayed.



3. Type the required values, then click **Execute**.

Exporting the simple command

With the Tivoli Intelligent Orchestrator, a simple command can be exported either to the screen for viewing purposes or to a specified file, as follows:

- If exported to the screen, the command's source code is displayed on your screen in the form of an unalterable XML file.
- If exported to a file, the command is saved at a specified location as an XML file that can be used for later environment builds or can be easily shared with other users. By default, the XML file is saved on the Tivoli Intelligent Orchestrator server.

Exporting a simple command to the screen

To export a simple command to the screen:

1. Click **System configuration and workflow management -> Simple Commands -> <Group name> -> <Simple command name>**. Click the **Export** tab.



2. Under **Destination**, select **Screen**. Select the **Include Java plug-in data** check box if you want to include Java plug-in information.
3. Click **Export**. The source code for the current simple command is displayed on the screen.

- Under **Source**, select **File**. In the nearby box, specify the full path and the name of the xml source file to be imported and click **Import**. The simple command is imported from the specified xml file.

Or:

Under **Source**, select **Screen**. Copy the content of the xml source file into the clipboard, and then paste it onto the **Screen** editable area.

Configuring logical device operations

This section describes the following procedures:

- Displaying logical device operations in simple view and detail view modes
- Adding a new logical device operation
- Editing the properties of the logical device operation
- Setting up variables for the logical device operation
- Executing the logical device operation
- Removing the logical device operation
- Exporting the logical device operation
- Importing the logical device operation

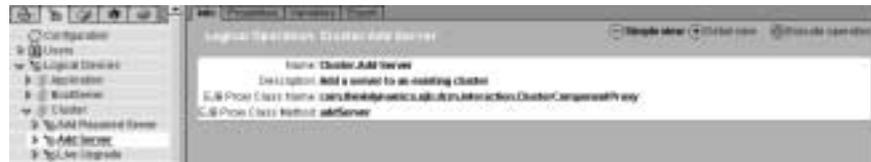
Displaying logical device operations in simple view and detail view modes

To display the logical device operation in simple or detail view modes, follow the steps:

1. Click **System configuration and workflow management -> Logical Devices**. The **Logical Operations** tab is displayed. In the list, select the logical device operation you are interested in. The **Info** tab is displayed.

Or:

Click **System configuration and workflow management -> Logical Devices -> <Logical device name> -> <Logical operation name>**. The **Info** tab is displayed.



Two displaying modes are available:

Simple view mode

Displays the logical operation's name, description, EJB Proxy class name, and EJB Proxy class method, as illustrated in the figure above.

Detail view mode

Lists the data displayed in the *simple view* mode and lists all the workflows that implement this logical device operation. This mode is illustrated in the following

figure.



Adding a new logical device operation

To add a new logical device operation:

1. Click **System configuration and workflow management -> Logical Devices**. The **Logical Operations** tab displays all the logical device operations currently defined in the system.



Or:

Click **System configuration and workflow management > Logical Devices > <Logical Device name>**. The Logical Operations by <Logical Device name> screen is displayed.



- Click the  **Add logical operation** button. The Logical Operation Data dialog box is displayed.



- Type a unique name and a description for the new logical operation, and then type the EJB Proxy class name and method for the new logical operation.
- Click **Save**.

To complete the process of creating the new logical device operation, you also need to edit the logical operation's properties and set up its variables, as described in the following sections.

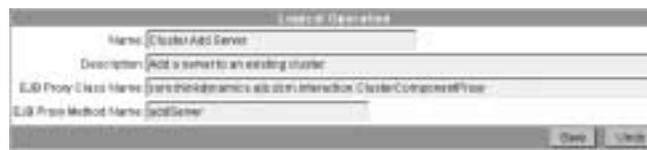
Editing the properties of the logical device operation

To edit the properties of the logical device operation:

- Click **System configuration and workflow management -> Logical Devices**. The **Logical Operations** tab is displayed. In the list, identify the logical operation whose properties you want to edit, and then click the  **Modify logical operation** button. The Logical Operation Data dialog box is displayed.

Or:

Click **System configuration and workflow management > Logical Devices > <Logical device name> > <Logical operation name>**. Click the **Properties** tab. The Logical Operation Data dialog box is displayed.



- Edit the logical operation's properties, as required.
- Click **Save**.

The next step in creating the new logical operation is to set up its variables, as described in the following section.

Setting up variables for the logical device operation

For every logical device operation defined in the system, you must manage a set of variables. Setting up variables for a logical operation means:

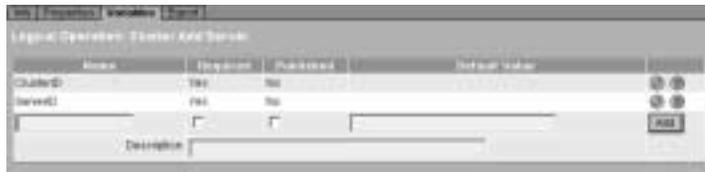
- Adding a new variable
- Editing the properties of an existing variable
- Removing the variable

These procedures are described in the following subsections.

Adding a new variable

To add a new variable to the current logical device operation:

1. Click **System configuration and workflow management > Logical Devices > <Logical device name> > <Logical operation name>**. Click the **Variables** tab.



2. In the corresponding text boxes, type the name and a brief description for the new variable you want to add.
3. Click **Add**. The newly defined variable is added to the logical operation's variable list.

Editing the properties of an existing variable

To edit the properties of an existing variable:

1. Click **System configuration and workflow management > Logical Devices > <Logical device name> > <Logical operation name>**. Click the **Variables** tab.
2. In the variable list, find the variable whose properties you want to edit, and then click its **Edit variable** button. The fields corresponding to this variable become editable, as shown below:



3. Edit the properties of the variable as required, and then click **Save**.

Removing the variable

To remove a variable from the current logical device operation:

1. Click **System configuration and workflow management > Logical Device > <Logical device name> > <Logical operation name>**. Click the **Variables** tab.
2. In the list, find the variable you want to remove, and then click its **Delete variable** button.
3. When prompted, click **OK**.

Executing the logical device operation

To execute the logical device operation:

1. Click **System configuration and workflow management > Logical Device > <Logical device name> > <Logical operation name>**. The **Info** tab is displayed.



2. Click the **Execute operation** button. The **Execute** tab displays a list with the input variables that are required for the execution.



3. Specify the variable values that are required for the operation execution, and then click **Execute**. Your request for executing the logical operation is processed, and the execution is initialized.

When the `Device.ExecuteCommand` and `Device.CopyFile` logical operation use secure shell when they are run, the Tivoli Intelligent Orchestrator server may not be running as the user that the remote machine is expecting according to its SSH key. To convert the identity to the proper user for accessing the remote machine, perform the following steps:

1. Select **System configuration and workflow management**
2. Select **Configuration**
3. Select **Variables**
4. Add a new variable with these properties:
 - Key: **DeploymentCurrentUser**
 - Component: **Deployment engine**

- Value: Type your current Tivoli Intelligent Orchestrator user name
5. Select **Inventory**
 6. Select **Servers**
 7. Select the server that the Tivoli Intelligent Orchestrator is installed on
 8. Choose the **Credentials** tab
 9. Create or select the client SAP and the password credentials to use
 10. Edit the password credentials as required
 11. In the **Username** field, type the user name that the remote host is expecting

Removing the logical device operation

To remove the logical device operation:

Note: You will not be allowed to remove a logical device operation that is referenced by other workflows, simple commands, or Java plug-ins.

1. Click **System configuration and workflow management -> Logical Devices**. The **Logical Operations** tab lists all the logical device operations currently defined in the system.



Or:

Click **System configuration and workflow management > Logical Devices > <Logical Device name>**. The Operations by Logical Device screen is displayed.



2. In the list, identify the logical operation that you want to remove, and then click its **Delete logical operation** button.
3. When prompted, click **OK**.

Exporting the logical device operation

A logical device operation can be exported either to the screen for viewing purposes or to a specified file, as follows:

- If exported to the screen, the logical operation's source code is displayed on your screen in the form of an unalterable XML file.

- If exported to a file, the logical device operation is saved as an XML file at a specified location. By default, the XML file is saved on the Tivoli Intelligent Orchestrator's server.

Exporting a logical device operation to the screen

To export a logical device operation to the screen:

1. Click **System configuration and workflow management > Logical Devices > <Logical device name> > <Logical operation name>**. Click the **Export** tab.



2. Under **Destination**, select **Screen**.
3. Click **Export**. The operation's source code is displayed on the screen.



Exporting a logical device operation to a file

To export a logical device operation to an xml file:

1. Click **System configuration and workflow management -> Logical Devices -> <Logical device name> > <Logical operation name>**. Click the **Export** tab.

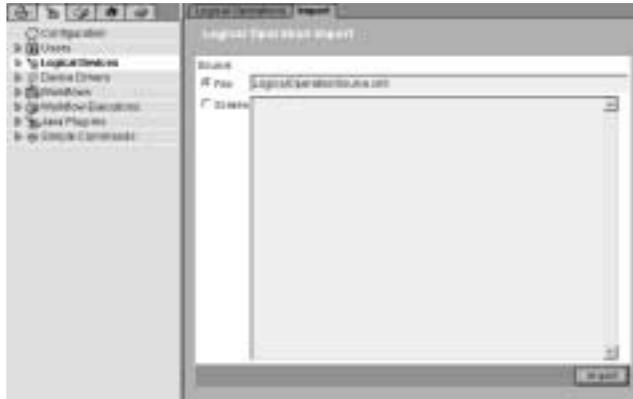


2. Under **Destination**, select **File**. Type the full path and the name of the xml file that the logical operation will be exported to.
3. Click **Export**. The logical operation is exported to the specified xml file.

Importing the logical device operation

You can import a logical device operation from a specified xml source file. To do this:

1. Click **System configuration and workflow management -> Logical Devices**. Click the **Import** tab.



2. Under **Source**, select **File**. Specify the full path and the name of the xml source file to be imported, and click **Import**. The logical operation is imported from the specified xml file.

Or:

Under **Source**, select **Screen**. Copy the content of the xml source file into clipboard, and then paste it onto the **Screen** editable area.

Chapter 6. Generating reports

The Tivoli Intelligent Orchestrator has the ability to generate a variety of reports that can be sent to the printer or into PDF files. The Tivoli Intelligent Orchestrator's reporting engine is used to generate dynamic content: its main purpose is to create page-oriented, ready-to-print reports, in a simple and flexible manner.

The Tivoli Intelligent Orchestrator organizes data retrieved from the Data Center Model database according to the report design defined in an XML file. The report design represents a template that is used by the reporting engine to deliver dynamic content to the printer or to the screen.

Note: For a practical example of the report design defined for the **Inventory** report, refer to the “*Defining the report design in an XML file: a case study*” appendix, described later in this guide.

This chapter elaborates on the following:

- *Displaying reports*
- *Report layout*
- *List of available reports*

Displaying reports

To display any of the generated reports:

1. Click **Reports** -> <**Report name**>. Adobe Acrobat is launched, and the selected report is displayed on the screen.

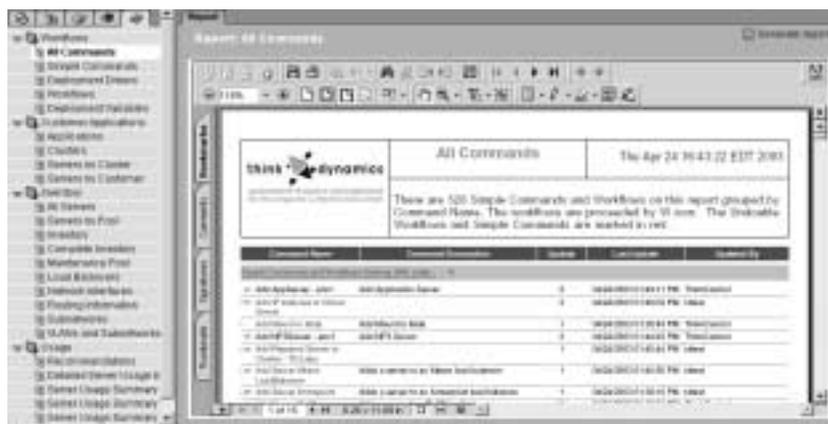


Figure 18. The Tivoli Intelligent Orchestrator displaying a report

Use the buttons on the Acrobat toolbar to send the displayed report to the printer or save it as a PDF file.

Report layout

The report design defines the content and the layout of the report's sections. The report layout includes the sections illustrated in the figure below:

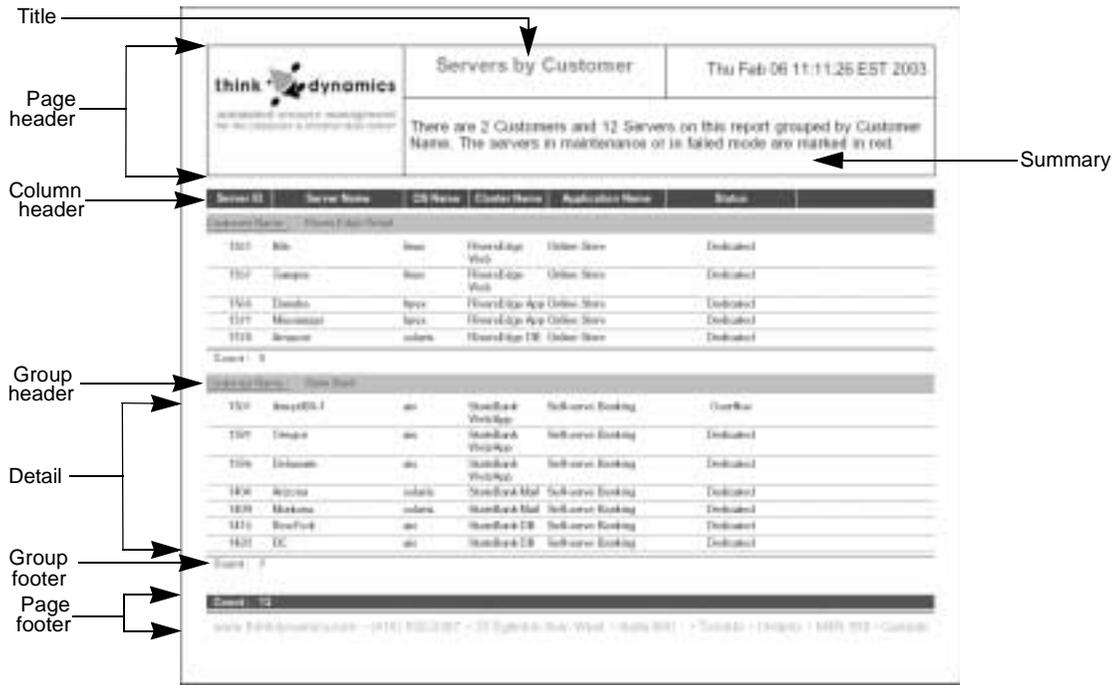


Figure 19. The report layout

List of available reports

This section provides a brief description of the content of each type of report the Tivoli Intelligent Orchestrator can generate. The following hierarchy of report categories is described:



- **Workflows**
 - *All Commands*
 - *Simple Commands*
 - *Deployment Drivers*
 - *Workflows*
 - *Deployment Variables*
- **Customer Applications**
 - *Applications*
 - *Clusters*
 - *Servers by Cluster*
 - *Servers by Customer*
- **Inventory**
 - *All Servers*
 - *Servers by Pool*
 - *Inventory*
 - *Complete Inventory*
 - *Maintenance Pool*
 - *Load Balancers*
 - *Network Interfaces*
 - *Routing Information*
 - *Subnetworks*
 - *VLANs and Subnetworks*
- **Usage**
 - *Recommendations*
 - *Detailed Server Usage by Application*
 - *Server Usage Summary by Application*
 - *Server Usage Summary by Cluster*
 - *Server Usage Summary by Pool*
 - *Detailed Server Usage Summary by Pool*

The reports are described in more detail in the following sections.

Workflows

All Commands

This report provides data on all of the currently available simple commands and workflows, grouped by name. The workflows are preceded by a **W** icon.



The following data is displayed:

Data	Description
Command Name	The simple command or workflow's unique name.
Command Description	A brief description for the simple command or workflow.
Update	The number of times the simple command or workflow has been updated.
Last Update	The date and time of the last update.
Updated By	The name of the system component or operator who performed the update.

Note: The **Generate reports**  button, which is provided on every report page, enables you to automatically generate all of the reports and update all PDFs with latest data.

Simple Commands

This report provides data on all of the currently available simple commands, grouped by command name. Simple commands that are not available are marked in red.



The following data is displayed:

Data	Description
Command Name	The simple command's unique name.
Command Description	A brief description for the simple command.
Update	The number of times the simple command has been updated.
Last Update	The date and time of the last update.
Updated By	The name of the system component or operator who performed the update.
Java Plug-in Used	The simple command's Java plug-in.

Deployment Drivers

This report provides data on all of the currently available drivers, grouped by driver name. Drivers that are not available are marked in red.



The following data is displayed:

Data	Description
Driver Name	The deployment driver's unique name.
Driver Description	A brief description for the driver.
Driver Type	The driver type.
Class Used	The driver class name.

Workflows

This report provides data on all of the currently available workflows, grouped by workflow name. Workflows that are not available are marked in red.



The following data is displayed:

Data	Description
Workflow Name	The workflow's unique name.
Workflow Description	A brief description for the workflow.
Update	The number of times the workflow has been updated.
Last Update	The date and time of the last update.
Updated By	The name of the system component or operator who performed the update.

The following data is displayed:

Data	Description
Application Name	The application's unique name.
SLO Description	A brief description for the application's Service Level Objective (SLO).
SLO Value	The SLO value.

Clusters

This report provides data on all of the currently available clusters, grouped by customer name.



The following data is displayed:

Data	Description
Cluster ID	The cluster's numeric identifier.
Cluster Name	The cluster's unique name.
Cluster Type	The cluster's type, for example, <i>iis</i> , <i>apache</i> , <i>weblogic</i> , <i>websphere</i> , <i>sendmail</i> , and so on.
App. Name	The name of the application that the cluster belongs to.
OS Name	The cluster's operating system.
Resource Pool	The resource pool for the cluster.
Min. Servers	The minimum number of servers for the cluster.
Max. Servers	The maximum number of servers for the cluster.
IP Address	The cluster's IP address.
VLAN	The cluster's VLAN.

Servers by Cluster

This report provides data on all of the currently available servers, grouped by cluster name. The servers in maintenance or in **failed** mode are marked in red.

Server ID	Server Name	OS Name	Cluster Name	Application Name	Status
100	Web	Win	BoardApp	WebApp	Dedicated
101	Server	Win	BoardApp	WebApp	Dedicated
102	Server	Win	BoardApp	WebApp	Dedicated
103	Server	Win	BoardApp	WebApp	Dedicated
104	Server	Win	BoardApp	WebApp	Dedicated
105	Server	Win	BoardApp	WebApp	Dedicated
106	Server	Win	BoardApp	WebApp	Dedicated
107	Server	Win	BoardApp	WebApp	Dedicated
108	Server	Win	BoardApp	WebApp	Dedicated
109	Server	Win	BoardApp	WebApp	Dedicated
110	Server	Win	BoardApp	WebApp	Dedicated

The following data is displayed:

Data	Description
Server ID	The server's numerical identifier.
Server Name	The server's unique name.
OS Name	The server's operating system.
Cluster Name	The name of the cluster the server is assigned to.
Application Name	The name of the application the cluster belongs to.
Status	The server's current status: available, dedicated, or overflow.

Servers by Customer

This report provides data on all of the currently available servers, grouped by customer name. The servers in maintenance or in **failed** mode are marked in red.

Server ID	Server Name	OS Name	Cluster Name	Application Name	Status
100	Web	Win	BoardApp	WebApp	Dedicated
101	Server	Win	BoardApp	WebApp	Dedicated
102	Server	Win	BoardApp	WebApp	Dedicated
103	Server	Win	BoardApp	WebApp	Dedicated
104	Server	Win	BoardApp	WebApp	Dedicated
105	Server	Win	BoardApp	WebApp	Dedicated
106	Server	Win	BoardApp	WebApp	Dedicated
107	Server	Win	BoardApp	WebApp	Dedicated
108	Server	Win	BoardApp	WebApp	Dedicated
109	Server	Win	BoardApp	WebApp	Dedicated
110	Server	Win	BoardApp	WebApp	Dedicated

The following data is displayed:

Data	Description
Server ID	The server's numerical identifier.
Server Name	The server's unique name.
OS Name	The server's operating system.
Cluster Name	The name of the cluster the server is assigned to.
Application Name	The name of the application that the cluster belongs to.
Status	The server's current status: available, dedicated, or overflow.

Inventory

All Servers

This report provides data on all of the currently available servers, grouped by pool name. The servers in maintenance or in **failed** mode are marked in red.



The following data is displayed:

Data	Description
Server ID	The server's numerical identifier.
Server Name	The server's unique name.
OS Name	The server's operating system.
Cluster Name	The name of the cluster that the server is assigned to.
Application Name	The name of the application that the cluster belongs to.
Status	The server's current status: available, dedicated, or overflow.

Servers by Pool

This report provides data on all of the currently available servers, grouped by pool name. The servers in maintenance or in **failed** mode are marked in red.

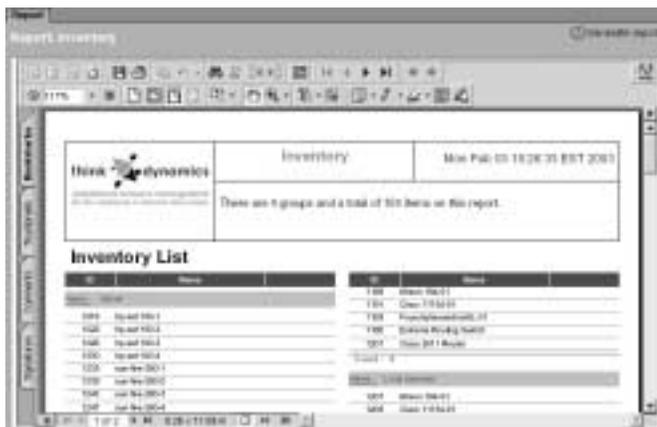


The following data is displayed:

Data	Description
Server ID	The server's numerical identifier.
Server Name	The server's unique name.
OS Name	The server's operating system.
Cluster Name	The name of the cluster that the server is assigned to.
Application Name	The name of the application that the cluster belongs to.
Status	The server's current status: available, dedicated, or overflow.

Inventory

This report provides data on all of the hardware assets and resources currently managed by the system, including servers, switches, load balancers, subnetworks, and VLANs.

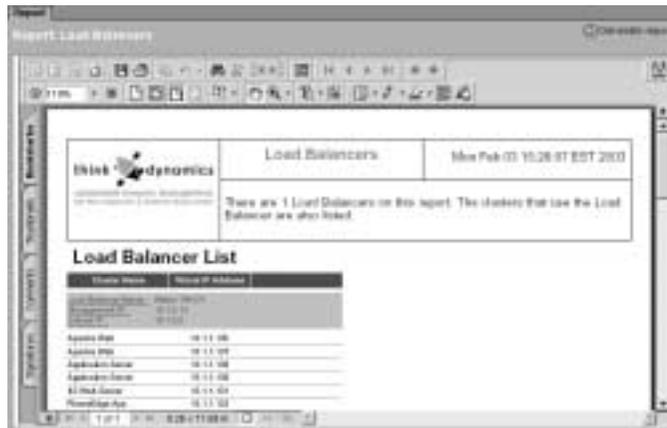


The following data is displayed:

Data	Description
ID	The numerical identifier of the item in maintenance.
Name	The unique name of the item in maintenance.

Load Balancers

This report provides data on all of the load balancers that are currently managed by the system. Also listed are the clusters that use the load balancers.



The following data is displayed:

Data	Description
Cluster Name	The name of the cluster that uses the load balancer.
Virtual IP Address	The load balancer's virtual IP address.

Network Interfaces

This report provides data on all of the currently available network interfaces, grouped by the server that uses them.

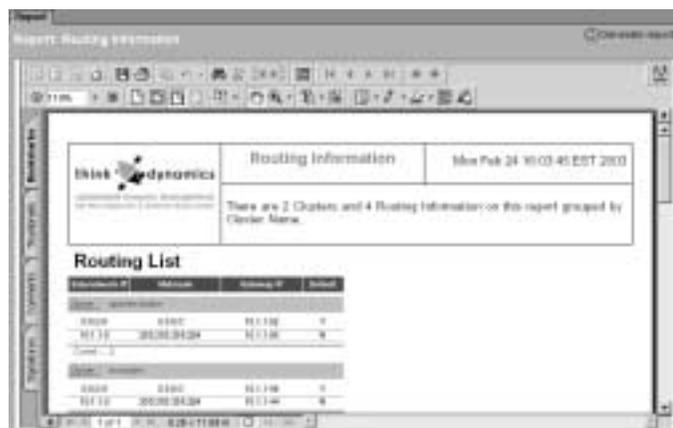


The following data is displayed:

Data	Description
Name	The network interface's name.
IP Address	The IP address for the network interface.
NIC	The NIC for the network interface.
Subnetwork	The subnetwork IP address for the network interface.
Managed	An indicator of whether the network interface is intended for management purposes.
Failed	An indicator of whether or not the network interface is handled by the Fault Management System.

Routing Information

This report provides routing information per cluster.

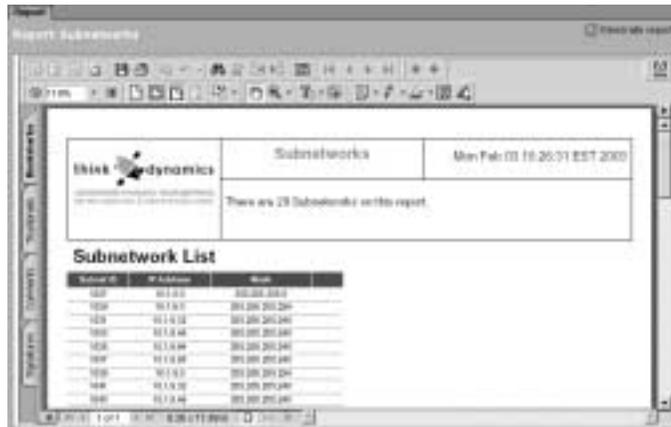


The following data is displayed:

Data	Description
Subnetwork IP	The subnetwork's IP address.
Netmask	The network's mask.
Gateway IP	The IP address of the gateway.
Default	An indicator of whether the gateway for this route is the default one.

Subnetworks

This report provides data on all the subnetworks currently managed by the system.



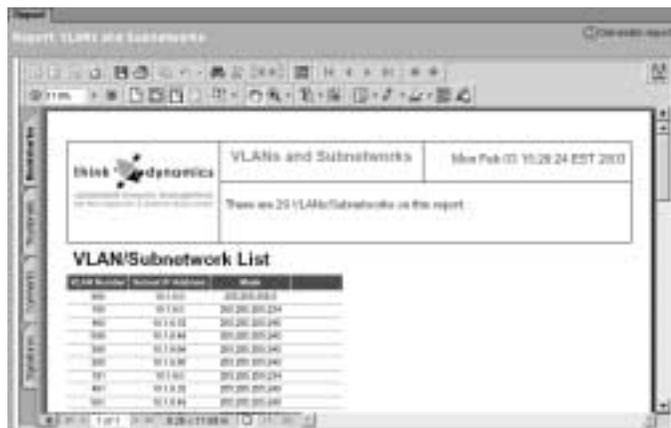
Subnet ID	IP Address	Mask
100	10.0.0.0	255.255.255.0
101	10.0.0.0	255.255.255.252
102	10.0.0.0	255.255.255.248
103	10.0.0.0	255.255.255.240
104	10.0.0.0	255.255.255.224
105	10.0.0.0	255.255.255.192
106	10.0.0.0	255.255.255.160
107	10.0.0.0	255.255.255.128
108	10.0.0.0	255.255.255.96
109	10.0.0.0	255.255.255.64

The following data is displayed:

Data	Description
Subnet ID	The subnetwork's numerical identifier.
IP Address	The subnetwork's IP address.
Mask	The subnetwork's mask.

VLANs and Subnetworks

This report provides data on all of the VLANs and subnetworks currently managed by the system.



VLAN Number	Subnet IP Address	Mask
100	10.0.0.0	255.255.255.0
101	10.0.0.0	255.255.255.252
102	10.0.0.0	255.255.255.248
103	10.0.0.0	255.255.255.240
104	10.0.0.0	255.255.255.224
105	10.0.0.0	255.255.255.192
106	10.0.0.0	255.255.255.160
107	10.0.0.0	255.255.255.128
108	10.0.0.0	255.255.255.96
109	10.0.0.0	255.255.255.64

The following data is displayed:

Data	Description
VLAN Number	The VLAN number.
Subnet IP Address	The IP address for the subnetwork.

Data	Description
Mask	The subnetwork's mask.

Usage

Recommendations

This report shows the recommendations made and conditions that resulted in the decisions. The conditions include the number of servers available in the pool, the number of servers used by the cluster, the application priority, the CPU utilization, and an Estimated Breach Probability (EBP) before and after the recommendation.

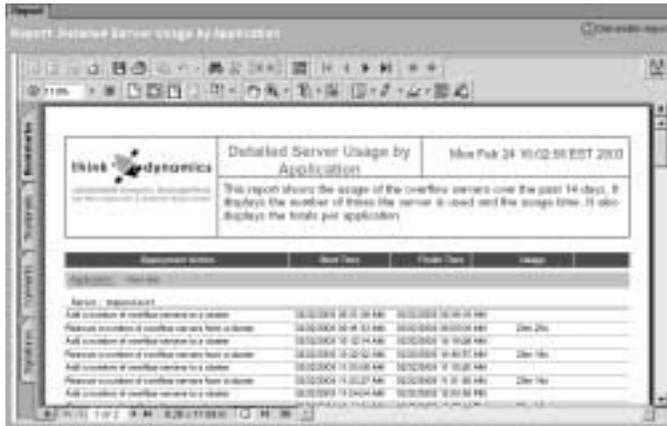


The following data is displayed:

Data	Description
Recommendation	The recommendation type, for example, add or remove server.
Delta	The number of deployed servers. Negative values are used for the removed servers.
Available	The number of servers available in the pool.
Cluster Name	The name of the cluster that the deployed servers belong to.
Servers Used	The number of used servers.
CPU Utilization	The CPU utilization.
EBP Before	The estimated breach probability before the recommendation
EBP After	The estimated breach probability after the recommendation.
Creation Time	The recommendation's date and time.

Detailed Server Usage by Application

This report shows the usage of the overflow servers over the past fourteen days. It displays the number of times the server is used, the usage time, and the totals per application.



The following data is displayed:

Data	Description
Deployment Action	The type of deployment action: overflow servers <i>added</i> to or <i>removed</i> from the application.
Start Time	The date and start time for the deployment action.
Finish Time	The date and end time for the deployment action.
Usage	The usage time per deployment action.
Number of Usages	Total number of usages per overflow server.
Usage Time	Total usage time per overflow server.
Total Number of Usages	Total number of overflow server usages per application.
Total Usage Time	Total usage time per application.

Server Usage Summary by Application

This report summarizes the usage of the overflow servers over the past fourteen days. For each server, the report displays the number of times the server is used, the usage time, the percentage of usage, maximal usage time, minimal usage time, and average usage time. Also displayed are the total number of overflow server usages and the total usage time per application.

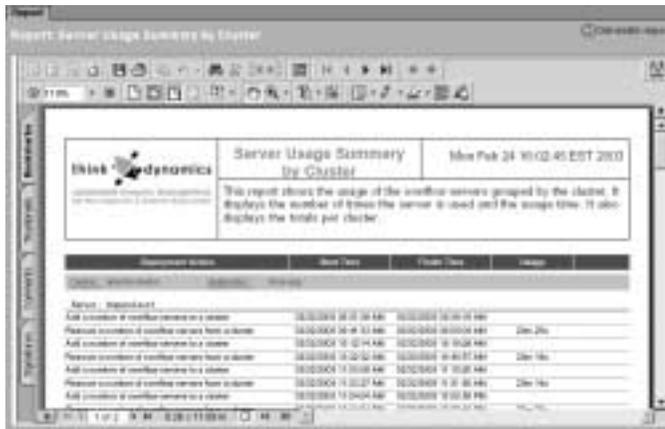


The following data is displayed:

Data	Description
Number of Usages	The number of times the overflow server has been used.
Usage Time	The usage time for the overflow server.
Percentage Used	The usage percentage for the overflow server.
Maximal Usage	The maximal usage time for the overflow server.
Minimal Usage	The minimal usage time for the overflow server.
Average Usage	The average usage time for the overflow server.
Total Number of Usages	The total number of overflow server usages per application.
Total Usage Time	The total usage time per application.
Percentage Used	The total percentage for overflow server usage per application.

Server Usage Summary by Cluster

This report shows the usage of the overflow servers, grouped by cluster. For each overflow server, the report displays the number of times the overflow server is used and the usage time. Also displayed are the total number of overflow server usages and the total usage time per cluster.



The following data is displayed:

Data	Description
Deployment Action	The type of deployment action: overflow servers <i>added</i> to or <i>removed</i> from a cluster.
Start Time	The date and start time for the deployment action.
Finish Time	The date and end time for the deployment action.
Usage	The usage time per deployment action.
Number of Usages	Total number of usages per overflow server.
Usage Time	Total usage time per overflow server.
Total Number of Usages	Total number of overflow server usages per cluster.
Total Usage Time	Total usage time per cluster.

Server Usage Summary by Pool

This report summarizes the usage of the overflow servers over the past fourteen days. For each overflow server, the report displays the number of times the server is used, the usage time, the percentage of usage, the maximal and minimal usage time, and the average usage time. Also displayed are the total number of overflow server usages and the total usage time per pool.

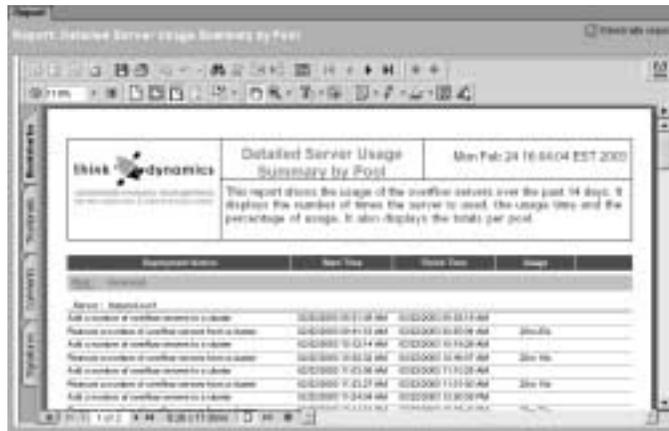


The following data is displayed:

Data	Description
Number of Usages	The number of times the overflow server has been used.
Usage Time	The usage time for the overflow server.
Percentage Used	The usage percentage for the overflow server.
Maximal Usage	The maximal usage time for the overflow server.
Minimal Usage	The minimal usage time for the overflow server.
Average Usage	The average usage time for the overflow server.
Total Number of Usages	The total number of overflow server usages per pool.
Total Usage Time	The total usage time per pool.
Percentage Used	The total percentage for overflow server usage per pool.

Detailed Server Usage Summary by Pool

This report shows the usage of the overflow servers over the past fourteen days. For each overflow server, the report displays the number of times the server is used, the usage time, and the percentage of usage. Also displayed are the total number of overflow server usages and the total usage time per pool.



The following data is displayed:

Data	Description
Deployment Action	The type of deployment action: overflow servers <i>added</i> to or <i>removed</i> from a cluster.
Start Time	The date and start time for the deployment action.
Finish Time	The date and end time for the deployment action.
Usage	The usage time per deployment action.
Number of Usages	Total number of usages per overflow server.
Usage Time	Total usage time per overflow server.
Percentage Used	Total percentage used per overflow server.
Total Number of Usages	The total number of overflow server usages per pool.
Total Usage Time	The total usage time per pool.
Percentage Used	The total percentage of overflow server usage per pool.

Chapter 7. SOAP commands available through command line

The Tivoli Intelligent Orchestrator provides a command-line interface for the data center operators who prefer to access the system's internal properties (attributes) and operations (methods) using the command line. This chapter provides reference information for the SOAP scripts that can be used with the command-line interface, and for the SOAP commands that can be used with the Web-based interface, through the **Command** box.

The Web interface allows short-names for the WebUI services, whereas the soapcli command line requires the full WDSL URL. The following URL can be used to retrieve a list of supported SOAP services and URLs:

```
http://hostname:port/tcSoap/wsd1
```

Note: For more information about the terms and abbreviations that are used in this chapter, refer to the "Glossary" provided later in this guide.

The following information is included:

- SOAP scripts available through the command-line interface
- SOAP commands available through the Web-based interface
- SOAP commands initiating device functions

SOAP scripts available through the command-line interface

The \$HOME/tools/soap directory contains a number of scripts, .cmd files for Windows environments and .sh files for UNIX environments, that execute predefined SOAP commands.

To execute these commands, you should:

1. Log on to the Tivoli Intelligent Orchestrator server as user **thinkcontrol**.
2. Open a Command prompt or Telnet window.
3. Click the tools/soap directory.
4. Run the script with specific parameters, as specified in the following table.

The commands can also be run on any machine other than the Tivoli Intelligent Orchestrator server. To do this, the SOAP client must include a set of lightweight java libraries and commands that can be used to invoke the service from any client machine.

The required libraries are:

- /j2ee.jar
- /mail.jar
- /qname.jar
- /soap.jar
- /wsdl4j.jar
- /wsif.jar
- /xerces.jar
- /tccommandline.jar

Use the following Web address to download the WSDL file from the Tivoli Intelligent Orchestrator server: `http://hostname:port/tcSoap/wSDL`.

If you invoke a SOAP client from any machine other than the Tivoli Intelligent Orchestrator server, you need to make sure you modify the endpoint of the WSDL to point to the Tivoli Intelligent Orchestrator server host rather than the local host. For details, see `XXX_Service.wsdl`.

Note: For details on the specific syntax of some of the available SOAP commands, refer to the “SOAP command syntax reference” appendix, later in this guide.

Table 2. SOAP scripts available through the command-line interface

Script (.sh or .cmd)	Parameters	Description
approve	<username> <password> <recommendation-id>	Approves the specified Resource Broker recommendation, while in semiautomatic mode.
auto	<username> <password>	Sets the global operating mode to automatic.
deploy	<username> <password> <cluster-id> <number-of-servers>	Adds or removes a specified number of servers to or from a cluster. You should use negative numbers for server removal.
failserver	<username> <password> <server-id>	Notifies the Tivoli Intelligent Orchestrator about server failures.
manual	<username> <password>	Sets the global operating mode to manual.
mode	<username> <password>	Queries the current global operating mode.
repairserver	<username> <password> <server-id>	Notifies the Tivoli Intelligent Orchestrator about repaired servers, as opposed to <code>failserver</code> .
semi	<username> <password>	Sets the global operating mode to semiautomatic.
soapcli	<username> <password> <wsdl location> <method> <parameters>	A basic SOAP invocation script. The WSDL location, the methods, and the parameters are described in the next section.

Most of the information in brackets, for example, <server-id>, <cluster-id>, and <lbalancer-id>, can be obtained through the Tivoli Intelligent Orchestrator’s Web-based interface, by holding the mouse cursor over the appropriate device. In the case of the <recommendation-id>, it can be obtained on the Deployments screen.

For example:

```
./failserver.sh tc pwd 1061
```

where `tc` is the <username>, `pwd` is the <password>, and `1061` is the <server-id>.

SOAP commands available through the Web-based interface

The **Command** box in the Web-based user interface accepts commands using the following syntax:

```
<WSDL location> <method-name> <param1> <param2> ...
```

Note: For null parameters, you should use “null” as parameter value. For example, in the following command, “null” is used to specify that the port module parameter is null.

```
http://hostname:port/tcSoap/wsd1/com/ibm/tc/soap/OperationsModeService.wsd1 movePortToVLAN 1039 1032 1034 null 11
```

The currently supported WSDL locations are presented in the following table and described in more detail in the following sections.

Table 3. WSDL locations available through the Web-based interface

WSDL location	Description
faultmanagement	Enables you to access commands for the Fault Management system.
OperationsModeService	Enables you to access commands for the Operations Mode logic.
recommendations	Enables you to access commands for creating and approving deployment requests.

You can access each of these WSDLs only if you are granted the corresponding access privileges. For more information on access privileges, refer to the “User Access Privilege Descriptions” section in the “Configuring Tivoli Intelligent Orchestrator” chapter, earlier in this guide.

FaultManagementService.wsd1

This WSDL location enables you to access commands for the Fault Management System. To access it, you need the **fault management** access privilege.

Table 4. SOAP commands for FaultManagementService.wsd1

Method	Parameters	Description
serverEvent	<server-id> “true” or “false”	Sets or clears the failed flag for the specified server.
adminServer-Event	<cluster-id> “true” or “false”	Sets or clears the failed flag for the specified cluster administrative server.
loadBalancer-Event	<lbalancer-id> “true” or “false”	Sets or clears the failed flag for the specified load balancer.
networkInterfaceEvent	<network-interface-id> “true” or “false”	Sets or clears the failed flag for the specified network interface.
switchEvent	<switch-id> “true” or “false”	Sets or clears the failed flag for the specified switch.

OperationsModeService.wsdl

This WSDL location enables you to access commands for the Operations Mode logic. To access it, you need the **operating mode query** and **operating mode change** access privileges.

Note: For more details on operating modes, refer to the “OperationsModeService.wsdl” section in *Appendix A*, described later in this guide.

Table 5. SOAP commands for OperationsModeService.wsdl

Method	Parameters	Description
getApplication-Mode	<app-id>	Returns the current operating mode for the specified application.
getClusterMode	<cluster-id>	Returns the current operating mode for the specified cluster.
getEffectiveApplicationMode	<app-id>	Returns the effective operating mode for the specified application.
getEffectiveClusterMode	<cluster-id>	Returns the effective operating mode for the specified cluster.
getGlobalMode		Returns the current global operating mode.
setMaintenance-Mode	<objectId> <typeId> <enable> <comment>	Sets the specified object to be in or out of maintenance mode. The object can be any Data Center Model (DCM) entity, including a switch, a server, or an entire application.
setApplication-MaintenanceMode	<app-id> <enable>	Sets the specified application to be in or out of maintenance mode.
setClusterMode	<cluster-id> <new-mode>	Sets the specified cluster to the specified operating mode.
setApplication-Mode	<app-id> <new-mode>	Sets the specified application to the specified operating mode.
setGlobalMode	<new-mode>	Sets the current global operating mode to a new mode.

The possible values of the <new-mode> argument are: **automatic**, **manual**, **semiautomatic**, and **default**.

RecommendationsService.wsdl

This WSDL location enables you to access commands for creating and approving deployment requests. It generates a *Recommendation* object, which causes deployment requests to be sent to the Deployment Engine. To access this WSDL location, you need the **recommendation query** and **recommendation change** access privileges.

Table 6. SOAP commands for RecommendationsService.wsdl

Method	Parameters	Description
addRepairedServer	<cluster-id> <server-id> <comment>	Adds the specified repaired server back to its cluster.
approveRecommendation	<recomm-id> <approved-by> <comment>	Approves a recommendation issued in <i>semiautomatic</i> mode.
declineRecommendation	<recomm-id> <approved-by> <comment>	Declines a recommendation issued in <i>semiautomatic</i> mode.
deploy	<cluster-id> <num-servers> <approve-by> <comment>	Adds or removes the specified number of servers to/from a cluster. Valid if cluster is in <i>manual</i> mode.
removeFailedServer	<cluster-id> <server-id> <comment>	Removes the specified failed server from the specified cluster.
undeploySpecificServer	<cluster-id> <server-id> <approved-by> <comment>	Removes the specified server from the specified cluster.

SOAP commands initiating device functions

A number of SOAP commands are implemented to initiate any of the device driver workflows using the command line. These commands are available through both the Web-based interface and the command-line interface. Each SOAP command offers access to a logical command that allows you to execute a specific deployment on one of the logical devices supported by the system.

Note: For details on the specific syntax to be used for accessing these SOAP commands through either the Web-based user interface or the command-line interface, refer to the previous sections.

The currently supported WSDL locations are presented in the following table, and described in more detail in the following sections.

Table 7. WSDL locations for initiating device functions

WSDL location	Description
boot-server-component	Offers access to logical commands for deployments on boot servers.
cluster-component	Offers access to logical commands for deployments on clusters.

Table 7. WSDL locations for initiating device functions

WSDL location	Description
device-component	Offers access to logical commands for initiating devices.
file-repository-component	Offers access to logical commands for managing file repository systems.
firewall-component	Offers access to logical commands for setting up firewall configurations.
ip-system-component	Offers access to logical commands for configuring the IP systems.
load-balancer-component	Offers access to logical commands for deployments on load balancers.
operating-system-component	Offers access to logical commands controlling operating systems.
power-unit-component	Offers access to logical commands controlling power units.
router-component	Offers access to logical commands for deployments on routers.
software-component	Offers access to logical commands for software deployments on managed devices.
software-stack-component	Offers access to logical commands for software stack deployments.
spare-pool-component	Offers access to logical commands for deployments on spare pools.
switch-component	Offers access to logical commands for deployments on switches.
switch-fabric-component	Offers access to logical commands for deployments on switch fabrics.

You can access each of these WSDLs only if you are granted the corresponding access privileges. For more information on access privileges, refer to the “User Access Privilege Descriptions” section in the “Configuring Tivoli Intelligent Orchestrator” chapter, earlier in this guide.

BootServerComponentService.wsdl

This WSDL location enables you to access logical commands for deployments on boot servers. To access it, you need the **Boot Server interaction** access privilege.

Table 8. SOAP commands for BootServerComponentService.wsdl

Method	Parameters	Description
installImage	<bootServerId> <software-StackImageId> <destinationDeviceId>	Installs the specified software stack image on the specified boot server or destination device, and updates the DCM to reflect the new IP address and VLAN of this device.
uploadImage	<bootServerId> <software-StackImageId>	Uploads the specified software stack image on the specified boot server, and updates the DCM to reflect the change.

ClusterComponentService.wsdl

This module enables you to access logical commands for deployments on clusters. To access this module, you need the **Cluster interaction** access privilege.

Table 9. SOAP commands for ClusterComponentService.wsdl

Method	Parameters	Description
addServer	<clusterId> <serverId>	Adds the specified server to the specified cluster. If a server ID is specified, the server must be locked before calling this method. If the server ID is not specified, the workflow calls the <i>RM Allocate Server</i> driver to allocate a server, which also locks the server.
removeServer	<clusterId> <serverId>	Removes the specified server from the specified cluster. If a server ID is specified, the server must be locked before calling this method. If the server ID is not specified, the workflow calls the <i>RM Choose Server for Removal</i> driver to choose which server to be removed, also locking the server.
addRepaired-Server	<clusterId> <serverId>	Adds the specified repaired server to the specified cluster.
remove-FailedServer	<clusterId> <serverId>	Removes the specified faulty server from the specified cluster.

DeviceComponentService.wsdl

This WSDL location enables you to access logical commands for initiating devices. To access it, you need the **Device interaction** access privilege.

Table 10. SOAP commands for DeviceComponentService.wsdl

Method	Parameters	Description
initialize	deviceId	Initializes the specified device.
hardware-Reboot	deviceId	Performs a hardware reboot on the specified device. If no <i>Device.Hardware Reboot</i> workflow is associated with this device or its device driver, the default hardware reboot workflow is executed, which turns the power supply, or supplies, off and on for this device.
software-Reboot	deviceId	Performs a software reboot on the specified device.
powerOff	deviceId	Powers off the specified device, by turning off the power supply, or supplies, for that device.
powerOn	deviceId	Powers on the specified device, by turning on the power supply, or supplies, for that device.

FileRepositoryComponentService.wsdl

This WSDL location enables you to access logical commands for managing file repository systems. To access it, you need the **File Repository interaction** access privilege.

Table 11. SOAP commands for FileRepositoryComponentService.wsdl

Method	Parameters	Description
getFile	<fileRepositoryId> <sourcePath> <sourceFileName> <destIPAddress> <destPath> <destFileName>	Retrieves a file from the specified remote file repository system.
putFile	<fileRepositoryId> <sourceIPAddress> <sourcePath> <sourceFileName> <destIPAddress> <destPath> <destFileName>	Uploads a file to the specified remote file repository system.
removeFile	<fileRepositoryId> <path> <fileName>	Removes a file from the specified remote file repository system.

FirewallComponentService.wsdl

This WSDL location enables you to access logical commands for setting up firewall configurations. To access it, you need the **Firewall interaction** access privilege.

Table 12. SOAP commands for FirewallComponentService.wsdl

Method	Parameters	Description
addACL	<firewallId> <aclId>	Adds the specified Access Control List (ACL) to the specified firewall, and updates the DCM to reflect the addition of the new ACL to the firewall configuration.
enableACL	<firewallId> <destinationNetworkInterfaceId> <aclId> <type>	Specifies whether the ACL is enabled on the specified firewall.
disableACL	<firewallId> <destinationNetworkInterfaceId> <aclId>	Specifies whether the ACL is disabled on the specified firewall.
removeACL	<firewallId> <aclId>	Removes the specified ACL from the specified firewall, and updates the DCM to reflect the removal of the ACL from the firewall configuration.

IpSystemComponentService.wsdl

This WSDL location enables you to access logical commands for configuring the IP systems. To access it, you need the **IP System interaction** access privilege.

Table 13. SOAP commands for *IpSystemComponentService.wsdl*

Method	Parameters	Description
addIP-Address	<deviceId> <operatingSystemId> <nicId> <destinationSubnetworkId> <networkInterfaceName> <ipAddress> <defaultGateway> <isManaged>	Adds the specified IP address (network interface) to the specified IP network device. If no <i>IPSystem.Add IP Address</i> workflow is associated with this device or its device driver, it forwards the request to the <i>addIPAddress</i> method in the <i>OperatingSystemComponentService.wsdl</i> . It also updates the DCM to reflect the addition of the new network interface to this device.
apply-Routing-Table	<deviceId> <operatingSystemId> <routingTableId> <networkInterfaceId> <removeExistingRoutes>	Applies the specified routing table to the specified IP network device. If no <i>IPSystem.Apply Routing Table</i> workflow is associated with this device or its device driver, it forwards the request to the <i>applyRoutingTable</i> method in the <i>OperatingSystemComponentService.wsdl</i> . It also updates the DCM to reflect the new routing table for this device.
removeIP-Address	<deviceId> <operatingSystemId> <destinationSubnetworkId> <ipAddress>	Removes the specified IP address (network interface) from the specified IP network device. If no <i>IPSystem.Remove IP Address</i> workflow is associated with this device or its device driver, it forwards the request to the <i>removeIPAddress</i> method in the <i>OperatingSystemComponentService.wsdl</i> . It also updates the DCM to reflect the removal of the network interface from this device.

LoadBalancerComponentService

This WSDL location enables you to access logical commands for deployments on load balancers. To access it, you need the **Load Balancer interaction** access privilege.

Table 14. SOAP commands for *LoadBalancerComponentService.wsdl*

Method	Parameters	Description
addServer	<loadBalancerId> <serverId> <destinationVIPId> <port> <protocol> <weight> <maxConnections>	Adds the specified server to the specified virtual IP, and updates the DCM to reflect the new real IP in the target virtual IP.
removeServer	<loadBalancerId> <serverId> <sourceVIPId>	Removes the specified server from the specified virtual IP, and updates the DCM to reflect the removal of the real IP from the target virtual IP.

Table 14. SOAP commands for LoadBalancerComponentService.wsdl

Method	Parameters	Description
createVirtualIP	<loadBalancerId> <clusterId> <virtualIpName> <virtualIpAddress> <inTcpPortFirst> <inTcpPortLast> <outTcpPort> <balancingAlgorithm>	Creates a virtual IP on the specified load balancer, and updates the DCM to reflect the addition of the new virtual IP.
removeVirtualIP	<loadBalancerId> <vipId>	Removes the specified virtual IP from the specified load balancer, and updates the DCM to reflect the removal of the virtual IP. The returned value is the numerical identifier of the request that triggered the workflow execution for removing the VIP.
AddRealIPToVirtualIP	<LoadBalancerId> <NetworkInterfaceId> <Port> <Protocol> <Weight> <MaxConnections> <DestinationVIPId>	Adds a real IP to a virtual IP on the specified load balancer.
RemoveRealIPFromVirtualIP	<LoadBalancerId> <RealIPId>	Removes a real IP from the virtual IP on the specified load balancer.

OperatingSystemComponentService.wsdl

This WSDL location enables you to access logical commands for controlling operating systems. To access it, you need the **Operating System interaction** access privilege.

Table 15. SOAP commands for OperatingSystemComponentService.wsdl

Method	Parameters	Description
addIPAddress	<deviceId> <operatingSystemId> <nicId> <destinationSubnetworkId> <networkInterfaceName> <ipAddress> <defaultGateway> <isManaged>	Adds an IP address (network interface) to the specified IP network device, and updates the DCM to reflect the addition of the new network interface to this device.
applyRoutingTable	<deviceId> <operatingSystemId> <routingTableId> <networkInterfaceId> <removeExistingRoutes>	Applies the specified routing table to the specified IP network device, and updates the DCM to reflect the new routing table for this device.
removeIP-Address	<deviceId> <operatingSystemId> <destinationSubnetworkId> <ipAddress>	Removes an IP address (network interface) from the specified IP network device, and updates the DCM to reflect the removal of the network interface from this device.

Table 15. SOAP commands for *OperatingSystemComponentService.wsdl*

Method	Parameters	Description
software-Reboot	<deviceId>	Reboots (software) the specified device.

PowerUnitComponentService.wsdl

This WSDL location enables you to access logical commands for controlling power units. It allows you to access logical commands for remotely turning on/off power outlets. Each device can be connected to one or more power outlets, which can be plugged into separate power units.

To access this WSDL location, you need the **Power Unit interaction** access privilege.

Table 16. SOAP commands for *PowerUnitComponentService.wsdl*

Method	Parameters	Description
turnOutletOff	<powerUnitID> <outlet>	Turns off the specified outlet on the specified power unit.
turnOutletOn	<powerUnitID> <outlet>	Turns back on the specified outlet on the specified power unit.

RouterComponentService.wsdl

This WSDL location enables you to access logical commands for deployments on routers. To access it, you need the **Router interaction** access privilege.

Table 17. SOAP commands for *RouterComponentService.wsdl*

Method	Parameters	Description
createRoute	<routerId> <networkInterfaceId> <subnetworkId> <gateway>	Adds the specified route to the specified router, and updates the DCM to reflect the addition of the new route.
removeRoute	<routerId> <routeId>	Removes the specified route from the specified router, and updates the DCM to reflect the route removal.

SoftwareComponentService.wsdl

This WSDL location enables you to access logical commands for software deployments on managed devices. To access it, you need the **Software interaction** access privilege.

Table 18. SOAP commands for *SoftwareComponentService.wsdl*

Method	Parameters	Description
checkStatus	<softwareId> <deviceId>	Queries the status of the specified software product on the specified device. The workflow updates the DCM status of the software product on the target device.
install	<softwareId> <deviceId>	Installs the specified software product on the specified device, and updates the DCM to reflect the installation of the software product on the target device.

Table 18. SOAP commands for *SoftwareComponentService.wsdl*

Method	Parameters	Description
start	<softwareId> <deviceId>	Starts the specified software product on the target device, and updates the DCM to reflect the software product's status on this device.
stop	<softwareId> <deviceId>	Stops the specified software product on the target device, and updates the DCM to reflect the software product's status on the this device.
uninstall	<softwareId> <deviceId>	Uninstalls the specified software product from the target device, and updates the DCM to reflect the removal of the software product from this device.

SoftwareStackComponentService.wsdl

This WSDL location enables you to access logical commands for controlling software stacks. To access it, you need the **Software Stack interaction** access privilege.

Table 19. SOAP commands for *SoftwareStackComponentService.wsdl*

Method	Parameters	Description
install	<softwareStackId> <deviceId>	Installs the specified software stack on the specified device. If no <i>SoftwareStack.Install</i> workflow is associated with this device or its device driver, the default software stack <i>install</i> workflow is used, which installs each software product in the software stack. The DCM is updated to reflect the new status of all software products in the target device's stack.
uninstall	<softwareStackId> <deviceId>	Uninstalls the specified software stack from the specified device. If no <i>SoftwareStack.Uninstall</i> workflow is associated with this device or its device driver, the default software stack <i>uninstall</i> workflow is used, which uninstalls each software product in the software stack. The DCM is updated to reflect the new state of all software products in the target device's stack.

SparePoolComponentService.wsdl

This WSDL location enables you to access logical commands for deployments on spare pools. To access it, you need the **Spare Pool interaction** access privilege.

Table 20. SOAP commands for *SparePoolComponentService.wsdl*

Method	Parameters	Description
cleanup- Server	<sparePoolId> <serverId>	Cleans up the specified server from the specified spare pool before initializing it. If no <i>SparePool.Cleanup Server</i> workflow is associated with the resource pool or its device driver, it forwards the request to the <code>uninstall</code> method in the <i>SoftwareStackComponentService.wsdl</i> , using the software stack associated with the resource pool.

Table 20. SOAP commands for *SparePoolComponentService.wsdl*

Method	Parameters	Description
initialize-Server	<sparePoolId> <serverId>	Initializes the specified server to match the required configuration for servers in a resource pool. If no <i>SparePool.Initialize Server</i> workflow is associated with the resource pool or its device driver, it forwards the request to the <code>install</code> method in the <i>SoftwareStackComponentService.wsdl</i> , using the software stack associated with the resource pool.

SwitchComponentService.wsdl

This WSDL location enables you to access logical commands for deployments on switches. To access it, you need the **Switch interaction** access privilege.

Table 21. SOAP commands for *SwitchComponentService.wsdl*

Method	Parameters	Description
createVLAN	<switchId> <vlanName> <vlanNumber> <subnetworkId>	Adds the specified VLAN to the specified switch, and updates the DCM to reflect the addition of the new VLAN to the target switch.
movePortTo-VLAN	<switchId> <sourceVLANId> <destinationVLANId> <portModule> <portNumber>	Moves the specified port from the source VLAN to the destination VLAN, and updates the DCM to reflect the new switch port to the VLAN association.
removeVLAN	<switchId> <vlanId>	Removes the specified VLAN from the specified switch, and updates the DCM to reflect the VLAN removal from the target switch.
turnPortOff	<switchId> <portModule> <portNumber>	Turns off the specified port on the target switch, and updates the DCM to reflect the new status of the switch port.
turnPortOn	<switchId> <portModule> <portNumber>	Turns the specified port on, and updates the DCM to reflect the new status of the switch port.

SwitchFabricComponentService.wsdl

This WSDL location enables you to access logical commands for deployments on switch fabrics. To access it, you need the **Switch Fabric interaction** access privilege.

Table 22. SOAP commands for *SwitchFabricComponentService.wsdl*

Method	Parameters	Description
createVLAN	<switchFabricId> <vlanName> <vlanNumber> <subnetworkId>	Creates the specified VLAN on all switches in the specified switch fabric. It calls the <code>createVLAN</code> method in the <i>SwitchComponentService.wsdl</i> for each switch in the switch fabric.
removeVLAN	<switchFabricId> <vlanId>	Removes the specified VLAN from all switches in the specified switch fabric. It calls the <code>removeVLAN</code> method in the <i>SwitchComponentService.wsdl</i> for each switch in the switch fabric.

Appendix A. SOAP command syntax reference

Certain Tivoli Intelligent Orchestrator internal properties and methods are available through implemented SOAP commands, which can be used with both the Web-based interface and the command-line interface. This appendix provides a reference to some of the SOAP commands that can be used through the Tivoli Intelligent Orchestrator's Web-based user interface, describing the command syntax and providing relevant examples.

The following WSDL locations are available through the Tivoli Intelligent Orchestrator's command line:

- FaultManagementService.wsdl
- OperationsModeService.wsdl
- RecommendationsService.wsdl
- MessageTranslator.wsdl

The OperationsModeService.wsdl and MessageTranslator.wsdl are described in detail in the following section.

OperationsModeService.wsdl

The *OperationsModeService.wsdl* can set the operating mode of the Tivoli Intelligent Orchestrator (the global operating mode) and the operating mode of any application or cluster under the Tivoli Intelligent Orchestrator's management using the following strings:

- *automatic*
- *manual*
- *semiautomatic*
- *default*

The operating mode affects the deployment of recommendations generated by the Resource Broker component and the Fault Management system as follows:

<i>manual</i> mode	Recommendations are generated but not acted upon by the Deployment Engine. The operator can review the recommendations, but must explicitly execute the resource reallocations manually.
<i>semiautomatic</i> mode	Recommendations are generated with null <i>authorized-by</i> field. The operator can view pending requests and must authorize a request before the Deployment Engine will act upon it.
<i>automatic</i> mode	Recommendations are automatically authorized by the Tivoli Intelligent Orchestrator and are subsequently acted upon by the Deployment Engine. For stability reasons, additional deployment requests for a given cluster will not be initiated while other deployment requests are pending.

The operating mode can be specified at the global, application, or cluster level. Specifying the operating mode for a cluster is possible only through the command-line interface. Setting a cluster's operating mode to *default* will cause it to use its parent application's operating mode. Setting an application's operating mode to *default* will cause it to use the global operating mode.

In addition to the above three modes, one additional attribute that affects the operating mode for specific applications, is as follows:

Maintenance

This attribute allows the operator to “disconnect” an application from the Tivoli Intelligent Orchestrator so that maintenance activities can be performed. When enabled, the Resource Broker treats all of the clusters in the application as if they are in *manual* mode, and the operator is free to make changes to the database entries for all components of that application.

Note: The *Syntax* and *Examples* sections outlined in the following methods are specific to the command line interface.

getApplicationMode

This method retrieves the current operating mode of the specified application.

Syntax

```
soapcli <username> <password> <WSDL location> getApplicationMode
<app-id>
```

Parameters

The following parameter is used:

Parameter	Description
app-id	The identifier of the specified application

Returned value

The value of the current operating mode for the specified application. It can be one of the following:

- **automatic:** The system makes all the resource change decisions pertaining to this application, and automatically executes the necessary allocations.
- **semiautomatic:** The system identifies the need for resource changes for the specified application, and issues requests that must be approved by the operator.
- **manual:** The operator has full control over the application's resource needs, and explicitly adds and removes servers manually.
- **default:** The application uses the current global operating mode.

Example

```
soapcli tc pwd
http://hostname:port/tcSoap/wsd1/com/ibm/tc/soap/OperationsModeService.wsd1
OperationsModeService getApplicationMode 1208
```

getClusterMode

This method retrieves the operating mode of the specified cluster.

Syntax

```
soapcli <username> <password> <WSDL location> getClusterMode  
<cluster-id>
```

Parameters

The following parameter is used:

Parameter	Description
cluster-id	The identifier of the specified cluster

Returned value

The value of the current operating mode for the specified cluster.

It can be one of the following:

- **automatic:** The system makes all the resource change decisions pertaining to the cluster, and automatically executes the necessary allocations.
- **semiautomatic:** The system identifies the need for resource changes for the specified cluster, and issues requests that must be approved by the operator.
- **manual:** The operator has full control over the cluster's resource needs, and explicitly adds and removes servers manually.
- **default:** The cluster uses the parent application's operating mode.

Example

```
soapcli tc pwd  
http://hostname:port/tcSoap/wsd/com/ibm/tc/soap/OperationsModeService.wsd  
l getClusterMode 1225
```

getEffectiveApplicationMode

This method calculates the effective operating mode of a specified application. If the application's operating mode is set to **default**, the global operating mode is applied.

Syntax

```
soapcli <username> <password> <WSDL location>  
getEffectiveApplicationMode <app-id>
```

Parameters

The following parameter is used:

Parameter	Description
app-id	The identifier of the specified application

Returned value

The value of the effective operating mode for the specified application.

It can be one of the following:

- **automatic:** The system makes all the resource change decisions pertaining to the application, and automatically executes the necessary allocations.

- **semiautomatic:** The system identifies the need for resource changes for the specified application, and issues requests that must be approved by the operator.
- **manual:** The operator has full control over the application's resource needs, and explicitly adds and removes servers manually.
- **failure:** A failure has been detected; therefore, the application will behave as if in **manual** mode.
- **maintenance:** The application has been removed from the system's control and is available for reconfiguration.

Example

```
soapcli tc pwd
http://hostname:port/tcSoap/wsd1/com/ibm/tc/soap/OperationsModeService.wsd1 getEffectiveApplicationMode 1208
```

getEffectiveClusterMode

This method calculates the effective operating mode of a specified cluster. If the cluster's operating mode is set to **default**, the operating mode of the parent application is used. If the application's operating mode is set to **default**, the global operating mode is used.

Syntax

```
soapcli <username> <password> <WSDL location>
getEffectiveClusterMode <cluster-id>
```

Parameters

The following parameter is used:

Parameter	Description
cluster-id	The identifier of the specified cluster

Returned value

The value of the current operating mode for the specified cluster.

It can be one of the following:

- **automatic:** The system makes all the resource change decisions pertaining to the cluster, and automatically executes the necessary allocations.
- **semiautomatic:** The system identifies the need for resource changes for the specified cluster, and issues requests that must be approved by the operator.
- **manual:** The operator has full control over the cluster's resource needs, and explicitly adds and removes servers manually.
- **failure:** A failure has been detected, therefore the cluster will behave as if in **manual** mode.
- **maintenance:** The cluster has been removed from the system's control and is available for reconfiguration.

Example

```
soapcli tc pwd
http://hostname:port/tcSoap/wsd1/com/ibm/tc/soap/OperationsModeService.wsd1 getEffectiveClusterMode 1225
```

getGlobalMode

This method retrieves the current operating mode of the entire system. The global operating mode can override the operating modes of individual applications and clusters provided these have been initially set to **default**.

Syntax

```
soapcli <username> <password> <WSDL location> GlobalMode
```

Parameters

No parameters are required.

Returned value

The value of the current global operating mode.

Example

```
soapcli tc pwd  
http://hostname:port/tcSoap/wsd/ibm/tc/soap/OperationsModeService.wsd GlobalMode
```

setApplicationMaintenanceMode

This method enables or disables the *maintenance* attribute that affects the operating mode for the specified application.

When enabled, all of the clusters in the application behave as in *manual* mode. All subsequent calls to `setApplicationMode()`, `setClusterMode()` change the underlying operating mode state, but will not have any effect until the *maintenance* attribute is disabled. When disabled, all of the clusters in the application have their operating mode restored to the “pre-maintenance mode” values.

The Application Controller will stop all data acquisition from all the components associated with the application in maintenance.

Disabling the *maintenance* attribute causes the Resource Broker and Application Controller to reread the application’s configuration from the database and re initialize the necessary components.

Syntax

```
soapcli <username> <password> <WSDL location>  
setApplicationMaintenanceMode <app-id> <enable> <comment>
```

Parameters

The following parameters are used:

Parameter	Description
app-id	The identifier of the specified application
enable	“True” to enable the <i>maintenance</i> attribute, “false” to disable it
comment	Text comment to be stored in the database

Returned value

None

Example

```
soapcli tc pwd
http://hostname:port/tcSoap/wsd1/com/ibm/tc/soap/OperationsModeService.wsd1 setApplicationMaintenanceMode 1166 true test_1234
```

setApplicationMode

This method sets the operating mode for the specified application.

Syntax

```
soapcli.sh <username> <password> <WSDL location> setApplicationMode
<app-id> <new-mode>
```

Parameters

The following parameters are used:

Parameter	Description
app-id	The identifier of the specified application
new-mode	The new operating mode to be used by the specified application

Returned value

The value of the current or new operating mode for the specified application.

Example

```
soapcli tc pwd
http://hostname:port/tcSoap/wsd1/com/ibm/tc/soap/OperationsModeService.wsd1 setApplicationMode 1166 manual
```

setClusterMode

This method sets the operating mode for the specified cluster.

Syntax

```
soapcli <username> <password> <WSDL location> setClusterMode
<cluster-id> <new-mode>
```

Parameters

The following parameters are used:

Parameter	Description
cluster-id	The identifier of the specified cluster
new-mode	The new operating mode to be used by the specified cluster

Returned value

The value of the current or new operating mode for the specified cluster.

Example

```
soapcli tc pwd
http://hostname:port/tcSoap/wsd1/com/ibm/tc/soap/OperationsModeService.wsd1 setClusterMode 1168 manual
```

setGlobalMode

This method changes the current global operating mode.

Syntax

```
soapcli <username> <password> <WSDL location> GlobalMode <new-mode>
```

Parameters

The following parameter is used:

Parameter	Description
new-mode	The new global operating mode

Returned value

The value of the new global operating mode.

Example

```
soapcli tc pwd
http://hostname:port/tcSoap/wsd1/com/ibm/tc/soap/OperationsModeService.wsd1 GlobalMode automatic
```

setMaintenanceMode

This method enables or disables the *maintenance* attribute that affects the operating mode for the specified object.

Syntax

```
soapcli <username> <password> <WSDL location> setMaintenanceMode <object-id> <type-id> <enable> <comment>
```

Parameters

The following parameters are used:

Parameter	Description
object-id	The identifier of the specified object
type-id	The type of the specified object
enable	"True" to enable the <i>maintenance</i> attribute, "false" to disable it
comment	Text comment to be saved in the database

Returned value

None

Example

```
soapcli tc pwd
http://hostname:port/tcSoap/wsd1/com/ibm/tc/soap/OperationsModeService.wsd1 setMaintenanceMode 1168 2 true test_1234
```

MessageTranslator.wsdl

This WSDL location provides the createDeploymentRequest operation, that creates a new deployment request using a soap call.

Table 1. SOAP commands for MessageTranslator.wsdl

Method	Parameters	Description
createDeploymentRequest	<requestTypeId> <workflowId> <requestParameters>	Adds a new deployment request of the specified type, that will execute the specified deployment workflow.

The requestParameters parameter takes a string representation of the request parameters that you pass into the call. The request parameters supports two kinds of formats:

- An XML format: This format is useful when it is called from a Java program)

```
<?xml version="1.0" encoding="UTF-8"?>
<requestParameters>
<param name="param1" value="value1" />
<param name="param2" value="value2" />
<param name="param3" value="value3" />
</requestParameters>
```
- An delimiter format: This format is useful when it is called from a command line.
"param1=value1,param2=value2,param3=value3"

The server side of the Web Services code will determine whether or not a passed-in request parameters are having the XML format or the delimiter format. If the request parameters that are passed in have unknown format, all the request parameters passed in will be ignored. The ID of the new deployment request will be returned from this SOAP call. The following are examples of calling the method from the command line:

```
soapcli userid password
http://localhost:9080/tcSoap/wsdl/com/ibm/tc/soap/MessageTranslatorService.wsdl
createDeploymentRequest 1640 1021 " "
```

```
http://localhost:9080/tcSoap/wsdl/com/ibm/tc/soap/MessageTranslatorService.wsdl
createDeploymentRequest 1640 1021
"param1=value1,param2=value2,param3=value3"
```

The following is an example of how to integrate the SOAP call from your own Java client program by taking the requestParameters as an XML string:

```
public class MyTest
public MyTest() throws Exception {
String wsdl= "http://localhost:9080/tcSoap/wsdl/com/ibm/tc/soap/MessageTranslatorService.wsdl";
String opname= "createDeploymentRequest";
List params= new ArrayList();
params.add("1640");
params.add("1021");
String requestParams =
"<?xml version=\"1.0\" encoding=\"UTF-8\"?>" +
```

```
"<requestParameters>" +
"<param name=\"param1\" value=\"value1\"/>" +
"<param name=\"param2\" value=\"value2\"/>" +
"<param name=\"param3\" value=\"value3\"/>" +
"<param name=\"param4\" value=\"value4\"/>" +
"</requestParameters>";
params.add(requestParams);
System.out.println("result: " + Invoker.invoke(wsdl, opname, params, "use-
rid", "password"));
}
```

Appendix B. Defining the report design in an XML file: a case study

This appendix provides a practical example of an Tivoli Intelligent Orchestrator report, the **Inventory** report, which is generated based on the report design defined in an XML file. The following sections are included:

- Defining the report design in the XML file
- Displaying the generated report

For more information on aspects of this XML report template, refer to:

<http://jasperreports.sourceforge.net/tutorial/index.html>

Defining the report design in the XML file

The example below shows the XML file that is used by Tivoli Intelligent Orchestrator's reporting engine to generate the **Inventory** report:

```
<jasperReport
  name="Inventory"
  columnCount="2"
  printOrder="Vertical"
  pageWidth="595"
  pageHeight="842"
  columnWidth="270"
  columnSpacing="15"
  leftMargin="20"
  rightMargin="20"
  topMargin="30"
  bottomMargin="30"
  isTitleNewPage="false"
  isSummaryNewPage="false">
  <reportFont name="Arial_Normal" isDefault="true"
fontName="Arial" size="8" pdfFontName="Helvetica" pdfEncoding="Cp1252"
isPdfEmbedded="false"/>
  <reportFont name="Arial_Bold" isDefault="false"
fontName="Arial" size="8" isBold="true" pdfFontName="Helvetica-Bold"
pdfEncoding="Cp1252" isPdfEmbedded="false"/>
  <reportFont name="Arial_Italic" isDefault="false"
fontName="Arial" size="8" isItalic="true" pdfFontName="Helvetica-Oblique"
pdfEncoding="Cp1252" isPdfEmbedded="false"/>
  <reportFont name="Comic_Normal" isDefault="false"
fontName="Comic Sans MS" pdfFontName="COMIC.TTF" pdfEncoding="Identity-H"
isPdfEmbedded="true"/>
  <reportFont name="Comic_Bold" isDefault="false" fontName="Comic
Sans MS" isBold="true" pdfFontName="COMICBD.TTF" pdfEncoding="Identity-H"
isPdfEmbedded="true"/>
  <parameter name="ReportTitle" class="java.lang.String"/>
  <parameter name="SummaryImage" class="java.awt.Image"/>
  <queryString><![CDATA[
    SELECT do.id as object_id, do.name as object_name, dot.name
    as object_type_name
    FROM DCM_Object do, DCM_Object_Type dot
    WHERE do.type_id = dot.type_id
    AND do.type_id in (3, 6, 7, 9, 10)
    ORDER BY do.type_id, do.id
  ]]></queryString>
  <field name="object_id" class="java.lang.Integer"/>
  <field name="object_name" class="java.lang.String"/>
```

```

    <field name="object_type_name" class="java.lang.String"/>
    <variable name="ReportVersion" class="java.lang.String"
    resetType="Report" resetGroup="" calculation="Nothing">
        <variableExpression>
            "0.2.5"
        </variableExpression>
    </variable>
    <variable name="ObjectName" class="java.lang.String"
    resetType="None" calculation="Nothing">
        <variableExpression>
            ${object_type_name}
        </variableExpression>
    </variable>
    <variable name="GroupCount" class="java.lang.Integer"
    resetType="Group" resetGroup="FirstLetterGroup" calculation="System">
        <variableExpression/>
        <initialValueExpression>
            ($V{GroupCount} != null) ? (new
Integer($V{GroupCount}.intValue() + 1)) : (new Integer(1))
        </initialValueExpression>
    </variable>
    <variable name="TitleString" class="java.lang.String"
    resetType="None" calculation="Nothing">
        <variableExpression>
            "Inventory"
        </variableExpression>
    </variable>
    <variable name="TitleDetail" class="java.lang.String"
    resetType="None" calculation="Nothing">
        <variableExpression>
            "There are " + String.valueOf($V{GroupCount}) + "
groups and a total of " + String.valueOf($V{REPORT_COUNT}) + " items on
this report."
        </variableExpression>
    </variable>
    <variable name="PageHeaderString" class="java.lang.String"
    resetType="None" calculation="Nothing">
        <variableExpression>
            "Inventory List"
        </variableExpression>
    </variable>
    <group name="FirstLetterGroup" isStartNewColumn="false"
isResetPageNumber="false" isReprintHeaderOnEachPage="false"
minHeightToStartNewPage="50">
        <groupExpression>
            $V{ObjectName}
        </groupExpression>
        <groupHeader>
            <band height="25">
                <rectangle>
                    <reportElement x="0" y="5" width="270" height="15"
forecolor="#B8B9FB" backcolor="#B8B9FB"/>
                    <graphicElement stretchType="NoStretch"/>
                </rectangle>
                <staticText>
                    <reportElement x="0" y="7" width="30" height="10"
forecolor="#D51B20"/>
                    <textElement textAlignment="Left">
                        <font reportFont="Arial_Normal"
isUnderline="true"/>
                    </textElement>
                    <text>Name : </text>
                </staticText>
                <textField isBlankWhenNull="true">
                    <reportElement x="35" y="7" width="235"
height="10" forecolor="#D51B20"/>
                    <textElement textAlignment="Left">
                        <font reportFont="Arial_Normal"/>
                    </textElement>
                    <textFieldExpression class="java.lang.String">
                        ${object_type_name}
                    </textFieldExpression>
                </textField>
            </band>
        </groupHeader>
    </group>

```

```

        </textFieldExpression>
    </textField>
</band>
</groupHeader>
<groupFooter>
    <band height="15">
        <line>
            <reportElement x="0" y="2" width="270" height="0"
forecolor="#FF0000"/>
            <graphicElement stretchType="NoStretch"/>
        </line>
        <staticText>
            <reportElement x="5" y="3" width="35" height="10"
forecolor="#FF0000"/>
            <textElement textAlignment="Left"
lineSpacing="Single">
                <font reportFont="Arial_Bold"/>
            </textElement>
            <text>Count : </text>
        </staticText>
        <textField>
            <reportElement x="40" y="3" width="25" height="10"
forecolor="#FF0000"/>
            <textElement textAlignment="Left"
lineSpacing="Single">
                <font reportFont="Arial_Bold"/>
            </textElement>
            <textFieldExpression class="java.lang.Integer">
                ${FirstLetterGroup_COUNT}
            </textFieldExpression>
        </textField>
    </band>
</groupFooter>
</group>
<!--
    <group name="BreakGroup">
        <groupExpression>new Boolean(${BreakGroup_COUNT}.intValue() >
5)</groupExpression>
        <groupHeader>
            <band height="5"/>
        </groupHeader>
        <groupFooter>
            <band height="5"/>
        </groupFooter>
    </group>
-->
<title>
    <band height="100">
        <rectangle>
            <reportElement x="0" y="0" width="555" height="100"
mode="Transparent"/>
            <graphicElement stretchType="NoStretch"/>
        </rectangle>
        <rectangle>
            <reportElement x="0" y="0" width="150" height="100"
mode="Transparent"/>
            <graphicElement stretchType="NoStretch"/>
        </rectangle>
        <image scaleImage="RetainShape">
            <reportElement x="5" y="10" width="140" height="80"/>
            <graphicElement/>
            <imageExpression>"tc_logo.jpg"</imageExpression>
        </image>
        <rectangle>
            <reportElement x="150" y="0" width="200" height="40"
mode="Transparent"/>
            <graphicElement stretchType="NoStretch"/>
        </rectangle>
        <textField isBlankWhenNull="true"
isStretchWithOverflow="true">
            <reportElement x="155" y="5" width="190" height="30"
forecolor="#7778DA"/>

```

```

        <textElement textAlignment="Center"
lineSpacing="Single">
        <font reportFont="Arial_Bold" size="15"/>
        </textElement>
        <textFieldExpression
class=" java.lang.String">${TitleString}</textFieldExpression>
        </textField>
        <rectangle>
        <reportElement x="350" y="0" width="205" height="40"
mode="Transparent"/>
        <graphicElement stretchType="NoStretch"/>
        </rectangle>
        <textField>
        <reportElement x="355" y="10" width="195" height="30"
forecolor="#111111"/>
        <textElement textAlignment="Right" lineSpacing="Single">
        <font reportFont="Arial_Normal" size="12"/>
        </textElement>
        <textFieldExpression class=" java.lang.String">(new
Date()).toString(</textFieldExpression>
        </textField>
        <textField evaluationTime="Report">
        <reportElement x="155" y="55" width="390" height="40">
        <printWhenExpression>new
Boolean(${GroupCount}.intValue() == 1)</printWhenExpression>
        </reportElement>
        <textElement textAlignment="Justified" lineSpacing="Single">
        <font reportFont="Arial_Normal" size="11"/>
        </textElement>
        <textFieldExpression
class=" java.lang.String">${TitleDetail}</textFieldExpression>
        </textField>
        <textField evaluationTime="Report">
        <reportElement x="155" y="55" width="390" height="40">
        <printWhenExpression>new Boolean(${GroupCount}.intValue() >
1)</printWhenExpression>
        </reportElement>
        <textElement textAlignment="Justified" lineSpacing="Single">
        <font reportFont="Arial_Normal" size="11"/>
        </textElement>
        <textFieldExpression
class=" java.lang.String">${TitleDetail}</textFieldExpression>
        </textField>
</band>
</title>
<pageHeader>
<band height="32">
        <textField>
        <reportElement x="5" y="5" width="550" height="25"/>
        <textElement textAlignment="Left">
        <font reportFont="Arial_Bold" size="18"/>
        </textElement>
        <textFieldExpression
class=" java.lang.String">${PageHeaderString}</textFieldExpression>
        </textField>
</band>
</pageHeader>
<columnHeader>
<band height="14">
        <rectangle>
        <reportElement x="0" y="0" width="270" height="14"
backcolor="#4C4C4C"/>
        <graphicElement pen="None"/>
        </rectangle>
        <line>
        <reportElement x="50" y="0" width="0" height="14"
forecolor="#FFFFFF"/>
        <graphicElement stretchType="NoStretch"/>
        </line>
        <line>
        <reportElement x="200" y="0" width="0" height="14"
forecolor="#FFFFFF"/>

```

```

        <graphicElement stretchType="NoStretch"/>
    </line>
    <staticText>
        <reportElement x="0" y="2" width="50" height="10"
forecolor="#FFFFFF"/>
        <textElement textAlignment="Center">
            <font reportFont="Arial_Bold"/>
        </textElement>
        <text>ID</text>
    </staticText>
    <staticText>
        <reportElement x="50" y="2" width="150" height="10"
forecolor="#FFFFFF"/>
        <textElement textAlignment="Center">
            <font reportFont="Arial_Bold"/>
        </textElement>
        <text>Name</text>
    </staticText>
</band>
</columnHeader>
<detail>
    <band height="13">
        <textField>
            <reportElement x="0" y="0" width="50" height="10">
                </reportElement>
                <textElement textAlignment="Center">
                    </textElement>
                    <textFieldExpression class="java.lang.Integer">
                        ${object_id}
                    </textFieldExpression>
                </textField>
            <textField isBlankWhenNull="true"
isStretchWithOverflow="true">
                <reportElement x="50" y="0" width="150" height="10"/>
                <textElement textAlignment="Left"/>
                <textFieldExpression class="java.lang.String">
                    ${object_name}
                </textFieldExpression>
            </textField>
        </line>
        <reportElement x="0" y="12" width="270" height="0"
forecolor="#999999" positionType="Float"/>
        <graphicElement stretchType="NoStretch"/>
    </line>
    </band>
</detail>
<columnFooter>
    <band height="12">
        <rectangle>
            <reportElement x="0" y="2" width="270" height="10"
backcolor="#555555"/>
            <graphicElement pen="None"/>
        </rectangle>
        <staticText>
            <reportElement x="5" y="2" width="40" height="10"
forecolor="#FFFFFF"/>
            <textElement textAlignment="Left" lineSpacing="Single">
                <font reportFont="Arial_Bold"/>
            </textElement>
            <text>Count : </text>
        </staticText>
        <textField>
            <reportElement x="45" y="2" width="25" height="10"
forecolor="#FFFFFF"/>
            <textElement textAlignment="Left" lineSpacing="Single">
                <font reportFont="Arial_Bold"/>
            </textElement>
            <textFieldExpression class="java.lang.Integer">
                ${COLUMN_COUNT}
            </textFieldExpression>
        </textField>
    </band>

```

```

        </columnFooter>
<pageFooter>
<band height="40">
  <line>
    <reportElement x="0" y="15" width="555" height="0"
mode="Transparent"/>
    <graphicElement stretchType="NoStretch"/>
  </line>
  <staticText>
    <reportElement x="5" y="20" width="110" height="12"
forecolor="#97D414"/>
    <textElement textAlignment="Left" lineSpacing="Single">
      <font reportFont="Arial_Normal" size="10"/>
      </textElement>
      <text>www.ibm.com</text>
    </staticText>
    <rectangle>
      <reportElement x="120" y="25" width="1" height="1"
forecolor="#97D414"/>
      <graphicElement stretchType="NoStretch"/>
    </rectangle>
    <staticText>
      <reportElement x="125" y="20" width="70" height="12"
forecolor="#97D414"/>
      <textElement textAlignment="Left" lineSpacing="Single">
        <font reportFont="Arial_Normal" size="10"/>
        </textElement>
        <text>(416) 932-9367</text>
      </staticText>
      <rectangle>
        <reportElement x="200" y="25" width="1" height="1"
forecolor="#97D414"/>
        <graphicElement stretchType="NoStretch"/>
      </rectangle>
      <staticText>
        <reportElement x="205" y="20" width="100" height="12"
forecolor="#97D414"/>
        <textElement textAlignment="Left" lineSpacing="Single">
          <font reportFont="Arial_Normal" size="10"/>
          </textElement>
          <text>20 Eglinton Ave. West</text>
        </staticText>
        <rectangle>
          <reportElement x="310" y="25" width="1" height="1"
forecolor="#97D414"/>
          <graphicElement stretchType="NoStretch"/>
        </rectangle>
        <staticText>
          <reportElement x="315" y="20" width="50" height="12"
forecolor="#97D414"/>
          <textElement textAlignment="Left" lineSpacing="Single">
            <font reportFont="Arial_Normal" size="10"/>
            </textElement>
            <text>Suite 600</text>
          </staticText>
          <rectangle>
            <reportElement x="370" y="25" width="1" height="1"
forecolor="#97D414"/>
            <graphicElement stretchType="NoStretch"/>
          </rectangle>
          <staticText>
            <reportElement x="375" y="20" width="35" height="12"
forecolor="#97D414"/>
            <textElement textAlignment="Left" lineSpacing="Single">
              <font reportFont="Arial_Normal" size="10"/>
              </textElement>
              <text>Toronto</text>
            </staticText>
            <rectangle>
              <reportElement x="415" y="25" width="1" height="1"
forecolor="#97D414"/>
              <graphicElement stretchType="NoStretch"/>
            </rectangle>

```

```

        </rectangle>
        <staticText>
            <reportElement x="420" y="20" width="40" height="12"
forecolor="#97D414"/>
            <textElement textAlignment="Left" lineSpacing="Single">
                <font reportFont="Arial_Normal" size="10"/>
            </textElement>
            <text>Ontario</text>
        </staticText>
        <rectangle>
            <reportElement x="460" y="25" width="1" height="1"
forecolor="#97D414"/>
            <graphicElement stretchType="NoStretch"/>
        </rectangle>
        <staticText>
            <reportElement x="465" y="20" width="45" height="12"
forecolor="#97D414"/>
            <textElement textAlignment="Left" lineSpacing="Single">
                <font reportFont="Arial_Normal" size="10"/>
            </textElement>
            <text>M4R 1K8</text>
        </staticText>
        <rectangle>
            <reportElement x="510" y="25" width="1" height="1"
forecolor="#97D414"/>
            <graphicElement stretchType="NoStretch"/>
        </rectangle>
        <staticText>
            <reportElement x="515" y="20" width="45" height="12"
forecolor="#97D414"/>
            <textElement textAlignment="Left" lineSpacing="Single">
                <font reportFont="Arial_Normal" size="10"/>
            </textElement>
            <text>Canada</text>
        </staticText>
    </band>
</pageFooter>
<summary>
    <band height="0"/>
<!--
    <band height="60">
        <rectangle>
            <reportElement x="0" y="10" width="555" height="50"
forecolor="#111111" backcolor="#DDDDDD"/>
            <graphicElement stretchType="NoStretch"/>
        </rectangle>
        <image scaleImage="RetainShape">
            <reportElement x="5" y="20" width="80" height="40"/>
            <graphicElement/>
            <imageExpression class="java.awt.Image">
                ${SummaryImage}
            </imageExpression>
        </image>
        <staticText>
            <reportElement x="200" y="15" width="200" height="40"/>
            <textElement textAlignment="Justified" lineSpacing="Single">
                <font size="12"/>
            </textElement>
            <text>End of Report!</text>
        </staticText>
    </band>
-->
</summary>
</jasperReport>

```

Displaying the generated report

Based on the design defined in the XML file presented above, the generated **Inventory** report has the following format:

The screenshot displays two side-by-side views of an 'Inventory List' report. The left view shows a detailed table with columns for ID, Name, and a third column containing various system details. The right view shows a similar table with columns for ID and Name, listing 'Global Fabric' items. The report includes a header with the IBM logo, the title 'Inventory', and the date 'Thu Feb 06 11:42:11 EST 2008'. A message states 'There are 3 groups and a total of 130 items in this report.'

Figure 20. The Inventory report generated based on the design defined in the XML file

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.

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 might 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.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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
2Z4A/101
11400 Burnet Road
Austin, TX 78758 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

AIX
DB2
DB2 Universal Database
IBM
The IBM logo
Netfinity
RS/6000
Tivoli
Tivoli Enterprise

WebSphere

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks, of Microsoft Corporation in the U.S. and other countries.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S., and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Glossary

Using the Tivoli Intelligent Orchestrator presupposes familiarity with and common understanding of a number of specific terms. Provided below are definitions for some of the most commonly used terms, as well as definitions that are adapted to reflect the Tivoli Intelligent Orchestrator's specific functionality.

Tivoli Intelligent Orchestrator-specific terms

A

application. An application can be defined as a logical assemble of hardware and software that work in conjunction to fulfill a business need. Typically, an application does not include the network infrastructure components, such as switches, load balancers, routers, VLANs, and so on.

application cluster. The Tivoli Intelligent Orchestrator makes use of clustering technology, which allows software applications to run on a collection of computing resources. Clustering is a way of achieving horizontal scalability, which implies that an application scales by adding computers to an existing cluster. The clustering technology is currently available from all major software vendors, for all three standard tiers of a Web-enabled application: Web server, application server, and database server.

R

resource. This is a computing resource, usually consisting of a computing unit with at least a CPU and some RAM. In its current implementation, a unit is a distinct computer/server, and the system currently supports various technologies such as PCs with Windows 2000 or Linux.

resource pool/spare pool. This is a transient collection of computer resources (units) that are currently not dedicated to an execution environment. The decision on how to organize resource pools is an economic and response-time one. You can allocate all the servers in one resource pool, or have two pools, one with servers+Windows 2000, and one with servers+Linux, and so on. The more specialized the pool, the better the response time for allocation (the resource is already in a very advanced state of readiness), however this arrangement prevents timely optimal reallocation "between" pools. With one generic pool only, the allocation time increases, as an operating system has to be installed first and so on. However, the allocation can be optimal among all running environments requiring dynamic allocation of resources.

S

server. This is a computer that can either belong to a particular cluster and run a load balanced instance of a particular application, or reside in a resource pool, waiting to be assigned to run a managed application. The server can also reside in a special state, such as out of service or in transition. Servers that belong to the same cluster are configured to be

in the same virtual network and can be served by dedicated network hardware (associated with a cluster), such as switches and load balancers. A cluster has a set of drivers to control the hardware.

switch. The Tivoli Intelligent Orchestrator is affecting the networking infrastructure by moving servers from one VLAN to another. The system manages network switches (both layer 2 and 3), routers, load balancers, and firewalls. The Tivoli Intelligent Orchestrator's technology relies on the existence of a layer 3 switch in the target environment, which is the only one allowing the system to programmatically create and modify VLANs, therefore moving computers from one environment to another. The Tivoli Intelligent Orchestrator manages the following switch types: Cisco, Extreme, and Foundry.

General terms

A

ASP. Abbreviation for Application Service Provider. ASPs are third-party entities that manage and distribute software-based services and solutions to customers across a wide area network from a central data center. In essence, ASPs are a way for companies to outsource some or almost all aspects of their information technology needs.

C

cluster. A logical group of servers, associated with one of the application tiers (e.g. Web, application, database) of a particular application.

CPU. Abbreviation for Central Processing Unit. The CPU is the brains of the computer. Sometimes referred to simply as the processor or central processor, the CPU is where most calculations take place. In terms of computing power, the CPU is the most important element of a computer system.

D

data center. A centralized storage facility used by an ASP to retain database information related to the decision-making processes of an organization. Data centers can house applications at all levels: Web, application, and database. Many MSPs and ASPs house some of the end customer's equipment, then use a colocation facility to provide the redundant power, cabling, security, and access to the Internet through one or more ISPs.

DHCP. Abbreviation for Dynamic Host Configuration Protocol. DHCP is a protocol for assigning dynamic IP addresses to devices on a network. With dynamic addressing, a device can have a different IP address every time it connects to the network. In some systems, the device's IP address can even change while it is still connected. DHCP also supports a mix of static and dynamic IP addresses. Dynamic addressing simplifies network administration because the software keeps track of IP addresses rather than requiring an administrator to manage the task. This means that a new computer can be added to a network without the hassle of manually assigning it a unique IP address. Many ISPs use dynamic IP addressing for dial-up users.

DMZ. Abbreviation for Demilitarized Zone. A DMZ is used by a company that wants to host its own Internet services without sacrificing unauthorized access to its private network. The DMZ sits between the Internet and an internal network's line of defense,

usually some combination of firewalls and bastion hosts. Typically, the DMZ contains devices accessible to Internet traffic, such as Web (HTTP) servers, FTP servers, SMTP (e-mail) servers and DNS servers.

DNS. Abbreviation for Domain Name System (or Service). A DNS is an Internet service that translates domain names into IP addresses. Because domain names are alphabetic, they are easier to remember. The Internet however, is really based on IP addresses. Therefore, every time you use a domain name, a DNS service must translate the name into the corresponding IP address. For example, the domain name *www.example.com* might translate to 198.105.232.4. The DNS system is, in fact, its own network. If one DNS server does not know how to translate a particular domain name, it asks another one, and so on, until the correct IP address is returned.

F

firewall. A system designed to prevent unauthorized access to or from a private network. Firewalls can be implemented in both hardware and software, or a combination of both. Firewalls are frequently used to prevent unauthorized Internet users from accessing private networks connected to the Internet, especially intranets. All messages entering or leaving the intranet pass through the firewall, which examines each message and blocks those that do not meet the specified security criteria.

FTP. Abbreviation for File Transfer Protocol. The FTP is the protocol used on the Internet for sending files.

H

HTTP. Abbreviation for HyperText Transfer Protocol. The HTTP is the protocol used by the World Wide Web to define how messages are formatted and transmitted, and what actions Web servers and browsers should take in response to various commands. For example, when you enter a URL in your browser, this actually sends an HTTP command to the Web server, directing it to fetch and transmit the requested Web page.

I

ICMP. Abbreviation for Internet Control Message Protocol, an extension to the Internet Protocol (IP) defined by RFC 792. ICMP supports packets containing error, control, and informational messages. The PING command, for example, uses ICMP to test an Internet connection.

IMAP. Abbreviation for Internet Message Access Protocol. IMAP is a protocol for retrieving e-mail messages. The latest version, IMAP4, is similar to POP (POP version 3) but supports some additional features. For example, with IMAP4, you can search through your e-mail messages for keywords while the messages are still on mail server. You can then choose which messages to download to your machine.

intranet. A network based on TCP/IP protocols belonging to an organization, usually a corporation. The intranet is accessible only by the organization's members, employees, or others with authorization, and used to share information. An intranet's Web sites look and act just like any other Web sites, but the firewall surrounding an intranet fends off unauthorized access.

ISP. Abbreviation for Internet Service Provider. An ISP is a company that provides access to the Internet. For a monthly fee, the service provider gives you a software package or instructions for configuring standard tools, username, password, and access phone number. Equipped with a modem, you can then log on to the Internet and browse the World Wide Web, and send and receive e-mail.

J

JMS. Abbreviation for Java™ Message System. By combining Java technology with enterprise messaging, the JMS API provides a new, powerful tool for solving enterprise computing problems. Enterprise messaging provides a reliable, flexible service for the asynchronous exchange of critical business data and events throughout an enterprise. The JMS API adds to this a common API and provider framework that enables the development of portable, message based applications in the Java programming language. The JMS API improves programmer productivity by defining a common set of messaging concepts and programming strategies that will be supported by all JMS technology-compliant messaging systems.

JMX. Abbreviation for Java™ Management Extensions. JMX represents a universal, open technology for management, ready to be deployed across all industries, wherever management is or will be needed. By design, this new standard is suitable for adapting legacy systems, implementing new management solutions and plugging into those of the future.

L

LAN. Abbreviation for Local-Area Network. A LAN is a computer network that spans a relatively small area. Most LANs are confined to a single building or group of buildings. However, one LAN can be connected to other LANs over any distance through communications lines and radio waves. Most LANs connect workstations and personal computers. Each node (individual computer) in a LAN has its own CPU with which it executes programs, but it also is able to access data and devices anywhere on the LAN. This means that many users can share devices, such as laser printers, as well as data. Users can also use the LAN to communicate with each other, by sending e-mail or engaging in chat sessions.

load balancing. Load balancing is the process of distributing processing and communications activity evenly across a computer network so that no single device is overwhelmed. Load balancing is especially important for networks where it's difficult to predict the number of requests that will be issued to a server. Busy Web sites typically employ two or more Web servers in a load balancing scheme. If one server starts to get swamped, requests are forwarded to another server with more capacity.

M

MSP. Abbreviation for Managed Service Provider, also called a Management Service Provider. An MSP is a company that manages information technology services for other companies through the Web. An MSP client may use internal operations or an ASP to run its business functions.

N

NIC. Abbreviation for Network Interface Card. An NIC is an expansion board you insert into a computer so the computer can be connected to a network. Most NICs are designed for a particular type of network, protocol, and media, although some can serve multiple networks.

NSH. Abbreviation for Network Shell. The NSH is a command line interface that offers distributed, cross-platform access to crucial systems across the enterprise. From NSH, system administrators have seamless and secure access to all remote servers from one central management workstation. NSH gives the system administrators the power to transfer and manipulate files on multiple servers, stop and start processes remotely, and execute commands or entire shell scripts on multiple systems in parallel.

P

POP. Abbreviation for Post Office Protocol. A POP is a protocol used to retrieve e-mail from a mail server. Most e-mail applications (sometimes called an e-mail client) use the POP protocol, although some can use the newer IMAP (Internet Message Access Protocol).

PXE. Abbreviation for Pre-Boot Execution Environment. PXE is one of the components of Intel's WfM specification, allowing a workstation to boot from a server on a network prior to booting the operating system on the local hard drive. A PXE-enabled workstation connects its NIC to the LAN through a jumper, which keeps the workstation connected to the network even when the power is off. Because a network administrator doesn't have to physically visit the specific workstation and manually boot it, operating systems and other software, such as diagnostic programs, can be loaded onto the device from a server over the network. PXE is a mandatory element of the WfM specification. To be considered compliant, the PXE must be supported by the computer's BIOS and its NIC.

R

router. A router is a device that connects any number of LANs. Routers use headers and a forwarding table to determine where packets go. Very little filtering of data is done through routers, as they do not care about the type of data they handle.

RSH. Abbreviation for Remote Shell. The RSH is a protocol that performs a back connect from a remote host which requires a direct connection to the end user's desktop. The RSH protocol lets you log in to your host without a password. Generally, RSH is used to connect to hosts where system security is not a concern.

S

server cluster. A server cluster is a group of networked servers that are housed in one location. A server cluster streamlines internal processes by distributing the workload between the individual components of the cluster and expedites computing processes by harnessing the power of multiple servers. The clusters rely on load-balancing software that accomplishes such tasks as tracking demand for processing power from different machines, prioritizing the tasks, and scheduling and rescheduling them depending on priority and demand that users put on the network. When one server in the cluster fails, another can step in as a backup.

Service Level Agreement (SLA). An SLA is a contract between an Application Service Provider (ASP) and the end user which stipulates and commits the ASP to a required level of service. An SLA should contain a specified level of service, support options, enforcement or penalty provisions for services not provided, a guaranteed level of system performance as relates to downtime or uptime, a specified level of customer support and what software or hardware will be provided and for what fee.

Service Level Management. The Service Level Management is a way of managing system performance (i.e. response time, compute power, storage space) by a contract in which a Service Provider agrees to deliver an agreed-upon minimum level of service.

SMTP. Abbreviation for Simple Mail Transfer Protocol. SMTP is a protocol for sending e-mail messages between servers. Most e-mail systems sending mail over the Internet use SMTP to send messages between servers. The messages can then be retrieved with an e-mail client using either POP or IMAP. In addition, SMTP is generally used to send messages from a mail client to a mail server, hence you need to specify both the POP/IMAP server and the SMTP server when you configure your e-mail application.

SNMP. Abbreviation for Simple Network Management Protocol. SNMP is a set of protocols for managing complex networks. It works by sending messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.

SSH. Abbreviation for Secure Shell. SSH is a program to log into another computer over a network, to execute commands in a remote machine, and to move files from one machine to another. It provides strong authentication and secure communications over insecure channels. SSH is available for Windows, Unix, Macintosh, and OS/2, and it also works with RSA authentication.

switch. A switch is a device that filters and forwards packets between LAN segments. Switches operate at the data link layer (layer 2) and sometimes the network layer (layer 3) of the OSI Reference Model and therefore support any packet protocol. LANs using switches to join segments are called switched LANs or, for Ethernet networks, switched Ethernet LANs.

T

TCP/IP. Abbreviation for Transmission Control Protocol/Internet Protocol, the suite of communications protocols used to connect hosts on the Internet. TCP/IP uses several protocols, the two main ones being TCP and IP. TCP/IP is built into the UNIX operating system and is used by the Internet, making it the de facto standard for transmitting data over networks. Even network operating systems that have their own protocols, such as Netware, also support TCP/IP.

TELNET. This is a protocol designed for terminal-oriented remote login, that operates using the TCP Protocol, and depends heavily on option negotiation. It provides a general, bi-directional, eight-bit byte oriented communications facility. Its primary goal is to allow a standard method of interfacing terminal devices and terminal-oriented processes to each other.

U

UDP/IP. Abbreviation for User Datagram Protocol. UDP/IP is a connection less protocol that, like TCP, runs on top of IP networks. Unlike TCP/IP, it provides very few error recovery services, offering instead a direct way to send and receive datagrams over an IP network. It is used primarily for broadcasting messages over a network.

V

VLAN. Abbreviation for virtual LAN. A VLAN is a network of computers that behave as if they are connected to the same wire even though they may actually be physically located on different segments of a LAN. VLANs are configured through software rather than hardware, which makes them highly flexible. One of the biggest advantages of VLANs is that when a computer is physically moved to another location, it can stay on the same VLAN without any hardware reconfiguration.

W

Web farm. A Web farm, or Web server farm, refers to either a Web site that runs off of more than one server or an ISP that provides Web hosting services using multiple servers.

Index

A

ACL

- adding a new 77
- adding a new access rule 78
- assigning workflows 79
- configuring 77
- editing an access rule 78
- editing properties 78
- editing user-defined variables 79
- removing the access rule 79

application

- adding a new cluster 117
 - adding a new dedicated server 120
 - bringing under the Intelligent Orchestrator's management 120
 - displaying cluster properties 139
 - editing properties 116
 - managing 115
 - manually moving to maintenance 121
 - removing a cluster 122
 - removing from a customer account 122
 - setting up 116
- application operating mode 109
- changing 109
 - changing using the Command box 110
 - changing using the command-line interface 111
 - changing using the Customer applications menu 110
 - changing using the Web-based interface 110

application topologies

- device driver 11, 12
- interactive user interface 14
- SOAP initiation of device functions 14
- software stack 12, 14
- switch fabric 11

B

blade server

- configuring 68
- removing 73

blade server chassis

- adding a new 69
- adding a new blade server 70
- assigning workflows to 70
- configuring credentials 70
- editing properties 69
- editing properties for a blade server 71
- editing user-defined properties 70
- removing a blade server 73

- setting up a network interface 72
- setting up a NIC 71
- setting up the routing configuration for the network interface 72

boot server

- adding a new 65
- adding an interface card 67
- assigning workflows 66
- configuring 65
- configuring credentials 68
- edit user-defined variables 66
- editing properties 65
- removing 68
- setting up a network interface 67
- setting up NICs 66
- setting up the routing configuration for the network interface 67

C

case study

- defining the report design in an XML file 247
- displaying the generated report 254

cluster

- adding a new dedicated server 120
- adding to an application 117
- assigning a software stack 119
- assigning workflows 118, 174
- assigning workflows to 171
- inheriting device driver workflows 171, 174, 176
- managing 115
- removing from an application 122
- setting up routing properties 118

cluster operating mode 111

- changing 111
- changing using the Command box 112
- changing using the command-line interface 113
- changing using the Customer applications menu 112
- changing using the Web-based interface 112

command variable

- adding new 186
- editing properties 187
- mapping 152, 187
- removing 187

command-line interface 19, 111, 113

- SOAP scripts 221

configuring

- ACLs 77
- blade servers 68
- boot servers 65
- firewalls 45

- global operating mode 18
- inventory resources 23
- Java plug-ins 181
- license pools 85
- load balancers 38
- logical device operations 191
- power units 80
- resource pool 104
- routers 44
- servers 48
- simple commands 183
- software patches 91
- software products 88
- software stacks 97
- subnets 24
- switch fabrics 35
- switches 28
- terminal servers 73
- users 19
- virtual IPs 46
- VLANs 26
- credentials
 - adding a new service access point 60
 - configuring 59
 - editing the properties of the service access point 60
 - removing the service access point 63
 - setting up RSA credentials for the SAP 63
 - setting up password credentials for the SAP 62
 - setting up SNMP credentials for the SAP 63
- customer
 - editing properties 115
 - managing 114
 - removing 115
 - setting up a new account 114

D

- data center device
 - assigning workflows to 176, 177
 - inheriting device driver workflows 176, 178
- Deployment Engine
 - managing through SOAP commands 223
 - managing through the Web-based interface 180
- device driver
 - adding a new 170
 - assigning workflows to 61, 171, 178, 179
 - configuring 169
 - editing properties 170
 - removing 171

F

- firewall
 - accessing other views 45
 - accessing view 45
 - configuring 45

G

- global operating mode 10, 18
 - changing 18
 - changing using the Command box 19
 - changing using the command-line interface 19
 - changing using the System configuration and workflow management menu 18
 - changing using the Web-based interface 18
- glossary
 - application 255
 - application cluster 255
 - ASP 256
 - cluster 256
 - CPU 256
 - data center 256
 - DMZ 256
 - DNS 257
 - firewall 257
 - FTP 257
 - general terms 256
 - HTTP 257
 - ICMP 257
 - IMAP 257
 - Intelligent Orchestrator-specific terms 255
 - intranet 257
 - ISP 258
 - JMS 258
 - JMX 258
 - LAN 258
 - load balancing 258
 - MSP 258
 - NIC 259
 - NSH 259
 - POP 259
 - PXE 259
 - resource 255
 - resource pool 255
 - router 259
 - RSH 259
 - server 255
 - server cluster 259
 - service level agreement (SLA) 260
 - service level management 260
 - SMTP 260
 - SNMP 260
 - SSH 260
 - switch 256, 260
 - TCP/IP 260
 - Telnet 260
 - UDP/IP 261
 - VLAN 261
 - Web farm 261

I

- Intelligent Orchestrator

- application topologies 11
- available user interfaces 6
- case study for the Inventory report 247
- command-line interface 221
- configuring 17
- configuring Java plug-ins 181
- configuring logical device operations 181
- configuring simple commands 181
- configuring workflows 143
- generating reports 199
- glossary 255
- introducing 1
- logging off from 4, 5
- logging on to 4
- main components 2
- managing 109
- operating modes 10
- SOAP commands 239
- user credentials 21
- what is 1
- interactive view
 - accessing other views from the firewall view 45
 - accessing other views from the NIC view 53
 - accessing other views from the router view 44
 - accessing other views from the switch fabric view 36
 - accessing other views from the switch view 30
 - accessing other views from the VLAN view 28
 - adding a server 136
 - creating a router connection 134
 - creating a VIP 135
 - creating a VLAN 132
 - firewall 45
 - hardware rebooting a device 135
 - initializing a device 135
 - moving port to VLAN 134
 - NIC for a server 52
 - performing deployment tasks 131
 - removing a router connection 134
 - removing a server 136
 - removing a VIP 136
 - removing a VLAN 133
 - router 44
 - show/hide capabilities on the switch fabric view 37
 - show/hide capabilities on the switch view 31
 - switch 30, 131
 - switch fabric 35
 - turning a port on 133
 - turning port off 133
 - VLAN 27, 131, 132
- interfaces
 - command-line 10
 - Web-based 6
- inventory resource
 - configuring 23

J

- Java plug-in
 - configuring 181
 - detail view mode 181
 - importing 183
 - removing 182
 - simple view mode 181

L

- license pool
 - adding a new 85
 - adding a new license key 86
 - configuring 85
 - editing user-defined variables 88
 - modifying the license key's parameters 87
 - removing 88
 - removing license keys 88
 - renaming license keys 87
 - setting up configuration 86
- load balancer
 - adding a new load balancer 39
 - adding a VIP 41
 - assigning workflows to load balancer 41
 - configuring 38
 - configuring the switch functionality 39
 - edit user-defined variables 40
 - editing properties 40
 - removing 43
 - setting up a network interface 42
 - setting up routing configuration 42
- logical device operation
 - adding a new 173
 - adding new 192
 - assigning to a workflow 61, 102, 175, 177, 179
 - configuring 191
 - detail view mode 191
 - editing properties 193
 - executing 195
 - exporting 196
 - exporting to screen 196
 - exporting to XML file 197
 - importing 197
 - removing 195
 - setting up variables 194
 - simple view mode 191
 - tagging a workflow 172
- logical device operation variable
 - adding a new variable 194
 - editing properties 194
 - removing 194

M

managing

- applications 115
- clusters 115
- customer accounts 114
- Deployment Engine through SOAP commands 223
- Deployment Engine through the Web-based interface 180
- servers 123
- switching operating modes at the application level 109
- switching operating modes at the cluster level 111

monitoring

- application servers 138
- customers 137
- displaying cluster properties for an application 139
- displaying deployment requests 140
- resource pool servers 138
- resource pools 137
- server performance 140

N

navigation tree 7

- Customer applications 9
- Data center assets and resources 8
- Realtime performance monitoring 9
- Reports 9
- System configuration and workflow management 8

NIC

- accessing other views 53
- accessing view 52

O

operating mode

- automatic 10
- manual 11
- semi-automatic 10

OperationsModeService

- getApplicationMode 240
- getClusterMode 240
- getEffectiveApplicationMode 241
- getEffectiveClusterMode 242
- getGlobalMode 242
- setApplicationMaintenanceMode 243
- setApplicationMode 243
- setClusterMode 244
- setGlobalMode 244
- setMaintenanceMode 245

P

power unit

- adding a new 80

- adding a new outlet 83
- adding an interface card 82
- assigning workflows to 81
- configuring 80
- configuring credentials 83
- edit user-defined variables 81
- editing properties 80
- editing the outlet properties 84
- removing 84
- removing outlet 84
- setting up a network interface 82
- setting up a NIC 81
- setting up configuration 83
- setting up the routing configuration for the network interface 83

R

report list

- All Commands 202
- All Servers 208
- Applications 205
- Clusters 206
- Complete Inventory 210
- Customer Applications reports 205
- Deployment Drivers 203
- Deployment Variables 205
- Detailed Server Usage by Application 215
- Detailed Server Usage Summary by Pool 219
- Inventory 209
- Load Balancers 211
- Maintenance Pool 210
- Network Interfaces 211
- Recommendations 214
- Routing Information 212
- Server Usage Summary by Application 216
- Server Usage Summary by Cluster 217
- Server Usage Summary by Pool 218
- Servers by Cluster 207
- Servers by Customer 207
- Servers by Pool 209
- Simple Commands 203
- Subnetworks 213
- Usage reports 214
- VLANs and Subnetworks 213
- Workflows 204
- Workflows reports 202

reports

- displaying 199
- layout 200
- list 201

resource pool

- adding a new pool 105
- assigning a software stack 106
- assigning workflows to 175
- assigning workflows to resource pool 106
- configuring 104

- edit user-defined variables 106
- editing properties 105
- inheriting device driver workflows 175
- removing 107

router

- accessing other views 44
- accessing view 44
- configuring 44
- creating a connection 134
- removing a connection 134

S

server

- accessing the NIC view 52
- accessing the software stack 130
- adding a new entry to the server stack 56
- adding a new server 49
- adding from the interactive view 136
- assigning a software stack 55
- assigning workflows 51
- asynchronous software reboot 128
- cleaning up within a resource pool 126
- configuring 48
- configuring credentials 59
- configuring software stack 55
- edit user-defined variables 51
- editing configuration for the stack entries 56
- editing properties 50
- generic reboot 127, 129
- hardware reboot 127, 129
- initializing 130
- initializing within a resource pool 125
- installing a patch 94
- installing software products 58, 90
- managing 123
- manually adding overflow servers to a cluster 123
- manually freeing overflow servers 124
- moving between resource pools 125
- performing management tasks 126
- powering off 127
- powering on 127
- querying for servers that require patching 93
- rebooting 127
- refreshing software status 59
- removing 64, 136
- removing a dedicated server from a cluster 124
- removing stack entry 57
- setting up network interfaces 54
- setting up NICs 52
- setting up the routing configuration for the network interface 54
- software reboot 127
- synchronous software reboot 128
- uninstalling software products 58, 91

service access point (SAP)

- adding a new 60

- assigning workflows to 61
- editing properties 60
- removing 63
- setting up RSA credentials 63
- setting up password credentials 62
- setting up SNMP credentials 63

simple command

- adding new 185
- cloning 188
- configuring 183
- detail view mode 184
- editing properties 185
- executing 188
- exporting 189
- exporting to screen 190
- exporting to XML file 190
- importing 191
- removing 187
- setting up variables 186
- simple view mode 184

SOAP

- command syntax 239
- commands available through command line 221
- commands available through the Web-based interface 222
- commands for message translator 237
- commands initiating device functions 228
- modules 222
- OperationsModeService 239
- scripts 221

SOAP module

- app-controller 223
- boot-server-component 229
- cluster-component 230
- deployment-engine 223
- device-component 230
- faultmanagement 224
- file-repository-component 231
- firewall-component 231
- ip-system-component 232
- load-balancer-component 232
- model 224
- OperationsModeService 224
- packager 225
- power-unit-component 233
- recommendations 226
- router-component 234
- simulator 227
- software-component 234
- software-stack-component 235
- spare-pool-component 235
- switch-component 236
- switch-fabric-component 236
- tc-driver-manager 227
- uiuc 228

software patch

- adding a new 92
- applying on a server from the Info tab of the patch 94
- applying to a server 94
- assigning workflows to 97

- configuring 91
- editing properties 96
- editing the user-defined variables 97
- querying for servers that require patching 93
- removing 97
- software product
 - assigning workflows to 90
 - configuring 88
 - edit user-defined variables 90
 - editing configuration 56
- software stack
 - adding a new 98
 - adding a new entry 56, 99
 - adding a new software product 89
 - assigning to a server 55
 - assigning workflows to 102
 - cloning 103
 - configuring 97
 - configuring stack entries 99
 - editing configuration of a nested stack 56
 - editing configuration of stack entry 100
 - editing properties 98
 - editing properties for the software product 89
 - editing user-defined variables 99
 - removing 104
 - removing software product 91
 - removing stack entry 57, 101
- subnet
 - adding a new subnet 24
 - assigning workflows to subnet 25
 - configuring 24
 - editing properties 24
 - removing 25
 - setting up blocking IPs 25
- switch
 - accessing other views 30
 - accessing the switch view 30
 - adding a new switch 29
 - adding a new switch port 33
 - assigning workflows to switch 32
 - configuring 28
 - edit properties 29
 - editing the use-defined variables 32
 - performing deployment tasks from the switch view 131
 - removing 35
 - setting up a network interface 33
 - setting up the routing configuration 34
 - showing/hiding devices, links, and linked objects 31
- switch fabric
 - accessing other views 36
 - accessing the switch fabric view 35
 - adding new switch to default fabric 37
 - configuring 35
 - showing/hiding devices, links, and linked objects 37
- system main components
 - Application Controller 2
 - Data Acquisition Engine 2
 - Data Center Model 2
 - Deployment Engine 2

- Global Resource Manager 3
- Management Interface 3

T

- TCP port
 - moving to VLAN 134
 - turning off 133
 - turning on 133
- terminal server
 - adding a new 74
 - adding an interface card 75
 - assigning workflows 75
 - configuring 73
 - configuring credentials 76
 - editing properties 74
 - editing user-defined variables 74
 - removing 77
 - setting up a network interface 75
 - setting up a NIC 75
 - setting up the routing configuration for the network interface 76
- transition
 - adding a new transition 151
 - editing properties 151
 - mapping variables 152
 - removing 153

U

- user
 - adding a new user 19
 - configuring 19
 - editing the profile 22
 - removing 23

V

- VIP
 - adding a new VIP 46
 - adding a new VIP to load balancer 41
 - assigning workflows to VIP 47
 - configuring 46
 - creating from the interactive view 135
 - editing properties 47
 - removing 47, 136

VLAN

- accessing other interactive views 28
- accessing the VLAN view 27
- adding a new VLAN 26
- assigning workflows to VLAN 27
- configuring 26
- creating a new VLAN 132
- editing properties 27

- moving port to 134
- performing deployment tasks from the VLAN view 131, 132
- removing 28, 133

W

- Web-based interface 18, 110, 112

workflow

- adding a new workflow 148
- assigning 171
- assigning to cluster 171, 174
- assigning to data center device 177
- assigning to device driver 61, 178, 179
- assigning to inventory resource 176
- assigning to resource pool 175
- assigning to software stack 102
- cloning 160
- cluster inheriting device driver workflows 174, 176
- configuring 143
- configuring and using 144
- data center device inheriting device driver workflows 178
- defining a new logical device operation 173
- deleting entire execution history at once 168
- deleting execution history 167
- deleting one execution record at a time 167
- displaying execution details 164
- displaying execution details for embedded workflow 166
- displaying execution history for all workflows 163
- displaying execution history for individual workflow 163
- displaying references 147
- displaying transitions 145
- editing properties 150
- executing 159
- exporting 160
- exporting to screen 161
- exporting to XML file 161
- importing 162
- inheriting device driver workflows 171
- introduction to 143
- removing 160
- setting up recovery path 153
- setting up transitions 150
- setting up variables 153
- tagging the workflow with a logical device operation 61, 102, 172, 175, 177, 179
- using 143

workflow variable

- adding a new variable 154
- assigning default value 156
- editing properties 156
- removing 158
- removing the variable mapping 158



Program Number: 5724-F75