Progress®



OpenEdge® Management:
Resource Monitoring

OpenEdge®

Table of Contents

Copyright....9

Preface....11

Purpose....11
Audience....11
Organization....12
Using ABL documentation....12
References to ABL compiler and run-time features....13
References to ABL data types....13
Typographical conventions....13
Examples of syntax descriptions....14
OpenEdge messages....16
Obtaining more information about OpenEdge messages....17

Chapter 1: Resource Monitoring with OpenEdge Management....19

Setting up OpenEdge Management....20
OpenEdge Management and OpenEdge Explorer....20
Understanding OpenEdge Management terminology....21
Commonly used terms....21
Using the OpenEdge Management or OpenEdge Explorer console....24
Identifying OpenEdge Management resources you can monitor....25
Monitoring a resource....25
Viewing resource monitor performance details....26
Creating and scheduling OpenEdge Management jobs....27
Exporting and importing OpenEdge Management monitoring components....27

Chapter 2: Creating Collections and Collection Views in OpenEdge Management....29

The Dashboard page....30
How collections and collection views differ....30
Collections....30
What is a collection....31
Collection terminology and hierarchy....31
Sample collection scenario....32
Working with collections....33
Creating a collection....33
Editing a collection....34
Deleting a collection....34

Working with collections as labels....35
Accessing and viewing collections as labels....35
Creating a collection from a resource's details page....35
Deleting a member of a collection from a resource's details page....36
Collection views....36
Working with collection views and viewlets....37
Viewlet characteristics....37
Creating a collection view....38
Customizing the view layout....38
Editing a collection view content....39
Deleting a collection view....39
Customizing a viewlet....39

Chapter 3: Setting up Resource Monitoring....41

Before you begin....42
Monitoring remote resources....42
Naming conventions....42
Monitoring a system, network, or file resource: an overview....43
Monitoring a database or other OpenEdge resource....43
Working with CPU and memory resources....43
Selecting a system, network, or file resource to monitor....44
Creating a resource monitoring plan....44
Trending considerations....45
Choosing a schedule for the monitoring plan....45
Adding rules to the monitoring plan....47
Sharing rules and rule sets....48
Calculating meaningful thresholds with the Configuration Advisor....50
Configuration Advisor baselining guidelines....50
Using the Configuration Advisor with a disk, CPU, or file system resource....51
Reviewing the recommended threshold settings....52
Comparing and selecting threshold settings....52
Submitting your threshold setting selections....53
Determining the effectiveness of your selections....53
Understanding alerts....53
Determining OpenEdge Management's action in response to an alert....54
Using default actions....54
Creating an e-mail action....54
Creating a log action....56
Creating a compound action....57
Creating a Simple Network Management Protocol (SNMP) trap action....57
Updating and deleting actions....59
Editing or copying a resource monitor....59
Deleting a resource monitor....60
Working with default values....60

Reviewing and changing global settings 60
Reviewing and changing resource type defaults 61
Using job and report templates 62
Using the OpenEdge Management Library....63

Chapter 4: Monitoring System Resources....65

System resource monitoring overview....65
Reviewing system resource default values....66
Reviewing system resource monitoring plans....66
Reviewing the CPU resource monitoring plan....67
Reviewing the memory resource monitoring plan....68
Reviewing a disk resource monitoring plan....69
Reviewing the file system resource monitoring plan....70

Chapter 5: Monitoring Network Resources....73

Network resource monitoring overview....73
Trending considerations for network resources 74
Reviewing network resource default values 74
Discovering and monitoring TCP and UDP ports 75
Reviewing the list of well-known ports 75
Discovering network resources....76
Adding resource monitors for discovered machines and ports 77
Creating a TCP or UDP resource monitor 78
Creating a PING (ICMP) resource monitor 79
Using the HTTP resource monitor to determine Web site status....79
General considerations 80
Creating an HTTP resource monitor 80

Chapter 6: Monitoring File Resources 85

File resource monitoring overview....85
Log file monitor features....86
File monitor features....86
Verifying mapped network drives 86
Reviewing log file and file resource default values....87
Creating a log file monitor 87
Defining properties 88
Using a bookmark....89
Creating a monitoring plan 90
Adding log file rules 91
Adding log file rule sets....92
Associating rule sets with log file monitors 94
Specifying search criteria....95
Creating search criteria 96

Enabling the log file monitor 100
Editing, copying, or deleting a log file monitor....101
Creating a file resource monitor....101
Trending performance data for file monitor resources....102
Creating a file monitor....102

Chapter 7: Monitoring Batch Program resources....105
Viewing and modifying the default monitoring plan....105
Viewing and modifying the default schedule....106

Chapter 8: Working with Jobs....107
Creating a custom job 108
Creating a job template 109
Using a predefined database maintenance job template 109
Creating a job instance....110
Setting up job completion actions and alerts 112
Updating job completion actions and alerts....114
Deleting job completion actions and alerts....114
Job chaining 114
Using environment variables 117
Editing a job....120
Copying a job....120
Killing a job....121
Deleting a job....121
Running a job....121
Scheduling a job 122
Using cron-based scheduling....122
Scheduling a job instance or a job template....123
Duration Execution Policy....124
Single Execution Policy....124
Job Execution Policy over time....125
Monthly By Date Execution Policy....125
Viewing scheduled jobs 125
Viewing running jobs 126
Viewing job history....126
Viewing database maintenance job instances for an individual database 127
Working with job templates 128
Creating a job template 129
Accessing existing job templates 131
Viewing debug details about jobs....131

Chapter 9: Exporting and Importing in OpenEdge Management.....133
Exporting and importing overview 133

General export and import considerations....134
Exporting OpenEdge Management Library components....134
Exporting referenced items 135
Importing OpenEdge Management Library components....135

Index....137

Copyright

© **2019 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved.**

These materials and all Progress® software products are copyrighted and all rights are reserved by Progress Software Corporation. The information in these materials is subject to change without notice, and Progress Software Corporation assumes no responsibility for any errors that may appear therein. The references in these materials to specific platforms supported are subject to change.

Corticon, DataDirect (and design), DataDirect Cloud, DataDirect Connect, DataDirect Connect64, DataDirect XML Converters, DataDirect XQuery, DataRPM, Deliver More Than Expected, Icenium, Kendo UI, Kinvey, NativeScript, OpenEdge, Powered by Progress, Progress, Progress Software Developers Network, Rollbase, SequeLink, Sitefinity (and Design), Sitefinity, SpeedScript, Stylus Studio, TeamPulse, Telerik, Telerik (and Design), Test Studio, and WebSpeed are registered trademarks of Progress Software Corporation or one of its affiliates or subsidiaries in the U.S. and/or other countries. Analytics360, AppServer, BusinessEdge, DataDirect Autonomous REST Connector, DataDirect Spy, SupportLink, DevCraft, Fiddler, JustAssembly, JustDecompile, JustMock, Kinvey Chat, NativeChat, NativeScript Sidekick, OpenAccess, ProDataSet, Progress Results, Progress Software, ProVision, PSE Pro, SmartBrowser, SmartComponent, SmartDataBrowser, SmartDataObjects, SmartDataView, SmartDialog, SmartFolder, SmartFrame, SmartObjects, SmartPanel, SmartQuery, SmartViewer, SmartWindow, and WebClient are trademarks or service marks of Progress Software Corporation and/or its subsidiaries or affiliates in the U.S. and other countries. Java is a registered trademark of Oracle and/or its affiliates. Any other marks contained herein may be trademarks of their respective owners.

February 2019

Last updated with new content: Release 12.0

Preface

For details, see the following topics:

- Purpose
- Audience
- Organization
- Using ABL documentation
- Typographical conventions
- Examples of syntax descriptions
- OpenEdge messages

Purpose

This guide introduces resource monitoring with OpenEdge® Management. The guide provides an introduction to monitoring resources and describes how to use the OpenEdge Management console; create collections and views; monitor system, network and file resources; use jobs and job templates; and import and export OpenEdge Management monitoring components.

Audience

This guide is designed for database administrators and end users of the OpenEdge Management product. This includes IT managers, ASP hosting companies, and others who are responsible for the day-to-day management and monitoring of resources and OpenEdge databases.

Organization

Resource Monitoring with OpenEdge Management on page 19

Provides an overview of resource monitoring and OpenEdge Management terminology.

Creating Collections and Collection Views in OpenEdge Management on page 29

Describes the **My Dashboard** page and provides information about creating collections of resources and collection views that focus specifically on OpenEdge Management monitoring data you find most valuable.

Setting up Resource Monitoring on page 41

Describes how to set up a resource monitor, including its properties, monitoring plan, schedule, and rules. The chapter also touches briefly on OpenEdge Management alerts and then describes OpenEdge Management actions that occur in response to alerts. Also included is information about default resource monitor values and using the Configuration Advisor with a disk, CPU, memory, or file system resource.

Monitoring System Resources on page 65

Describes the default values OpenEdge Management provides for system resources and the resource rules.

Monitoring Network Resources on page 73

Describes the default values OpenEdge Management provides for network resources and how to create network resource monitors.

Monitoring File Resources on page 85

Describes how to create a log file resource monitor (to monitor file contents) and a file resource monitor (to monitor file characteristics).

Creating Jobs and Job Templates

Describes how to create, edit, copy, schedule, run, or delete a OpenEdge Management job. The chapter provides information about creating job templates, reviewing scheduled jobs, reviewing running jobs, and reviewing job history.

Exporting and Importing in OpenEdge Management on page 133

Describes how to export and import sharable OpenEdge Management Library components, which can help reduce duplication of effort and improve consistency among machines.

Using ABL documentation

OpenEdge provides a special purpose programming language for building business applications. In the documentation, the formal name for this language is *ABL (Advanced Business Language)*. With few exceptions, all keywords of the language appear in all `UPPERCASE`, using a font that is appropriate to the context. All other alphabetic language content appears in mixed case.

For the latest documentation updates go to https://docs.progress.com.

References to ABL compiler and run-time features

ABL is both a compiled and an interpreted language that executes in a run-time engine. The documentation refers to this run-time engine as the *ABL Virtual Machine (AVM)*. When the documentation refers to ABL source code compilation, it specifies *ABL* or *the compiler* as the actor that manages compile-time features of the language. When the documentation refers to run-time behavior in an executing ABL program, it specifies *the AVM* as the actor that manages the specified run-time behavior in the program.

For example, these sentences refer to the ABL compiler's allowance for parameter passing and the AVM's possible response to that parameter passing at run time: "ABL allows you to pass a dynamic temp-table handle as a static temp-table parameter of a method. However, if at run time the passed dynamic temp-table schema does not match the schema of the static temp-table parameter, the AVM raises an error." The following sentence refers to run-time actions that the AVM can perform using a particular ABL feature: "The ABL socket object handle allows the AVM to connect with other ABL and non-ABL sessions using TCP/IP sockets."

References to ABL data types

ABL provides built-in data types, built-in class data types, and user-defined class data types. References to built-in data types follow these rules:

- Like most other keywords, references to specific built-in data types appear in all `UPPERCASE`, using a font that is appropriate to the context. No uppercase reference ever includes or implies any data type other than itself.
- Wherever *integer* appears, this is a reference to the `INTEGER` or `INT64` data type.
- Wherever *character* appears, this is a reference to the `CHARACTER`, `LONGCHAR`, or `CLOB` data type.
- Wherever *decimal* appears, this is a reference to the `DECIMAL` data type.
- Wherever *numeric* appears, this is a reference to the `INTEGER`, `INT64`, or `DECIMAL` data type.

References to built-in class data types appear in mixed case with initial caps, for example, `Progress.Lang.Object`. References to user-defined class data types appear in mixed case, as specified for a given application example.

Typographical conventions

This documentation uses the following typographical and syntax conventions:

Convention	Description
Bold	Bold typeface indicates commands or characters the user types, provides emphasis, or the names of user interface elements.
Italic	Italic typeface indicates the title of a document, or signifies new terms.
SMALL, BOLD CAPITAL LETTERS	Small, bold capital letters indicate OpenEdge key functions and generic keyboard keys; for example, **GET** and **CTRL**.

Convention	Description
KEY1+KEY2	A plus sign between key names indicates a **simultaneous** key sequence: you press and hold down the first key while pressing the second key. For example, **CTRL+X**.
KEY1 KEY2	A space between key names indicates a **sequential** key sequence: you press and release the first key, then press another key. For example, **ESCAPE H**.
Syntax:	
`Fixed width`	A fixed-width font is used in syntax, code examples, system output, and file names.
`Fixed-width italics`	Fixed-width italics indicate variables in syntax.
`Fixed-width bold`	Fixed-width bold italic indicates variables in syntax with special emphasis.
`UPPERCASE fixed width`	ABL keywords in syntax and code examples are almost always shown in upper case. Although shown in uppercase, you can type ABL keywords in either uppercase or lowercase in a procedure or class.
Period (.) or colon (:)	All statements except `DO`, `FOR`, `FUNCTION`, `PROCEDURE`, and `REPEAT` end with a period. `DO`, `FOR`, `FUNCTION`, `PROCEDURE`, and `REPEAT` statements can end with either a period or a colon.
[ ]	Large brackets indicate the items within them are optional.
[ ]	Small brackets are part of ABL.
{ }	Large braces indicate the items within them are required. They are used to simplify complex syntax diagrams.
{ }	Small braces are part of ABL. For example, a called external procedure must use braces when referencing arguments passed by a calling procedure.
|	A vertical bar indicates a choice.
. . .	Ellipses indicate repetition: you can choose one or more of the preceding items.

Examples of syntax descriptions

In this example, `ACCUM` is a keyword, and *`aggregate`* and *`expression`* are variables:

Syntax

```
ACCUM aggregate expression
```

FOR is one of the statements that can end with either a period or a colon, as in this example:

```
FOR EACH Customer NO-LOCK:
  DISPLAY Customer.Name.
END.
```

In this example, STREAM *stream*, UNLESS-HIDDEN, and NO-ERROR are optional:

Syntax

```
DISPLAY [ STREAM stream ] [ UNLESS-HIDDEN ] [ NO-ERROR ]
```

In this example, the outer (small) brackets are part of the language, and the inner (large) brackets denote an optional item:

Syntax

```
INITIAL [ constant [ , constant ] ]
```

A called external procedure must use braces when referencing compile-time arguments passed by a calling procedure, as shown in this example:

Syntax

```
{ &argument-name }
```

In this example, EACH, FIRST, and LAST are optional, but you can choose only one of them:

Syntax

```
PRESELECT [ EACH | FIRST | LAST ] record-phrase
```

In this example, you must include two expressions, and optionally you can include more. Multiple expressions are separated by commas:

Syntax

```
MAXIMUM ( expression , expression [ , expression ] ... )
```

In this example, you must specify MESSAGE and at least one *expression* or SKIP [(*n*)], and any number of additional *expression* or SKIP [(*n*)] is allowed:

Syntax

```
MESSAGE { expression | SKIP [ ( n ) ] } ...
```

In this example, you must specify {*include-file*, then optionally any number of *argument* or &*argument-name* = "*argument-value*", and then terminate with }:

Syntax

```
{ include-file
    [ argument | &argument-name = "argument-value" ] ... }
```

OpenEdge messages

OpenEdge displays several types of messages to inform you of routine and unusual occurrences:

- **Execution messages** inform you of errors encountered while OpenEdge is running a procedure; for example, if OpenEdge cannot find a record with a specified index field value.
- **Compile messages** inform you of errors found while OpenEdge is reading and analyzing a procedure before running it; for example, if a procedure references a table name that is not defined in the database.
- **Startup messages** inform you of unusual conditions detected while OpenEdge is getting ready to execute; for example, if you entered an invalid startup parameter.

After displaying a message, OpenEdge proceeds in one of several ways:

- Continues execution, subject to the error-processing actions that you specify or that are assumed as part of the procedure. This is the most common action taken after execution messages.
- Returns to the Procedure Editor, so you can correct an error in a procedure. This is the usual action taken after compiler messages.
- Halts processing of a procedure and returns immediately to the Procedure Editor. This does not happen often.
- Terminates the current session.

OpenEdge messages end with a message number in parentheses. In this example, the message number is 200:

```
** Unknown table name table. (200)
```

If you encounter an error that terminates OpenEdge, note the message number before restarting.

Obtaining more information about OpenEdge messages

In Windows platforms, use OpenEdge online help to obtain more information about OpenEdge messages. Many OpenEdge tools include the following Help menu options to provide information about messages:

- Choose **Help** > **Recent Messages** to display detailed descriptions of the most recent OpenEdge message and all other messages returned in the current session.
- Choose **Help** > **Messages** and then type the message number to display a description of a specific OpenEdge message.
- In the Procedure Editor, press the **HELP** key or **F1**.

On UNIX platforms, use the OpenEdge `pro` command to start a single-user mode character OpenEdge client session and view a brief description of a message by providing its number.

To use the pro command to obtain a message description by message number:

1. Start the Procedure Editor:

   ```
   OpenEdge-install-dir/bin/pro
   ```

2. Press **F3** to access the menu bar, then choose **Help** > **Messages**.
3. Type the message number and press **ENTER**. Details about that message number appear.
4. Press **F4** to close the message, press **F3** to access the Procedure Editor menu, and choose **File** > **Exit**.

1

Resource Monitoring with OpenEdge Management

OpenEdge® Management provides you with the means to monitor your network resources and your local and remote system, file, database, and OpenEdge® server (AppServer™, NameServer, and WebSpeed® Transaction Server) resources. You receive feedback on each resource OpenEdge Management is monitoring, based on performance criteria that you establish. If a resource's performance does not meet the criteria you set, you determine the action OpenEdge Management takes.

For details, see the following topics:

- Setting up OpenEdge Management
- Understanding OpenEdge Management terminology
- Using the OpenEdge Management or OpenEdge Explorer console
- Identifying OpenEdge Management resources you can monitor
- Viewing resource monitor performance details
- Creating and scheduling OpenEdge Management jobs
- Exporting and importing OpenEdge Management monitoring components

Setting up OpenEdge Management

When you open OpenEdge Management in a browser after installation, you select some initial configuration settings, such as the password you want to use for the management console and whether or not you want OpenEdge Management to start automatically when the AdminServer starts. You decide where you want the OpenEdge Management Trend Database to be located and what Web server port number to use.

You also specify an e-mail server and a default operator who receives e-mail when OpenEdge Management generates an alert and you have set *alert action* as the Default_Action. The alert action is the action that should occur, by default, when an alert triggers. Other initial configuration choices you make affect SNMP adapter settings (if you have the SNMP adapter installed), and if you want OpenEdge Management to monitor and trend file systems and disks.

Once you make these initial decisions, you can set up OpenEdge Management to monitor remote resources. You also add OpenEdge Management users in one of two roles: as administrators, with access to all OpenEdge Management functionality, or as operators, with limited functionality access.

You work with OpenEdge Management through the browser-based OpenEdge Management console. If your user role in OpenEdge Management is that of an operator, you might find that some management console links or options are gray and unavailable to you. Note that all descriptions and graphics provided in this guide are presented with the assumption that you have full access to OpenEdge Management functionality.

OpenEdge Management also provides a command-line interface that allows you to:

- Start, query, and stop OpenEdge Management
- Dump the contents of the OpenEdge Management configuration database to a readable form (an XML file) and, in the event of a catastrophic failure, use the backup dump file to restore the database
- Enable and disable polling
- Work with alert commands
- Access command-line help

OpenEdge Management also includes the Database Administration Console, which you can use to manage and work with databases enabled for multi-tenancy. For more information, see *OpenEdge Management and OpenEdge Explorer: Configuring Multi-tenancy.*

OpenEdge Management and OpenEdge Explorer

OpenEdge Explorer runs in the OpenEdge Management console, enabling you to set configuration properties, to start or stop, and to view the status or log files for the following OpenEdge resources:

- AppServer
- AppServer Internet Adapter
- Databases
- DataServers (Oracle and MS SQL Server)
- NameServer
- Progress Application Server for OpenEdge (PAS for OpenEdge)
- SonicMQ Adapter

- WebSpeed Messengers
- WebSpeed Transaction Server
- Web Services Adapter
- OE Web Server

All configuration settings that you establish or change in OpenEdge Explorer are automatically reflected in the corresponding property file—`config.mgr` for databases and `ubroker.properties` for AppServer, WebSpeed Server, NameServer resources, DataServers, and the supported adapters.

When working with OpenEdge Explorer, you can also add users in the following roles:

- Administrators who have access to all OpenEdge Explorer functionality
- Operators who only have access to limited functionality

You work with OpenEdge Explorer through the browser-based management console. If your user role is that of an operator, you might find that some management console links or options are gray and unavailable to you.

The focus of this guide is on using OpenEdge Management to monitor OpenEdge resources. For details about using OpenEdge Management or OpenEdge Explorer to set properties, see *OpenEdge Management and OpenEdge Explorer: Configuration*.

For details about using the Database Administration Console to manage and work with databases enabled for multi-tenancy, see *OpenEdge Management and OpenEdge Explorer: Configuring Multi-tenancy*.

Understanding OpenEdge Management terminology

As you work with OpenEdge Management, there are several terms with which you will become familiar. These terms are described in the following sections.

Commonly used terms

Terms commonly used in OpenEdge Management resource monitoring include:

- Resource
- AdminServer
- Collection
- Monitor and schedule
- Rules, alerts, and actions
- Resource monitoring plan
- OpenEdge Management Trend Database

Resource

A *resource* is a specific component of your configuration. A resource can be:

- A database or database log file

- A system resource, such as CPU, memory, disk, or file system
- A network resource, such as TCP or UDP port, Ping (ICMP), or HTTP communication
- A file resource, such as a log file or other file you identify for monitoring
- An OpenEdge server component, such as an AppServer, NameServer, or WebSpeed Transaction Server
- DataServer for Oracle and MS SQL Server
- AppServer Internet Adapter
- SonicMQ® Adapter
- Web Services Adapter
- OE Web Server
- WebSpeed® Messenger

Note: OpenEdge Management supports monitoring and managing the WebSpeed® Transaction Server product. Throughout this guide, references to the WebSpeed Transaction Server and WebSpeed are used interchangeably.

When you install OpenEdge Management, resource monitors for your CPU resource and memory resource are automatically created and assigned default values, which you can change. OpenEdge Management also creates monitors for the following resources registered with the AdminServer: databases, AppServers, NameServers, and WebSpeed Transaction Servers.

OpenEdge Management neither creates resource monitors for nor supports monitoring of the following resources: AppServer Internet Adapters, DataServers (Oracle, MS SQL Server), SonicMQ Adapters, WebSpeed Messengers, OE Web Server, or Web Services Adapters. You can, however, set and modify property settings for these resources, as well as for databases, AppServers, NameServers, and WebSpeed Transaction Servers. You can also start and stop them (if applicable) or view their log files. For more information, see *OpenEdge Management and OpenEdge Explorer: Configuration*.

AdminServer

An *AdminServer* represents a named instance of an AdminServer that either is running OpenEdge Management or has been configured to be monitored by OpenEdge Management. There is typically a one-to-one relationship between the host name and the AdminServer name, unless there are multiple AdminServers running OpenEdge Management on the same host.

AdminServers are also available in OpenEdge Explorer.

Collection

A *collection* is a user-defined group of resources.You can create and use a collection to better organize and operate on resources. For example, you might create a collection known as Collection A that includes all resources on which a particular application depends. You might then create another collection known as Collection B, which is also dependent on one of the resources in Collection A. If the resource common to both Collection A and Collection B fails, the failure is reflected in the status of both collections, enabling you to determine quickly the extent and the impact of the failure.

An administrator can create a *private* collection or a *shared* collection. An operator can create a private collection and can see, but not create, a shared collection.

A collection can include any number of the following resources:

- Monitored resources, such as CPU, disk, memory, and database
- Non-monitored resources, such as AppServer Internet Adapters, DataServers (Oracle, MS SQL Server), SonicMQ Adapters, WebSpeed Messengers, OE Web Server, or Web Services Adapters
- Jobs
- Reports
- Other collections

Monitor, schedule, rules, and rule sets

A *monitor* in OpenEdge Management is defined as the combination of a resource, schedules, and rules. The *schedule* defines a block of monitoring time, and the *rules* determine how a resource's performance is judged while it is being monitored. *Each* rule verifies if a resource complies with its performance criteria. Rules are considered broken when a resource is not in compliance with the criteria that you set. When you are working with database, log file, AppServer, NameServer, and WebSpeed monitors, you can also use *rule sets* to monitor performance.

When you monitor a resource, you set up criteria by which you can keep track of the resource's performance. You can adjust the criteria against which performance is measured to meet your expectations as necessary. For example, you might want to monitor a database twenty-four hours per day, seven days a week, and receive notification if the database shuts down abnormally.

While you cannot create a resource monitor for AppServer Internet Adapters, DataServers (Oracle, MS SQL Server), SonicMQ Adapters, WebSpeed Messengers, OE Web Server, or Web Services Adapters, you can create rules and use rule sets for instances of them. This allows you to establish conditions under which OpenEdge Management issues an alert to let you know that particular rule has been violated.

Alerts and actions

A broken rule generates an *alert*, which notifies you that some criteria have been violated.

OpenEdge Management allows you to set up *actions* that trigger automatically in response to alerts. For example, you can set up an action so that the system administrator will receive an e-mail if a database experiences an abnormal shutdown.

Resource monitoring plan

A *resource monitoring plan,* also known more simply as a *monitoring plan,* defines a block of time in which a specific resource is monitored and identifies the rules to be checked during the defined time frame. A resource that you create in OpenEdge Management must have one or more monitoring plans in place before OpenEdge Management can perform monitoring.

You can create monitoring plans for all OpenEdge Management resources with the exception of AppServer Internet Adapters, WebSpeed Messengers, and remote NameServers.

OpenEdge Management uses monitoring plans to determine when to monitor a resource and which rules and criteria to use when evaluating the resource's compliance within these defined parameters. By defining more than one monitoring plan, you can specify different criteria (rules) for different times. For example, you might evaluate the rules for a database resource every five minutes Monday through Friday and only every thirty minutes on Saturday and Sunday.

OpenEdge Management Trend Database

OpenEdge Management allows you to store *trend data*, which is the monitoring information OpenEdge Management collects, in either a local or remote OpenEdge Management Trend Database. If you choose to send trend data to a local database (the default), you specify the trend database location and the port used to connect to that database. If you choose to use a remote database, you specify the hostname and Web server port of the remote OpenEdge Management Web server. The trend database must be locally configured at the remote location. All values you enter for either option are validated. If you set up OpenEdge Management to collect trend data about a resource and you later delete that resource monitor, the monitor's trend data already collected in the Trend Database is not deleted.

You can collect and store trend data for the following resources:

- Databases
- AppServers
- NameServers
- WebSpeed Transaction Servers
- System resources

Using the OpenEdge Management or OpenEdge Explorer console

OpenEdge Management and OpenEdge Explorer share the same graphical user interface, the management console. The links you see in the console are the same whether you are using OpenEdge Management or OpenEdge Explorer.

However, while certain links are always visible, they are only active and available if you install OpenEdge Management. Within OpenEdge Management, any links related to the SNMP Adapter are active only if you have a license for and install the adapter.

The management console also contains the Database Administration Console, which you can use to manage and work with databases enabled for multi-tenancy.

After you install OpenEdge Management or OpenEdge Explorer, you can do one of the following to access the management console logon window:

- To access OpenEdge Management or OpenEdge Explorer from a Web browser, enter the URL `http://host:port` in the address or location field. The host is the name of the machine on which OpenEdge Management or OpenEdge Explorer is installed, and the port is the Web server port (by default, this port is 9090). Note that you can also type `localhost`.
- If you have OpenEdge Management/OpenEdge Explorer installed, choose **Start** > **Programs (or All Programs)** > **Progress** > **OpenEdge** > **Management Console** in the Windows platform.
- If you have only OpenEdge Explorer installed, choose **Start** > **Programs (or All Programs)** > **Progress** > **OpenEdge** > **OpenEdge Explorer** in the Windows platform.

The first time you log on to OpenEdge Management or OpenEdge Explorer, you use the default user name `admin` and the default password `admin`.

For a complete description of the management console, see *OpenEdge Management and OpenEdge Explorer: Getting Started*.

Identifying OpenEdge Management resources you can monitor

With OpenEdge Management, you can monitor the performance of local and/or remote resource types, provided you run the OpenEdge Management Remote Configuration Utility. The utility sets up monitoring of the remote resources shown in the following table.

You can also establish or modify property settings for remote resource types, once you run the OpenEdge Management Remote Configuration Utility.

For details about the Utility, see *OpenEdge Management and OpenEdge Explorer: Getting Started*.

Table 1: Resources for local and/or remote monitoring

Resource	Local	Remote
Log file	Yes	Yes
Other file	Yes	Yes
Network components (TCP, UDP, PING, HTTP)	Yes	Yes
CPU	Yes	Yes
Memory	Yes	Yes
OpenEdge server components (AppServer brokers, NameServers, WebSpeed Transaction Servers)	Yes	Yes
Disks	Yes	Yes
File systems	Yes	Yes

You can also monitor managed databases (recognized by the AdminServer and OpenEdge Management) or scripted databases, and you can migrate scripted databases to become managed databases.

Monitoring a resource

You select the resource you want to monitor, and create its monitoring plan. The **overall** steps, which are covered in greater detail throughout this guide, are provided here.

Once you install OpenEdge Management and set the initial configuration options, you can configure OpenEdge Management to monitor certain remote resources as well as local ones.

You can monitor the following resources in OpenEdge Management:

- A database or database log file
- A system resource, such as CPU, memory, disk, or a file system
- A network resource, such as TCP or UDP port, Ping (ICMP), or HTTP communication
- A file resource, such as a log file or other file you identify for monitoring

- An OpenEdge server component, such as an AppServer, a NameServer, or a WebSpeed Transaction Server

To monitor a resource in OpenEdge Management:

1. Create a resource monitor (for a database, file, network component, OpenEdge server component, or system component).
2. Create a monitoring plan for the resource monitor. The monitoring plan defines all the criteria you want to use to assess the resource's performance.

   The monitoring plan consists of the following parts:

   - One or more schedules that identify the time frame you want to monitor.
   - One or more rules that identify the performance standards you set for the resource. If a standard established by a rule is not met, the rule is considered broken and OpenEdge Management generates an alert to notify you of the performance issue.
   - An action that occurs when OpenEdge Management generates an alert about a performance issue. The following table shows a list of the available actions.

Table 2: OpenEdge Management actions

In response to this action . . .	OpenEdge Management . . .
Default action	Contains the e-mail action.
Email action	Generates an e-mail to a recipient you identify, notifying the recipient about the alert.
Log action	Records details about the alert in a log file.
SNMP trap action	Sends notification of the alert to your SMNP management console.
Compound action	Executes an action that consists of one or more other actions. A created job, for example, is a compound action.

Viewing resource monitor performance details

You can acquire information in OpenEdge Management about resource performance in the following ways:

- Create a collection of resources. The collection functions as a group, and each resource is a member. When you create a collection, you typically choose resources with a common characteristic. For example, you might create a collection of databases that span multiple hosts but support one particular application.
- Create one or more custom collection views that display only the details you want to see. For example, you can create a custom view of all monitoring plans you have in place, or find out a number of specifics about a particular database. The details that the custom view contains are determined by you and appear in individual panes knows as *viewlets*. Each viewlet provides details in a graphical format and allows you to get a quick status of only those pieces of information you find key at any particular time.

  OpenEdge Management provides you with a default custom view, known as the **Dashboard** page.

- Create and run a report that includes performance data (in both graphical and HTML formats) about one or more resources. When you create a report, you can use one of the report templates that OpenEdge Management provides, or you can create your own report template. Creating a report and running it do not have to be done at the same time; you can create a number of reports and schedule them to run at a particular time. You can review a list of reports you have scheduled to run and a list of reports that are currently running.

  Using the **Realtime Reports** option you can also run a real-time report with a specific focus, such as a database summary or statistics about open alerts (in detail or summary format). See *OpenEdge Management: Reporting* for details about creating and running OpenEdge Management reports.

Creating and scheduling OpenEdge Management jobs

As you monitor certain resources with OpenEdge Management, you can define tasks for execution at intervals you determine. Each task is known as a *job*.

OpenEdge Management supports three kinds of local and remote jobs:

- **Custom job** — A task that you define for execution at a scheduled interval
- **Job template** — A set of default values that you use when creating a job instance
- **Predefined database maintenance job template** — A tailored template that performs a maintenance activity for an OpenEdge database

Exporting and importing OpenEdge Management monitoring components

You can export or import certain OpenEdge Management resource monitoring components among machines, which allows you to establish common monitoring guidelines without having to recreate manually the monitoring structure you want. These exportable and importable components reside in the OpenEdge Management Library, making them available for sharing.

You can export or import the following monitoring components:

- Actions
- Report templates
- Job templates
- Database rule sets
- Log file rule sets
- NameServer rule sets
- AppServer rule sets
- WebSpeed rules sets
- Log file search criteria

2

Creating Collections and Collection Views in OpenEdge Management

In OpenEdge Management you can create one or more groups of resources, known as *collections*. When you create a collection, you typically group resources together that have a common element. For example, you can group all the database resources used by one division in your company, or all the resources that support one particular application. When you group resources, you can more easily examine their details and assess their status.

You can also optionally create one or more custom views within a collection and specify exactly what types of information you want to see. The information provides a customized view, sometimes in a graphical format, of your various resources' status.

For details, see the following topics:

- The Dashboard page
- Collections
- Collection terminology and hierarchy
- Sample collection scenario
- Working with collections
- Working with collections as labels
- Collection views

The Dashboard page

OpenEdge Management creates a default private home collection for each OpenEdge Management user. Click **Dashboard** > **Go to Dashboard** in the management console menu to see the default view named **Default** in the default private home collection named **My Collections.Home**.

The **Dashboard** page shows the **Default** view tab in the **Home** collection and consists of the following individual groups of data, known as *viewlets*:

- **Adminserver status** — Displays the number of AdminServers that are currently running and offline.
- **Resources with alerts** — Displays the resources that have alerts.
- **CPU** — Provides the CPU usage.
- **Memory** — Provides the memory usage.
- **Resource status legend** — Describes what the resource status icons on the page represent.
- **Alert severity legend** — Identifies the various icons used to indicate the severity of an alert.

Hover over the **Default** view tab to:

- Edit Collection
- Edit View Content
- Edit View Layout

Note that you cannot delete the **Default** view.

You can create a new view by clicking the **+** icon next to the **Default** view tab.

How collections and collection views differ

A collection is a logical grouping of resources and views, and a view is a display of user-selected data. The visualization of any collection is one or more views.

Collections and collection views differ in the following ways:

- You can create a collection with an arbitrarily large number of members. However, displaying them all in a single OpenEdge Management page might not be practical. Instead, you can create multiple views of a collection and each view can focus on a different aspect of the collection.
- A collection is hierarchical. You choose the parent for each collection when you create it. This allows you further flexibility in isolating specific performance data.
- You can create views that include nonresource viewlets, such as resources with alerts or active monitoring plans. Neither of these views corresponds to a single resource and would not equate to a member in a collection, which consists of a group of individual resources.

Collections

You create, edit, and delete collections from the **Dashboard** page.

What is a collection

A collection is a group of resources (known also as members) that you define. You group resources into a collection based on criteria that you find valuable when grouped and viewed together.

You can include the following resources in a collection:

- All monitored resources, such as CPU, disks, databases, or WebSpeed Transaction servers. A resource can be a member of multiple collections.
- Jobs
- Reports
- Other collections

When you create a collection, you assign it a name and provide its definition. You can identify the collection members when you create the collection, and later you can add or remove members. However, when adding members to a new or an existing collection, you cannot add a private collection to a shared collection and vice-versa. You can include resources from multiple AdminServers in a collection. All collections are displayed in the **Resources** grid frame.

The OpenEdge Management console displays a list of collections that appear as labels in the details page of each resource. You can also associate members to other collections or create a new collection from a resource's details page.

When you select members for a collection, you can use a wildcard feature that enables you to select, for example, all resources, all resources of a particular type, or all resources on all AdminServers. Using the wildcard is a powerful way to manage the members of a collection. For example, if you select all databases on all AdminServers as members of a collection, databases added to AdminServers after the collection was created are automatically added as members of the collection. Likewise, any database that is removed from an AdminServer is also removed automatically from the collection.

When you look at the status of a collection, you see a rollup of the status of all members of the collection. For example, consider a scenario in which you host an application that is used by multiple customers. The application uses multiple databases and OpenEdge brokers, and it spans several hosts. If your customers contract for premium services that include dedicated hosts, databases, and/or brokers, you can group resources based on the customers using them. In so doing, you can quickly assess the status of the application for any of the supported customers.

Collection terminology and hierarchy

Collections are categorized as follows:

- **My Collections** —Private collections that you create and that are available only to you.
- **Shared Collections** — Collections that are created by administrator users. Administrator users can view, reference, edit, and delete shared collections. Operator users can only view and reference shared collections.
- **Referenceable collection** — Any collection that you tie as a child to either **My Collections** or **Shared Collections** as a parent. When tied to **My Collections**, the collection becomes a private referenceable collection. When tied to **Shared Collections**, the collection becomes a shared referenceable collection.

The advantage of creating a referenceable collection is that you can reference the collection from more than one place. For example, consider the case of a system resources collection whose members include CPUs, memory, and file systems for three different AdminServers. You can create this system resources collection once, and tie it to a Shared Collection and a private collection. You create it once, in one collection, and create a reference to it from another collection. It is unnecessary to create the collection twice.

The referenceable collection can be referenced without being copied. This provides a single point of collection maintenance.

- **Child collection** — A collection whose parent was chosen through the **Parent** property. For details, see Creating a collection on page 33.

Sample collection scenario

When you create a collection, consider which resources it makes sense to group. Remember that a resource can be a member of more than one collection, so you have significant flexibility in how you organize collection members.

Consider the case of a large company that has OpenEdge Management monitoring multiple OpenEdge-based applications, such as Payroll, Finance, Accounts Receivable, and Shipping. These applications also span multiple hosts, and in some instances, the applications also share resources.

By using OpenEdge Management's collections feature, an administrator at the company can group resources based on the applications that use them, as follows:

- Payroll Collection:
  - Finance database on Host A
  - Payroll database on Host B
  - WebSpeed Broker on Host B
  - System resources on Hosts A and B
- Finance Collection:
  - Finance database on Host A
  - Accounts Receivable database on Host C
  - AppServer Broker on Host A
  - System resources on Hosts A and C
- Accounts Receivable Collection:
  - Accounts Receivable database on Host C
  - AppServer Broker on Host C
  - System resource on Host C
- Shipping Collection:
  - Shipping database on Host D
  - Accounts Receivable database on Host C

- WebSpeed Broker on Host D
- System resources on Hosts C and D

By looking at a particular collection, an administrator can easily determine the status of any application. Each application collection contains the set of resources that the application depends on. A failure in a resource is reflected in the status of the collections that use that resource.

Working with collections

You can create, edit, or delete a collection, as described in the following sections:

- Creating a collection on page 33
- Editing a collection on page 34
- Deleting a collection on page 34

Creating a collection

An administrator can create a private or shared collection. An operator can create only a private collection.

You can create a collection using the **Dashboard** option in the management console or the collection editor in a resource's details page. For more information on creating collections using the collection editor, refer to Creating a collection from a resource's details page on page 35.

To create a collection of OpenEdge Management resources using the **Dashboard** option:

1. Click **Dashboard** > **New** > **Collection** in the OpenEdge Management console menu.

   The **Create new Collection** page appears.

2. Provide the following properties:

   - **Name** — Type the name of the collection
   - **Parent** — From the drop-down list, choose one of the parent collections (thereby making the newly created collection a child) from the following table:

Value	Description
My Collections	Collections that are private and available only to you
My Collections.Home	A collection created by default for a user
Shared Collections	Collections available to all other users of OpenEdge Management

   - **Description** — An optional description of the collection.

3. Select the members of the collection by following these steps:

a) In the **List resource of type** field, you can choose resource types by either choosing from the available resources or selecting the wildcard (*). If you select the wildcard, all resource types are selected. In addition, any newly created resource will automatically be added as a member of the collection, and any deleted resource will automatically be deleted from the collection.

b) In the **for adminserver** field, you can choose the AdminServers either by choosing from those available or by choosing the wildcard. If you choose the wildcard, all AdminServers are selected. Any newly created AdminServer is automatically added as a member of the collection, and any deleted AdminServer is automatically deleted from the collection.

Note that you can also choose all available referenceable collections by choosing **collections** in the **for AdminServer** field.

4. Click **Apply Filter**. The list of collections that match the filter specifications appear in the **Available** list.
5. Select the resources, as follows:

   - To select all the databases as members of the collection, click **Select All**. Then click the right arrow.
   - To choose one or more individual databases, highlight and click each database you want to include. Then click the right arrow.

   The databases appear in the **Selected** list.

6. Click **Save**. The default view page for that new collection appears.

Note: OpenEdge Management displays collections for each its members, in the form of labels. You can view the list of collections in the **Resources** page and also as labels in the resource details page. For more information on accessing and viewing collections appearing as labels, refer to Accessing and viewing collections as labels on page 35.

Editing a collection

You can edit a private collection, and if you are an administrator you can also edit a shared collection.

To edit a collection:

1. On the OpenEdge Management console menu bar, select the collection from **Dashboard** > **Collections**. You can also select the collection from the **Collection** tab from the menu on the left.

   The **Dashboard** page displays the collection in a tab.

2. Hover over the collection tab and select **Edit Collection**.

   The **Edit Collection** page appears.

3. Edit the collection as required and then click **Save**.

Note: Removing a member from a collection also dissociates the collection name from the list displayed in the **Resources** section and the details page of that member.

Deleting a collection

You can delete a collection you no longer need.

To delete a collection:

1. On the OpenEdge Management console menu bar, select the collection from **Dashboard** > **Collections**. You can also select the collection from the **Collection** tab from the menu on the left.

   The **Dashboard** page displays the collection on a tab.
2. Hover over the collection tab and select **Delete Collection**.
3. Click **Yes** in the confirmation dialog box to confirm deletion.

Note: Deleting a collection removes all resources that are members of that collection. Performing this action also removes the collection name from the list of collections displayed for all members of that collection.

Working with collections as labels

Collections appearing as labels in OpenEdge Management enable you to access and view a private or shared collection's details page. Using the collection editor icon on a resource's details page, you can:

- Create a new private or shared collection
- Associate a resource as a member of an existing collection

You can delete a private or shared collection from the details page of a resource.

Note: Similar to creating and deleting collections from the **Dashboard** page, an administrator can create and delete collections from the details page of a resource. An operator must be provided proper authorization to perform such an operation. Neither the administrator nor the operator can delete the default collection. For authorizing users to create or delete collections, refer to *OpenEdge Management and OpenEdge Explorer: Getting Started*.

Accessing and viewing collections as labels

You can view both the default and user-defined, private or shared collections associated with a resource either on the **Resources** page or the resource's details page.

If you select the database resource on the **Resources** section, the collections are displayed on the right of the resource's details page.

You can access the details page for a particular collection by clicking its corresponding name either from the **Resources** page or from the resource's details page instead of accessing it from the **Dashboard** page of your console.

Creating a collection from a resource's details page

You can create an entire new collection from a resource's details page. You can create only shared collections using this approach.

To create a collection from the details page of a resource:

1. From the **Resources** page, click the resource's name to go its details page where you want to create a collection.
2. Click  to add the resource to a new or existing collection. A blank field appears along with a list of existing collections.
3. Do one of the following:
   - Type an alphanumeric text of your choice in the field for collection editor. The console filters the list with the matching keywords.
   - Click a collection name from the drop-down list.
4. Press **Enter**. The collection name appears on the details page.

Note: If you type a name that does not exist, the console automatically creates a collection with that name and the current resource will be a member of that collection.

You can view the details page of the added collection by clicking the collection name.

Note: You can follow the same procedure to associate other resources to an already existing resource by choosing an existing collection name from the list of collections that appear when creating a collection. You can also add members to the collection by editing a collection as described in Editing a collection on page 34.

Deleting a member of a collection from a resource's details page

Deleting a collection from the details page of a resource dissociates it from being a member of that collection. You can only delete the user-defined shared collections.

To delete a collection from the details page of a resource:

1. From the **Resources** page, click the resource's name to go to its details page where you want to delete the collection.
2. Hover over the collection you want to delete and click . The collection is removed from the details page as well as from the **Resources** page.

Note: If the red cross does not appear when you move your cursor over a collection name, it indicates that you do not have the privileges to remove the current member from the collection.

If you delete a collection from the **Dashboard** page, the collection is removed from the **Resources** page and the resource's details page. For more information, see Deleting a collection on page 34.

Collection views

You can create, edit, and delete a collection view from the detail menu bar. A collection view consists of individual viewlets. Each viewlet summarizes particular content, such as which resources have alerts or which active monitoring plans exist.

Working with collection views and viewlets

Each collection view contains one or more viewlets. There are three types of viewlets: resource, standard, and view panel.

Resource viewlets

A resource viewlet is specific to a single instance of a resource.

Standard viewlets

A standard viewlet is predefined by OpenEdge Management for displaying content from multiple resources or other aspects of OpenEdge Management.

Standard viewlets that OpenEdge Management provides are as follows:

- **Running reports** — The reports currently running shown in tabular format. You can customize the view by selecting the fields to display in each report.
- **Active monitoring plans** — A list of all active resource monitoring plans.
- **Resources with alerts** — A list of resources with open alerts as well as summary information, such as the date and the time of the last occurrence. Each resource name is preceded by its AdminServer name.

  You can customize the view by selecting the fields to display for each resource.
- **Resource status legend** — An explanation of what certain data in the viewlet represents.
- **Running jobs** — The jobs currently running shown in tabular format. You can customize the view by selecting fields to display for each job, or by limiting the display to jobs running longer than some period of time.
- **Alert severity legend** — An explanation of the various icons used to indicate the severity of an alert.
- **Adminserver status** — The number of AdminServers that are currently running and offline.
- **Collection members** — A list of all members of the collection, including other collections. A description of the member also appears, if available.

View panel viewlets

A view panel viewlet displays a view for any other collection being referenced by the active collection. Your ability to define views in one collection that you can display in another collection contributes to centralized, streamlined maintenance. This ability also eliminates the need for creating duplicate viewlets of the same data.

Viewlet characteristics

Many viewlets share common characteristics. Some of those characteristics are as follows:

- The title bar of a resource viewlet provides basic status information about the resource. This includes the most recent polling status (**Passed** or **Failed**, for example) as well as whether the resource has any outstanding alerts.
- If the resource does have any outstanding alerts, the title bar color is set to the worst-case severity of the open alerts for that resource. For example, if a resource has an outstanding alert with a severity of **Error**,

the title bar is yellow as you view it in your browser. Note that it is possible for a resource to have a status of **Passed** but still have outstanding alerts.

- The viewlet provides information in both tabular and graphical format.
- As the legend indicates, the percentage of time the resource is busy for the represented time frame.
- The **Customize Viewlet** icon on the top-right of the viewlet allows you to choose how the viewlet displays data. For details, see Customizing a viewlet on page 39.

Creating a collection view

You can create a new collection view by selecting the collection in which you want to create a new view from **Dashboard** > **Collections** in the OpenEdge Management console menu.

To create a new collection view in the collection you selected:

1. Click the **+** icon next to the **Default** collection tab.

   The **Create new Collection View** page and a **New View** collection tab appears.
2. Type the name of the view in the **Name** field.
3. Select one or more options in the **Standard viewlets to show** area to identify the standard viewlets you want to display in the view.
4. Choose the resource viewlets you want to display for each individual database resource by selecting one or more of the options in the **Resource viewlets to show** area.

   Each of these viewlets supplies detailed information about some aspect of a database.
5. Choose any other resource viewlets you want to include in the collection view.
6. Choose the view panel viewlets you want to appear (if available).
7. Click **Save** when you finish making your selections.

   The selected resource and standard viewlets appear in the collection.

Customizing the view layout

You have the option of overriding the default view column layout and the order in which the viewlets appear in a column.

To customize the layout of a collection view:

1. Select the view tab you want to customize in the collection.
2. Hover over the view tab and select **Edit View Layout**.The **Customize Layout** page appears.

   The **Left Panel** box lists the viewlets that appear in the left panel of the layout, and the **Right Panel** box lists the viewlets that appear in the right panel of the layout. The order of the viewlets as listed in a panel corresponds to the order in which the viewlets appear in the collection view.

   The width of the panel in the layout is determined by its contents.
3. Use the directional buttons to move viewlets either within a panel or from one panel to the other, as described in the following table.

Table 3: Moving the viewlets

Click this button . . .	To move the viewlet . . .
	Up one position in the panel
	Down one position in the panel
	From the left panel to the right panel
	From the right panel to the left panel
	Out of the collection view (which deletes the viewlet)

4. Click **Save** when you finish organizing the viewlets. The collection view reappears with the re-organized viewlets.

Editing a collection view content

You can change the content of a collection view.

To edit a collection view:

1. Hover over the view tab and select **Edit View Content**.
2. Make the changes you want.
3. Click **Save**. The view reappears with the updated viewlets content.

Deleting a collection view

You can delete a view you no longer need. However, be very cautious when you do this, and be sure you are deleting the currently active view and not one that you were using previously.

To delete a collection view:

1. Select the view you want to delete from the collection.
2. Hover over the view tab and select **Delete View**.
3. Click **Yes** in the confirmation dialog box to confirm deletion.

Customizing a viewlet

In addition to choosing the content and layout for a view, you can also customize many of the individual viewlets.

A customizable viewlet has the **Customize Viewlet** icon  in the upper right corner. For example, when you click the **Customize Viewlet** icon on a **Disk** viewlet, the **Customize Viewlet** page appears.

To customize the viewlet:

1. In the **Select viewlet components to display** area, select only the viewlet components you want to appear. Clear any components you do not want to include. You can choose any or all of the following options:

- **Resource name in title** — The name of the resource
- **Resource status in title** — The status of the resource
- **Tabular data** — The data displayed in a table
- **Graph** — Whether to display a graph or not

2. In the **Select graph properties** area, select the following graph properties:
   - **Historical** — **line**, **column**, **area**, **stacked column** or **stacked area**

     This option sets the viewlet graph style for showing the graph points in time.
   - **Size** — **medium**, **large**, or **very large**
   - **Max display time**— Choose the display time from the drop-down list. The higher you set the maximum display time, the more processing time will be required to produce the data and the fewer details will be visually evident.

     For details about the graph cache, see *OpenEdge Management and OpenEdge Explorer: Getting Started*.
3. Set the following **Graph Visuals** properties:
   - **Legend** — An explanation of what the colors in the graph represent
   - **Grid** — A grid to help you examine data more closely
   - **Threshold** — Whether rule thresholds (if available and currently in use) are displayed on the graphs
4. Click **Save**. The **Dashboard** page appears with the customized viewlets.

3

Setting up Resource Monitoring

This chapter describes how to create a resource to monitor; define the resource's monitoring plan (including schedules and rules); establish meaningful threshold values for rules; and create actions that occur in response to OpenEdge Management alerts.

For details, see the following topics:

- Before you begin
- Monitoring remote resources
- Naming conventions
- Monitoring a system, network, or file resource: an overview
- Working with CPU and memory resources
- Selecting a system, network, or file resource to monitor
- Creating a resource monitoring plan
- Calculating meaningful thresholds with the Configuration Advisor
- Understanding alerts
- Determining OpenEdge Management's action in response to an alert
- Using default actions
- Editing or copying a resource monitor
- Deleting a resource monitor
- Working with default values

- Using job and report templates
- Using the OpenEdge Management Library

Before you begin

Examining current resource considerations and future operational needs helps you to make preliminary decisions about resource monitoring prior to setting up specific monitoring criteria in OpenEdge Management.

For example, consider the following as you evaluate your current resource-related operations and processes:

- Which employee at your company is typically notified that a resource is down, and what method do you typically use to contact that individual? Is there too much downtime between a resource trouble report and the response from personnel assigned to handle a given problem? Do problems stack up because your system does not prioritize and queue these problems to the correct person's attention?
- What is your typical weekly resource load? Can you define the peak load times for your applications? If so, do you have procedures in place to address the higher-than-average strain these events place on your resources? Similarly, do you also experience idle time? Do you adequately plan this slower-than-normal time into your schedule so that you can routinely run other processes at this time?
- Do you have a scheduling process that enables you to run your applications, and automatically notify interested personnel when these jobs are done? Can this process also alert you to a problem and take predefined, corrective action?

As these questions suggest, there are considerations you must address as you move from your present approach for determining your resource needs to systematically analyzing and fulfilling those needs using OpenEdge Management.

Monitoring remote resources

You can monitor remote disks, CPU, memory, log files, file systems, databases, and OpenEdge server components, such as AppServer, NameServer, and WebSpeed resources, after you run the Remote Monitoring Configuration Utility. See *OpenEdge Management and OpenEdge Explorer: Getting Started* for details about the utility.

Note: This guide does not describe monitoring databases remotely. See *OpenEdge Management: Database Management* for details about remote databases.

Naming conventions

Observe the following naming conventions when you create a resource monitor:

- You cannot have two resources of the same type with the same name.
- When you name a resource, the name must be HTML-compliant. HTML-compliant means that the value you enter in the **Name** field follows the formatting rules and conventions that the Hyper Text Markup Language allows. For example, you cannot enter spaces between words or use special characters such as an asterisk (*), an ampersand (&), or a period.

- When you delete a resource, the name used for that resource becomes available for use again.

Monitoring a system, network, or file resource: an overview

Using OpenEdge Management, you can monitor specific resources.

To monitor a resource:

1. Select the type of resource you want to monitor.
2. Define the resource monitor's properties. The properties you define depend upon whether you select a system, network, or file resource and can include name, description, file name, host, or port, for example.
3. Create a monitoring plan for the resource, which includes:
   - **A monitoring plan definition** — Identifies the schedule you want OpenEdge Management to follow when monitoring, how often you want OpenEdge Management to poll the resource, and whether you want to enable alerts for the resource. Alerts notify you of potential problems with the resource's performance.
   - **Rule definitions** — Specifies under what circumstances OpenEdge Management generates an alert for the resource and what action OpenEdge Management should take, if any, as a result of the alert. You can choose from several different actions, such as sending an e-mail to a specific recipient, or adding an entry to a particular log file.
   - **Rule sets** (for log files and OpenEdge server components) — Groups one or more rules in a logical sequence
   - **Schedules** — Schedules a block of monitoring time
4. Optionally, use the Configuration Advisor to calculate more meaningful threshold values for resource rules.

Monitoring a database or other OpenEdge resource

The information presented in this guide pertains primarily to system, network, and file resource types, unless otherwise noted. Much of the information, however, also applies to database and OpenEdge server resource monitors. For details about creating database resource monitors, see *OpenEdge Management: Database Management*. For details about creating resource monitors for OpenEdge server components, see *OpenEdge Management: Servers, DataServers, Messengers, and Adapters*.

Working with CPU and memory resources

Monitoring a CPU resource or a memory resource is essentially the same as monitoring any other resource with OpenEdge Management. The one difference, however, is that OpenEdge Management automatically sets up one of each of these resources with default values when it initially starts. You cannot delete either the CPU resource or the memory resource, nor can you create additional instances of either resource type.

To access detailed information about each of these predefined resources:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the required CPU or Memory resource. The resource details page appears and shows the default values.

   For more information about selecting a custom system, network, or file resource to monitor, see Selecting a system, network, or file resource to monitor on page 44.

Selecting a system, network, or file resource to monitor

OpenEdge Management automatically creates one CPU resource and one memory resource for each AdminServer. You use the management console to create additional system, network, or file resources to monitor.

Note: For the purpose of illustration, the procedures in this section describe how to create a disk resource monitor and then define its properties and monitoring plan.

To create a disk resource monitor:

1. In the OpenEdge Management console menu, select **Resources** > **New** > **Disk**.

   The **Create new Disk** page appears.

   If there are remote AdminServers online, a drop-down list of AdminServer names also appears. You can now specify the resource's properties.

2. Enter the name of the resource in the **Name** field. For all resources, this value must be no more than 32 characters. You cannot use spaces, periods, or special characters.
3. Enter a description of the resource in the **Description** field. You can enter a maximum of 256 characters.
4. From the **Disk** drop-down list, select the disk for which you are creating this resource monitor. The drop-down list displays the reference that your operating system uses to recognize each disk drive available on your system.
5. Select the **Enabled** option to set the resource monitor to begin monitoring immediately. This option is checked by default. If you want to create a monitor that does not immediately begin monitoring, clear the option. The monitor is then disabled until you enable it.
6. Click **Save**.

   The **Create Monitoring Plan** page appears. You can now create the resource's monitoring plan.

Creating a resource monitoring plan

When you create a monitoring plan, you provide both a monitoring definition and one or more rule definitions.

The monitoring definition includes:

- The schedule selected from all the schedules that are available. The schedule specifies when the plan is active, and the available schedules include only those schedules that will not overlap with schedules already in use for a resource.
- A polling interval, which indicates the frequency with which this resource's rules are checked by OpenEdge Management.
- An option to enable alerts for this resource. Enabling alerts means that alerts will be generated while the plan is active.
- An option to store, or *trend*, the performance data collected by this resource to the OpenEdge Management Trend Database while the plan is active. See Trending considerations on page 45 for details.
- The number of polls after which performance data will be trended to the OpenEdge Management Trend Database.

The rule definition identifies:

- When to generate an alert for the resource
- What level of alert severity will cause OpenEdge Management to generate an alert
- The number of failed polls after which OpenEdge Management generates an alert
- The action that OpenEdge Management performs when an alert is generated
- The number of successful polls after which OpenEdge Management clears an alert
- The action OpenEdge Management performs when an alert is cleared

Trending considerations

You can choose to have OpenEdge Management trend performance data to the OpenEdge Management Trend Database. Consider the following points when setting trending options:

- Different trend expectations exist for different resource types. Therefore, each trend interval has a different meaning depending on the data to be trended. Refer to the specific resource type in this guide for detailed information about data trending expectations for each resource type.
- If you set a trending value in the **Trend Performance Data every** field, OpenEdge Management will use the value you provide in combination with the value in the **Polling Interval** field to calculate the trending interval.

  For example, if you initially set a **Polling Interval** of 30 seconds and also set the **Trend Performance Data every** field to trend data every 3 polls, OpenEdge Management will automatically calculate the trending interval as 90 seconds. You will see this result when you view the summary of the resource monitor's performance.

Choosing a schedule for the monitoring plan

A schedule defines a block of time for which a set of monitoring rules is active for a resource. When you add a monitoring plan to a resource, you specify the schedule to indicate when the monitoring plan will be active.

OpenEdge Management provides several predefined schedules. You can use the schedules as they are, or you can modify them to suit your operating needs. You can also use an existing schedule as the basis for a new schedule, changing its name and characteristics.

Note: The schedules discussed in this section refer to schedules you associate with resource monitors. Schedules that define frequency of an activity are associated with jobs, job templates, reports, and report templates. For specific information about job schedules, see Creating Jobs and Job Templates. For specific information about report schedules, see the relevant sections in *OpenEdge Management: Reporting*.

Each schedule you create, along with the OpenEdge Management-provided schedules, appears under **Library** > **View** > **Schedules**.

General scheduling considerations

Note the following points regarding schedules:

- Before creating a new schedule, review the existing schedules to determine if you can use or modify an existing one.
- You must create two separate schedules if you want to span an overnight period. For example, if you want to create a schedule that includes 8:00 PM through 8:00 AM, you must create the following two schedules:
  - 8:00 PM to 12:00 AM
  - 12:00 AM to 8:00 AM

  You must also select the appropriate days for each schedule.
- You can schedule different rules to apply to a resource at different times and on different days. For example, you might want to monitor your OpenEdge databases more frequently during the week than on the weekends. To accomplish this, you can create and associate different schedules with each monitoring plan.
- You can associate more than one schedule with a resource, but you cannot associate overlapping schedules with a single resource.
- The selection list of available schedules for a monitoring plan does not include schedules that overlap the time or day criteria with other schedules or schedules that are already in use by the resource.

Using predefined schedules

OpenEdge Management provides several predefined schedules that you can use in your monitoring plan. You can also use any predefined schedule as the basis for a new schedule, which you can then modify to suit your needs. The following table describes the predefined schedules.

Table 4: Predefined schedules

Name	Block schedule defined for . . .
Default_Schedule	Sunday through Saturday, 24 hours per day. See Using the Default_Schedule on page 47 for more details.
Weekdays	Monday through Friday, 24 hours per day.
Weekdays_Early_Morning	Monday through Friday from 12 AM to 8 AM.
Weekdays_Evenings	Monday through Friday from 5 PM to 12 AM.
Weekdays_Office_Hours	Monday through Friday from 8 AM to 5 PM.
Weekends	Saturday through Sunday, 24 hours per day.

When you select **Library** > **View** > **Schedules** in the management console menu, the **Library.Schedules** page appears and displays a list of predefined and user-generated schedules.

Once you create a resource monitor, you can edit an existing schedule or create a new schedule to associate with it at any time.

Using the Default_Schedule

The **Default_Schedule** is a reusable 24/7 schedule that you can use or modify just as you can any other schedule, whether predefined by OpenEdge Management or created by you. However, the **Default_Schedule** is unique among all other schedules because:

- OpenEdge Management automatically associates it with new resources when you create them.
- You can edit it, but you cannot delete or rename it.

Updating and deleting schedules

Once your resource monitor exists, you can update its schedule at any time. This task is accomplished by accessing the monitoring plans for the resource and selecting another schedule from the drop-down list of currently available schedules.

You cannot delete a schedule that is currently being used by a resource monitor.

Creating a new schedule

You can create a new schedule from the **OpenEdge Management Component Library** page.

To create a new schedule:

1. In the OpenEdge Management console menu, choose one:
   - Select **Library** > **New** > **Schedule**.
   - Select **Library** > **Go to Library** > **Create Schedule**.

     The **Create new Schedule** page appears.
2. Enter a name for the schedule in the **Name** field. You must not use spaces or special characters such as an asterisk (*), ampersand (&), or period (.) in the **Name** field.
3. Enter a description of the schedule in the **Description** field. The **Description** field is for reference only. The description appears in the schedule's title when you view the schedule after saving it.
4. Under **Schedule period**, select the days and times you want the schedule to include (note that **Days** is a required field).
5. Click **Save** to save or **Cancel** to discard the schedule. If you save the schedule, the days and times of the created schedule are displayed.

   You can edit the schedule you created, copy and use it as the basis for another new schedule, or delete it.

Adding rules to the monitoring plan

Rule definitions identify the specific attributes of a resource that you want to monitor, the severity of an alert OpenEdge Management generates in response to performance, and the action that occurs when an alert is generated. Rule definitions for system, network, and file resource types are composed of the following elements:

- Specific rules you set for each resource type. For example, the rule definition for a disk resource includes setting an alert that triggers when disk activity exceeds a specified percentage.

  Refer to the appropriate sections in Monitoring System Resources on page 65, Monitoring Network Resources on page 73 or Monitoring File Resources on page 85 for details about the specific rules for system, network, or file resources.

- Depending on the resource type and rule set, the **Threshold** field allows you to set the threshold limits in terms of units, such as number of polling cycles or percentage.
- **Alert severity**, which is one of four alert levels: **Information**, **Warning**, **Error**, or **Severe**.
- The **Throw alert after failed poll(s)** and **Clear alert after successful poll(s)** fields, which allow you to monitor changing resource circumstances over any number of polling cycles rather than set an alert to occur immediately after a single failure of an active rule. These can also automatically clear the alert after a single successful cycle. Monitoring resource successes and failures over a number of polling intervals allows you to reduce noise that would cause erroneous alerts due to normal fluctuation in resource behavior.
- The **Throw additional alerts** field, which lets you choose when the alerts are thrown. The choices are either after a clear or after a specified number of failures.
- The **On alert perform action** and **On clear perform action** fields, which indicate whether the alert initiates an action. Possible actions include sending an e-mail to a designated individual, logging the alert in a log file, or a combination of several different actions into one compound action.

Using default rules

OpenEdge Management provides default rule values for system, network, and file resources. A rule indicates a performance standard that you expect and one about which you want notification when broken. Default values provided by OpenEdge Management suggest performance settings you can use when creating your resources.

Default rules help you to create your resources and associated values as quickly as possible. You can change default rule values to more specifically reflect your resource monitoring objectives at any time.

For more detailed information about the default rules and their default values, see the relevant sections in Monitoring System Resources on page 65, Monitoring Network Resources on page 73 and Monitoring File Resources on page 85.

Using resource-specific rules

Each resource type has specific rules you can set to ensure that the resource monitoring data provides you with meaningful details about the resource. For example, you can set a CPU resource monitor to monitor CPU-related performance characteristics, and you can set a disk resource monitor to monitor disk-related performance characteristics. Resource-specific rules reflect how trend data intervals have different meaning to the various resource types.

Sharing rules and rule sets

You can logically group multiple rule definitions for certain (complex) rule types, such as log file monitors, AppServers, NameServers, WebSpeed Transaction Servers, and databases. These rule combinations, called *rule sets*, are established and stored in the **OpenEdge Management Component Library**. A rule set comprises one or more rule definitions.

One of the advantages of using rule sets is that they are late-binding. A change to the rule set (such as adding, deleting, or modifying a rule definition) will propagate to all monitoring plans that use the rule set.

By sharing groups of rule definitions, you can manage a large number of resource monitors consistently and efficiently. For example, once a rule set exists, you can associate it with a log file monitor's monitoring plan. When the monitoring plan is active and the log file monitor is polled, OpenEdge Management can evaluate each rule in the set. See Monitoring File Resources on page 85 for details about how log file monitors use rule sets. See *OpenEdge Management: Database Management* for details about database resources and rule sets, and *OpenEdge Management: Servers, DataServers, Messengers, and Adapters* for details about AppServer, NameServer, and WebSpeed rule sets.

Choosing a rule's values

To review and change rule definitions or other default values for a sample disk resource:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the disk, whose rule definitions you want to review or modify.

   The **Disk: <disk name>** page appears.

3. Under **Monitoring Plans**, click either a schedule or its **Edit** button. (The **Edit** button is located to the right of the plan's trending information.)

   The **Monitoring Plan** page for the disk appears.

4. Review the rule definition as follows:

   - Accept or change the default value that displays in the **Alert if disk activity exceeds** field. This is the only rule unique to setting up a disk resource monitor. For example, you might want to perform disk monitoring to determine if your disk usage is greater than 75 percent busy at certain peak times during the work week. In this field, you would override the default value of 90.0 with `75.0`.
   - Accept the default alert severity, or choose another severity option from the drop-down list associated with the **Alert severity** field.
   - Identify alert behavior by accepting or changing values in the fields listed in the following table:

Field	Description
Throw alert after: failed poll(s)	Indicates how many times in a row the polling interval must fail before an alert is generated
Throw additional alerts	Lets you choose when the alerts are thrown—either after a clear or after a specified number of failures
On alert perform action	Triggers the action you specify to occur when an alert is generated
Clear alert after: successful poll(s)	Indicates the number of sequential successful polls that must occur before an alert is cleared
On clear perform action	Triggers the action you specify when an alert is cleared

Note: By default, the **On alert perform action** is set to **Default_Action**. The **On clear performance action** is set by default to **Default_Clear_Action**; selecting an action for this field is optional.

5. When you finish, click **Save**.

   The **Disk: <disk name>** page reappears. All values you entered for the plan and rule definitions are now in effect. Any future alerts generated for this resource are based on the new criteria you set.

Calculating meaningful thresholds with the Configuration Advisor

The Configuration Advisor calculates suggested threshold values for rules by analyzing trend data. When you apply a recommended rule threshold setting, the alerts triggered will provide a more meaningful indication of your resource's performance than if you were to arbitrarily set rule values.

The Configuration Advisor analyzes a rule's past performance for a specified period of time and, based on this data, calculates a baseline value. A *baseline value* is a threshold setting that is based on your system's past activity for the particular rule.

You must be an administrator user to use the Configuration Advisor.

Configuration Advisor baselining guidelines

In order to receive the most meaningful feedback, consider these guidelines when setting up baselines with the Configuration Advisor:

- Make sure you set up OpenEdge Management to store trend data. If you did not set this option initially, you can set it by choosing **Options** > **Trend database** in the management console.
- It is important that you collect data and trend data on a 1:1 basis; in other words, it is recommended that you trend on every poll. Whether you set the polling interval to five minutes or ten minutes is not significant for baselining. However, it is important that you trend the performance data on every poll, regardless of the poll interval.
- Carefully consider the timeframe for which you request baselining recommendations. For example, if your resource typically sees little or no activity between the hours of 5 PM and 9 AM, you would not want to get a baseline for an entire 24-hour period. If you did, you would be mixing times of high activity with times of low or nonexistent activity, which would skew the baselining results. Instead, it would make sense to choose a timeframe of 9 AM to 5 PM.
- Typically, one week of data collection, using the OpenEdge Management-provided defaults, provides enough information for the Configuration Advisor to provide a meaningful threshold. This means that the calculating done by the Configuration Advisor will be more time-consuming than if you had chosen a shorter time frame.
- Each rule you want to analyze must have at least 32 data samples stored in the OpenEdge Management Trend Database. This sampling provides sufficient data from which the Configuration Advisor can determine a baseline value and subsequently perform a successful analysis of each rule's data.
- In order for the Configuration Advisor to provide a recommendation, a polled rule must currently be associated with a defined monitoring plan for that resource.

Using the Configuration Advisor with a disk, CPU, or file system resource

You can use the Configuration Advisor with the following resources: databases, disks, CPU, file systems, AppServers, NameServers, and WebSpeed Transaction Servers. For details about using the Configuration Advisor with the database or AppServer, NameServer, or WebSpeed Transaction Server, see *OpenEdge Management: Database Management* or *OpenEdge Management: Servers, DataServers, Messengers, and Adapters*.

The way in which you use the Configuration Advisor with a disk, CPU, or file system resource is similar because each resource has only one rule. For each rule, the Configuration Advisor calculates recommended values based on a particular formula, as listed in the following table.

Table 5: CPU, disk, and file system rules and formulas

Resource	Rule	Formula
CPU	CPU High	CPU_Usr + CPU_System + CPU_Wait
Disk	Disk busy	Dev_Pct_Busy
File system	Filesystem used	FileSys_UsedPct

In the case of the CPU rule CPU High, the Configuration Advisor looks at the sum of CPU_Usr, CPU_System, and CPU_Wait when calculating recommended values; for the disk resource, the Configuration Advisor looks just at the Dev_Pct_Busy value.

The following procedure describes the process of creating meaningful threshold settings for a disk resource. However, the steps also pertain to the CPU and file system resources.

To use the Configuration Advisor with a disk resource:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the disk whose rule definitions you want to review or modify.

   The **Disk: <disk name>** page appears.

3. Click **Configuration Advisor**.

   The **Configuration Advisor** page for the disk appears.

4. In the **Start Date** and **End Date** fields, define a date range that OpenEdge Management will use to collect data from the OpenEdge Management Trend Database. (The default data range covers one week.)

5. Under **Choose time period to analyze**, identify the time frame that defines a period in which the rules are generally active, or being used. This time frame is the period against which you want OpenEdge Management to calculate your baseline activity. (The default time period is Monday through Friday, 9:00 AM to 5 PM.)

   In the **Select rules (for analysis)** section, the **Disk Busy** rule, which is the only rule available for analysis for a disk resource, is preselected.

6. Click **Submit**. The **Configuration Advisor** begins to calculate the threshold settings, and one of the following occurs:

   - If sufficient data exists for analysis, the Configuration Advisor displays the calculated result.

- You can now review the Configuration Advisor's recommendation and choose whether or not to update your threshold values. See Reviewing the recommended threshold settings on page 52 for details.
- If insufficient data exists for analysis, a message appears (under **Recommend Values**) informing you that there is insufficient data for analysis. For example, you might have fewer than 32 data samples stored in the OpenEdge Management Trend Database. In this case, consider expanding the time between the start date and the end date to capture more data. Then rerun the Configuration Advisor.

Reviewing the recommended threshold settings

The recommendations that the Configuration Advisor presents include a range of possible values from which to select. The asterisked value that appears in the **Recommend Values** field indicates the primary recommended threshold setting, based on the data analysis. Click the drop-down list to see the additional recommended values.

If there were more than one rule being evaluated by the Configuration Advisor for this resource, each rule would appear in its own row, with its own recommended values, update, current threshold, and **Detail** button.

Each recommended value is expressed as a pair of numbers. The first number specifies the recommended threshold setting. The second number, displayed in brackets, identifies the number of times this threshold value, if used with the collected data, would have broken the rule and triggered an alert. As you review the recommended threshold settings, keep in mind the rule behavior and alert notification frequency you want to establish for a resource.

Reviewing threshold calculation details

Each row has an associated **Detail** button that displays threshold calculation information about the rule's analysis.

The following details are provided:

- **Number of useable samples** — The number of data samples extracted from the OpenEdge Management Trend Database
- **Min Value** — The minimum value derived from the data set
- **Max Value** — The maximum value derived from the data set
- **Std Deviation** — The root mean squared deviation
- **Average** — The average value derived from the data set

When you finish looking at the details, click **OK** to close the window.

Comparing and selecting threshold settings

By default, the Configuration Advisor assumes that you are going to select and submit one of the recommended threshold settings.The **Update** field for each rule is selected by default. However, you have options concerning the selection process. As you compare the existing and recommended values, you can elect to change none, some, or all values for a rule and each individual monitoring plan.

To compare and select threshold settings:

1. For the one disk rule, note the value that displays in the **Current Threshold** field under a specific schedule.
2. Click the **Recommend Values** drop-down arrow to display all recommended values for the associated rule.
3. Compare the possible recommended values that display with the value in the **Current Threshold** field. As you determine the best threshold rule setting, keep your goals for this rule in mind. Also, consider any additional selection criteria when you compare values.

Additional selection criteria

Reasons for selecting one value and not another might include:

- How often you want alerts generated
- Your knowledge of your system's operational needs and goals
- Any other factors unique to your resource's performance that you want to consider when making your selection

Submitting your threshold setting selections

When you click **Update Selected Rules**, OpenEdge Management applies your selections. There is no undo option associated with this submission. To reset any values back to a previously defined setting, you must access the resource's monitoring plan, display the individual rule, and set (manually) the displayed value.

Determining the effectiveness of your selections

The most effective way to determine if your threshold adjustment is serving your needs is to review your alert notifications. Strive for a threshold setting that is consistent with your resource and business needs. If you receive alerts too frequently or too infrequently to suit your operational needs, you might want to refine your threshold settings.

Understanding alerts

An *alert* is an indication that a resource is not performing in compliance with the rules you established in the resource's monitoring plan. When OpenEdge Management triggers an alert for a resource, the following events occur:

- The actions you assigned to the resource monitor are initiated. For example, you might define an action to send an e-mail or page an operator when an alert occurs.
- The alert is added to OpenEdge Management's open alert table. To see the current open alerts, click **Alerts** > **Go to Alerts** from the menu bar.

The level of severity you set for an individual resource in its rules definition determines the type of alert that OpenEdge Management generates: **Information**, **Error**, **Warning**, or **Severe**. For more information, see *OpenEdge Management: Alerts Guide and Reference*.

All alerts (with the exception of information alerts) cause the resource status to change to **Failed** and the icon becomes red.

Note: You can also set up actions and alerts to respond to job status details and job execution time expectations. For more information about job execution, actions, and alerts, see Creating Jobs and Job Templates.

Determining OpenEdge Management's action in response to an alert

An *action* is an automatic, user-defined process that occurs in response to a rule violation on a monitored resource. You create actions and associate them with rules to ensure that certain activities occur when the rule triggers.

You create and store actions in the **OpenEdge Management Component Library**. You use these actions to send e-mails, record the occurrence of an alert in a log file, send an SNMP trap, or perform a combination of these activities.

When you select **Library** > **View** > **Actions** from the menu bar, all actions—with the exception of jobs—appear in the **Library.Actions** page.

Using default actions

OpenEdge Management provides five default action definitions:

- **Default_Action** — A compound action that contains a single action
- **Default_Clear_Action** — Sends mail to the default email recipient
- **Default_Mail_Action** — Sends email to the user you specified when you initially configured OpenEdge Management
- **Default_Pager_Action** — Sends a message to small display devices, such as a pager
- **Default_Task_Action** — Send a message to the default email recipient when a job finishes

The **Default_Action** is set in the **On alert perform action** field when you initially configure a monitor. If you prefer, you can choose a different action for OpenEdge Management to perform when an alert is generated.

Each default serves as a reusable, initial action definition that you can use or modify just as you can any action that you create. These default actions are provided to get OpenEdge Management resources created and running quickly. You can edit these default actions, but you cannot delete or rename them.

Creating an e-mail action

You can associate an e-mail action with a rule violation. The e-mail action defines the name, description, e-mail addresses of the sender and recipient, and a user-defined subject and message that are initiated in response to an alert.

To create an e-mail action:

1. In the OpenEdge Management console menu, choose one:
   - Select **Library** > **New** > **E-mail Action**.

- Select **Library** > **Go to Library** > **Create E-mail Action**.

The **Create new E-mail Action** page appears.

2. In the **Name** field, enter the name of your action.

   You must not use spaces or special characters such as an asterisk (*), ampersand (&), or period in this field.

3. In the **Description** field, enter a description of the action.
4. Enter the e-mail address of the intended alert recipient in the **Send to** field. (This is a required field.) Consider the following points when entering this value:

   - You should explicitly enter e-mail addresses (for example, `smith@`*`company.com`*). However, depending on your mail server implementation, OpenEdge Management might be able to determine the mail server domain name and automatically provide the `@`*`company.com`* portion of the address.
   - You can also enter names using quotation marks. The quotation marks will display when the recipient receives the e-mail. Also, you can enter a comma-delimited list of recipients. For example: *"Some User"* <*someuser@somedomain.com*>, *"Some Other User"* <*Someotheruser@somedomain.com*>.

5. Enter the e-mail address of the sender in the **Send from** field. OpenEdge Management supports either of the following formats to enter a value for the **Send from** field:

   - *someuser@somedomain.com*

     When this format is received in the recipient's inbox and appears in the FROM: header text of the message, the entire entry name and e-mail address, as entered in this format, appear.

   - *"Some User"* <*someuser@somedomain.com*>

     When the recipient receives an e-mail that is sent using this format, the recipient's inbox displays only *"Some User"* with quotation marks, as the sender of the e-mail. The address does not appear. Similarly, when you open the message, the **FROM:** header appears as: **From: "Some User."** Only when you display a received e-mail will you see both the user name and the domain name as entered using this format.

6. Enter a description of the e-mail's topic in the **Subject** field.

   In the **Subject** field, the phrase **OpenEdge Management Alert** automatically appears, followed by the e-mail variable **%AdminServerNAME%** that identifies the specific AdminServer to which the alert applies. However, you can override any of this information. You can also use any of the e-mail variables described in the following table in the **Subject** field to clarify the content of the e-mail.

Table 6: Variables for an e-mail action

Variable	Description
%ALERTID%	The ID of the current alert
%ALERTNAME%	The name of the current alert
%ALERTSEVERITY%	The alert severity (`Severe`, `Error`, `Warning`, or `Informational`)
%HOST%	The host on which OpenEdge Management is running

Variable	Description
%CONTAINERNAME%	The name of the container that is used to transport the resource
%RESOURCENAME%	The name of the resource that has caused the alert to generate
%ALERTLASTOCCURRENCE%	The date and time of the previous occurrence of the alert
%ALERTREASON%	A text description of the alert
%ALERTCOUNT%	The number of times that the alert has occurred
%ALERTTRIGVALUE%	The value that has triggered the alert
%HOSTLINK%	The HTTP path of the host name
%RESOURCELINK%	The full path of the resource that has generated the alert causing the e-mail to be sent
%ALERTLINK%	The link to the description of the alert that has caused the e-mail to be sent

Note: The percent symbol (%) is used before and after each variable. If you need to include the percent symbol in the text of a message, you must enter two percent symbols (%%) in the message body.

7. In the **Send As** field, select either **Text** (the default option) or **HTML** as the format option for your message. If you select **Text**, no additional work to your message is required. However, if you select **HTML**, you must use proper HTML-formatting techniques to ensure that your message will be readable by the recipient.
8. Choose the priority level of the e-mail from the **EMail Priority** list: **High**, **Normal**, or **Low**.
9. Use the default message variables that appear in the **Message** area, or edit these variables with the values that will appear in the specific e-mail action that you create. Refer to the above table for a listing of the variables you can define for an e-mail message.
10. Click **Save** to save or **Cancel** to discard the action. When you save the e-mail action, the properties of the new action appear.

    You can now edit, copy, delete, or test the e-mail action.

Creating a log action

A log action allows you to record a predefined message in a log file when an alert is generated.

To create a log action:

1. In the OpenEdge Management console menu, choose one:
   - Select **Library** > **New** > **Log Action**.

- Select **Library** > **Go to Library** > **Create Log Action**.

The **Create new Log Action** page appears.

2. Enter the name in the **Name** field. You must not use spaces or special characters such as an asterisk (*), ampersand (&), or period (.) in the Name field.
3. Enter the description in the **Description** field.
4. Enter the name of the log file in the **File** field.
5. Click **Save**. The **Log Action: <log action name>** page for the newly created log action appears.

   You can now edit, copy, delete, or test the log action.

Creating a compound action

When you create a compound action, you define one single action based on a number of individual actions that you frequently perform as a series of notification and operational steps in response to an alert. OpenEdge Management provides a default compound action you can use called Default_Action.

To create a compound action:

1. In the OpenEdge Management console menu, choose one:

   - Select **Library** > **New** > **Compound Action**.
   - Select **Library** > **Go to Library** > **Create Compound Action**.

   The **Create new Compound Action** page appears.

2. Enter the name in the **Name** field. (This is a required field.) You must not use spaces or special characters such as an asterisk (*), ampersand (&), or period (.) in the Name field.
3. Enter the description in the **Description** field.
4. Under **Select actions**, select an action in the **Available** list and click the right arrow. The action you selected moves to the **Selected** list and is added to the compound action. You can also select an action in the **Selected** list and click the left arrow. The action you selected is now removed from the compound action.

   Repeat these additions and/or removals until you have only the actions you want in the compound action in the **Selected** list area.

5. Click **Save**. A summary of the actions that compose the compound action you created appears.

   You can now edit, copy, delete, or test the compound action.

Creating a Simple Network Management Protocol (SNMP) trap action

You use the SNMP trap action if you want OpenEdge Management resource-related event notifications to be sent to your SNMP management console. You define properties to send event notifications, referred to as *traps*, to SNMP management consoles such as BMC's PATROL Enterprise Manager Connect SNMP and Tivoli Enterprise Console.

The SNMP trap action generates an asynchronous message to an SNMP management console.

Note the following information about SNMP trap actions:

- To define SNMP trap actions, you must have previously installed the OpenEdge Management SNMP Adapter product. You must also define the OpenEdge Management SNMP Adapter option on the **OpenEdge Management Configuration** page.
- Since OpenEdge supports SNMP versions 2c and 3, the SNMP trap action properties are set based on the SNMP version.
- All OpenEdge Management-specific SNMP variables and tables are located in the OpenEdge Management Information Base (MIB). The MIB file, `PSC-FM-MIB.txt`, is located in the config directory of your OpenEdge Management installation folder. By default, this location is `Progress\oemgmt\config`.

  Review the contents of this file and focus specifically on the `fmalertTable`, which contains information specific to each alert. The alert ID information in this file allows you to understand the contents of this table.

  The following table contains two examples that show how you can obtain alert details from the `fmalertTable`.

Table 7: Alert details obtained from the fmalertTable

To obtain an alert reason associated with Alert ID . . .	Issue an SNMP GET request using this Object Identifier (OID) . . .
3	1.3.6.1.4.1.1730.1.2.1.1.6.1.6.3
4	1.3.6.1.4.1.1730.1.2.1.1.6.1.6.4

For more information about the SNMP Adapter and its configuration, refer to the *Changing the OpenEdge Management SNMP Adapter settings* section in *OpenEdge Management and OpenEdge Explorer: Getting Started*. For information about initiating SNMP GET requests from an SNMP agent, see the SNMP documentation for your particular SNMP management console product.

To create an SNMP Trap Action:

1. In the OpenEdge Management console menu, choose one:
   - Select **Library** > **New** > **SNMP Trap Action**.
   - Select **Library** > **Go to Library** > **Create SNMP Trap Action**.

   The **Create new SNMP Trap Action** page appears.
2. Enter the name in the **Name** field. You must not use spaces or special characters such as an asterisk (*), ampersand (&), or period (.) in the **Name** field.
3. Enter a description in the **Description** field.
4. Enter the name of the host machine on which the SNMP management agent resides in the **Host** field. Traps will be sent to this host machine. (If you set a default value for the **Host** field on the **OpenEdge Management Configuration** page, that value appears in this field.)
5. Enter the port number of the host machine on which the SNMP management agent resides in the **Port** field. Traps will be sent to this port number on the host machine. (If you set a default value for the Port field on the **OpenEdge Management Configuration** page, that value appears in this field.)
6. Select the SNMP adapter settings to be used for the trap by setting the **Settings to use** option.
   - To use the SNMP adapter settings, select **Use preference settings**.
   - To manually set the SNMP adapter settings, select **Use trap settings**.

- If the SNMP version is set as version 2c (in **Options** > **SNMP Adapter**), the read community value of the SNMP adapter is displayed in the **Read Community** field. This is the community that will have read access to the trap information. You can customize the value by entering `public`, `private`, or any other community that you have set up for access.
- If the SNMP version is set as version 3, the security configuration of the SNMP adapter is displayed. You can customize these settings for the SNMP trap. For more information on the SNMP Adapter settings for version 3, refer to the *Changing the OpenEdge Management SNMP Adapter settings* section in *OpenEdge Management and OpenEdge Explorer: Getting Started*.

Note: While creating a trap action for SNMP version 3, if you do not set the authentication and privacy passwords, the SNMP adapter passwords are used as default. If you do enter new passwords, these passwords take precedence over the SNMP adapter passwords and are henceforth used as default.

7. Click **Save**.

   Once you have set up SNMP trap actions, you can associate them with resources just as you would any OpenEdge Management action. When the resource generates an alert, the SNMP trap action is thrown.

   You can review the contents of the trap (known as variable bindings) as they are defined in the MIB. You can query these variables any time the OpenEdge Management SNMP Adapter agent is running.

Updating and deleting actions

Once a resource monitor exists, you can update actions assigned to it at any time. However, you cannot delete an action that is in use by a resource. If you try to delete an action that is currently in use, OpenEdge Management displays a message box to notify you that the action is being used. OpenEdge Management also provides a list of resources that are using the action.

Editing or copying a resource monitor

You can edit and copy the properties, monitoring plans, or rule definitions for most resources. There are some limits to the types of editing you can perform for the OpenEdge Management-provided default resources such as CPU, Memory, and disk. For example, you cannot rename, copy, or delete any default resource monitors. However, you can change monitoring plan and rule definitions associated with each of them.

Note that if two or more people simultaneously edit a resource, the last changes submitted overwrite all previous changes.

To edit or copy a resource:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the system, network, or file resource.

   The resource's summary data, which includes properties, monitoring plans, and rule definitions, appears on the detail page.

3. Choose one of the following:

   - Click **Edit**. The property summary data reappears, allowing you to change editable fields.

- Click **Copy**. You can change the properties, monitoring plans, or rule definitions as necessary. Note that when you copy a resource, the copy displays the same name as the original resource. You must change this name before saving the copy.

4. Click **Save** when you finish.

Deleting a resource monitor

Once you delete a resource, you cannot undelete it. To temporarily disengage a monitor, you might consider disabling an **Enabled** resource option rather than deleting it.

You cannot delete the OpenEdge Management-provided CPU default resource or Memory default resource. Similarly, you cannot delete or rename Default_Action, Default_Clear_Action, or Default_Mail_Action. For details about deleting database resource monitors, see *OpenEdge Management: Database Management*. For details about deleting WebSpeed broker, AppServer broker, or NameServer resource monitors, see *OpenEdge Management: Servers, DataServers, Messengers, and Adapters*.

To delete other resource monitors:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the system, network, or file resource.

   The resource's summary data, which includes properties, monitoring plans, and rule definitions, appears on the detail page.

3. Click **Delete** and then click **Yes** in the confirmation dialog box to confirm deletion.

   A message appears informing you that the resource has been deleted.

Working with default values

OpenEdge Management provides different opportunities for you to set up and use default values that allow you to:

- Quickly and easily update default information at global and general resource levels. You can override these values on individual instances of each resource, as necessary.
- Minimize duplication of error in creating resource monitors.

Using default values can also help you to become operational quickly with OpenEdge Management and standardize your resource performance criteria. You can establish default values:

- At the OpenEdge Management global settings level
- From specific default pages associated with the different OpenEdge Management resource monitoring types

Reviewing and changing global settings

During the OpenEdge Management installation process, you have the option of choosing global settings for various resource options. At any time after the initial installation, you can review and update these settings. Specifically, you can enable and disable global settings for the following OpenEdge Management elements:

- Polling activities for all resources
- Trending activities data for all resources
- Alert generation

To update the global settings:

1. Click **Options** in the OpenEdge Management console.
2. In the list frame, select **Resource Monitoring**.

   The **OpenEdge Management Resource Monitoring Configuration** page appears.
3. To enable or disable global resource monitoring settings, select or clear the appropriate options.

   The resource monitoring settings are independent of each other and are useful for diagnostic purposes because they allow you to disable specific OpenEdge Management functionality. For example, if you know that your OpenEdge Management Trend Database is going to be taken down for maintenance, you can elect to turn off trending, but leave the rest of OpenEdge Management running.

   If you disable polling, trending, or alerts, an associated icon appears in the upper-right corner of the menu bar to indicate that the option is currently disabled. When polling is disabled at this configuration level, a message indicating that the polling has been disabled appears on each individual resource monitor-related page.

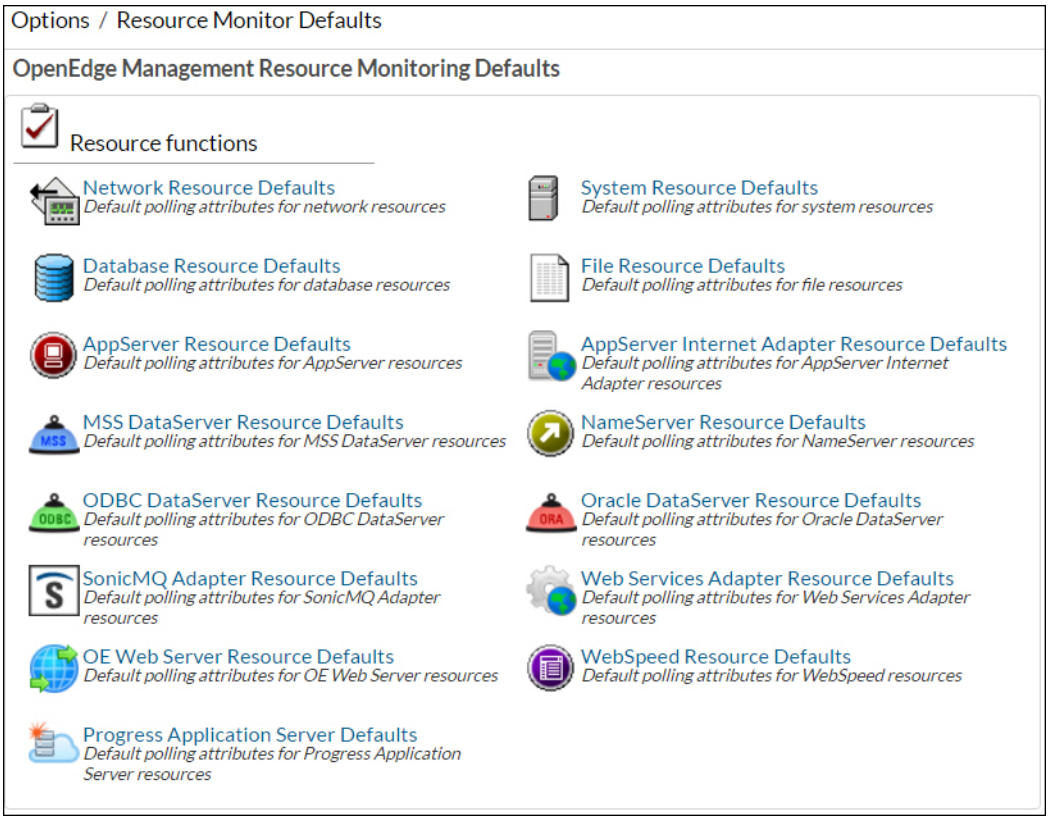
   For further information about updating global resource monitoring settings, see Working with default values on page 60. For details about setting actions for OpenEdge Management internal alerts, see *OpenEdge Management: Alerts Guide and Reference*.

Reviewing and changing resource type defaults

Each of the general OpenEdge Management resource types has predefined default values. You can change these defaults as necessary.

To see current default values:

1. Click **Options** in OpenEdge Management console, and select **Resource Monitor Defaults**. The **OpenEdge Management Resource Monitoring Defaults** page appears.



2. To view the current defaults for a resource type, click the appropriate link.

   For more information about system, network, or file resource defaults, see Monitoring System Resources on page 65, Monitoring Network Resources on page 73, or Monitoring File Resources on page 85. For information about database resource default values, see *OpenEdge Management: Alerts Guide and Reference*. For details about OpenEdge server and adapter component default values, see *OpenEdge Management: Servers, DataServers, Messengers, and Adapters*.

Using job and report templates

OpenEdge Management provides a template feature to help you create and maintain jobs and report data. Some benefits of these templates are:

- An OpenEdge Management job template helps you standardize and quickly create numerous individual job instances using predefined, common values.

  For details about working with job templates, see Creating Jobs and Job Templates. For details about the OpenEdge Management-supplied database maintenance job templates, see *OpenEdge Management: Database Management*.

- An OpenEdge Management report template helps you to standardize and quickly create numerous individual reports using predefined, common values. For information about working with report templates, see *OpenEdge Management: Reporting*.

Using the OpenEdge Management Library

The OpenEdge Management Library contains resource monitoring components that you define, share, and, in many instances, transfer among your company's various work sites. Two of the primary benefits of employing library components are:

- Minimizing duplication of effort when creating numerous and similar resources
- Standardizing implementation of your company's business policies through the definition and distribution of these components.

You can define schedules, actions, search criteria, and rule sets at the library level. The import and export activities that are also available in the library allow you to distribute all library components (except schedules) that you create on one machine to other machines located throughout the company.

To access the OpenEdge Management Library, click **Library** > **Go to Library** from the menu bar. The **OpenEdge Management Component Library** appears.

The components you build, maintain, and access through the **OpenEdge Management Component Library** page are independent of each other. This independence allows you the flexibility to:

- Combine and reuse these components
- Minimize redundant resource setup work for numerous items that are very similar
- Help standardize work flow processes across an enterprise

The following table briefly highlights the remaining library function options, according to the type of resource with which they are associated.

Table 8: Additional OpenEdge Management Library features

To help . . .	Use these library components . . .
Create log file monitors	**Search Criteria** and **Log File Rule Set**. These components allow you to define rules-related data that is to be shared among log file monitors. For example: • **Create Search Criteria** — Defines rule properties from which you can select and associate with other rule properties to create individual rules to run within your log file monitors • **Create Log File Rule Set** — Defines a set of log file monitor rules that you can associate with one or more log file monitor resources For details, see Monitoring File Resources on page 85.
Create database monitors	The database rule set, which allows you to define a set of database rules that you can use with one or more database resources. For details, see *OpenEdge Management: Database Management*.
Create OpenEdge monitors	The NameServer, AppServer, or WebSpeed rule set, which allows you to define a set of rules that you can use with each type of resource. For details, see *OpenEdge Management: Servers, DataServers, Messengers, and Adapters*.

To help . . .	Use these library components . . .
Identify existing OpenEdge Management components to export to other computers	**Export Components**, which allows you to identify various components that you create and share by importing from and/or exporting to another machine. For details, see Exporting and Importing in OpenEdge Management on page 133.
Identify existing OpenEdge Management components to import to other computers	**Import Components**, which allows you to identify various components that you create and share by importing from and/or exporting to another machine. For details, see Exporting and Importing in OpenEdge Management on page 133.

4

Monitoring System Resources

For details, see the following topics:

- System resource monitoring overview
- Reviewing system resource default values
- Reviewing system resource monitoring plans

System resource monitoring overview

You can monitor the following system resources in OpenEdge Management: CPU, disk, file system, and memory. OpenEdge Management automatically creates a CPU resource monitor and a memory resource monitor and assigns default values. You can accept the default values, or you can change them. However, you cannot delete or rename either the CPU or the memory resource.

You can choose to monitor all file systems and disks on your machine. If you make this choice during initial configuration, OpenEdge Management automatically creates a default resource monitor for each disk and file system. For details about selecting this configuration option, see the section about initial configuration settings in *OpenEdge Management and OpenEdge Explorer: Getting Started*.

Reviewing system resource default values

OpenEdge Management provides various default values for system, network, file, database, and OpenEdge server resources. You see these default values when you create a resource monitor and define its monitoring plan.

If the default values are acceptable to you, you can use them with the resource monitors you create. If, however, you want to change the default values, you can do so as follows:

- For each resource monitor you create, adjust the values on a per-resource basis so that you get the feedback data you want for that resource. Each new resource monitor is still initially created with the OpenEdge Management default.
- Change the actual defaults provided by OpenEdge Management so that each new resource monitor you create already has the most useful default values.

Changing default values does not affect any resource monitor already created.

To display default values for system resources:

1. Select **Options** > **Resource Monitor Defaults** in the OpenEdge Management console.

   The **OpenEdge Management Resource Monitoring Defaults** page appears.

2. Click **System Resource Defaults**.

   The **System Resource Defaults** page appears.

3. Review the following defined default values:

   - **Polling Interval** — The frequency with which this resource's rules are checked by OpenEdge Management—for example, every 300 seconds
   - **Alert if CPU activity/disk activity/file system usage/system memory usage/virtual memory usage exceeds** — The threshold percentage of activity or usage that causes OpenEdge Management to generate an alert
   - **Trend** — Whether you want performance data about this resource stored in the OpenEdge Management Trend Database
   - **Trend every** — How often you want the performance data stored in the OpenEdge Management Trend Database

4. Change the default values, as necessary, and click **Submit**.

   The new defaults will be in place when you create a new disk or file system resource. If the new resource the default values are still not what you want for a particular resource, you can change the values of the individual resource at that time.

   Note that you can revert back to the original OpenEdge Management-supplied default values at any time by selecting **Restore Defaults**.

Reviewing system resource monitoring plans

The CPU resource and the memory resource are automatically created when OpenEdge Management initially starts up. However, you can change the default monitoring plan and rules associated with each of these system resources.

Reviewing the CPU resource monitoring plan

You do not have to set up monitoring plans and rules for your CPU resource. OpenEdge Management automatically provides a default resource named CPU when you initially start up OpenEdge Management. You cannot delete this OpenEdge Management-provided default resource, nor can you change its name. However, you can change the values in the CPU resource's monitoring plan definition and rule definition.

To review the default CPU resource monitoring plan:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the default CPU resource you whose monitoring plan you want to review.

   The CPU details page appears.

3. To review the values set in the CPU's monitoring plan definition and rule definition, click the monitoring plan in the **Name** column. The monitoring plan appears.

   The default monitoring plan definition consists of the following values:

   - **Schedule** — Identifies the system-defined 24/7 default schedule used when the plan is active. The default is **Default_Schedule**.
   - **Polling Interval** — Identifies the polling cycle, which is the frequency at which the resource's rules are checked. The default is **5 minutes**.
   - **Alerts Enabled** — Indicates whether alerts are active and will be generated when the plan is active. The default is **true**.
   - **Trend Performance Data** — Indicates whether you want to store performance data in the OpenEdge Management Trend Database. The default is **true**.
   - **Trend every** — Indicates how often performance data is trended. The default is **1 poll(s)**.

   The default monitoring plan's rule definition consists of the following values:

   - **Alert if CPU usage exceeds** — Identifies when OpenEdge Management issues an alert on CPU usage. The default is 80.0%.

     OpenEdge Management monitors the percentage of CPU usage according to the value you enter in the **Alert if CPU usage exceeds** field within the rule definition. OpenEdge Management polls the CPU according to the **Polling Interval** value. This value reflects the percentage of time that a processor is busy executing a nonidle thread. You can view this percentage as the time spent doing useful work. You can also extrapolate the idle time based on this statistic.

     This statistic is presented as an average percentage based on snapshots taken at the beginning and end times of the defined sample period. Expressing the value as an average reduces the possibility of false alerts for random spikes in usage. For example, if your CPU monitor has a polling interval of 15 minutes, the average CPU time is calculated over this period of time.

   - **Alert Severity** — Identifies the severity level at which OpenEdge Management triggers an alert. The default is **Error**.
   - **Throw alert after** — Identifies at which point OpenEdge Management triggers an alert. The default is **2 failed poll(s)**.
   - **Throw additional alerts** — Indicates if OpenEdge Management will throw additional alerts. Possible values are **true** or **false**.

- **On alert perform action** — Identifies the action that OpenEdge Management takes when an alert is generated. The default is **Default Action**.
- **Clear alert after** — Identifies the number of successful polls after which OpenEdge Management clears the alert. The default is **0 successful poll(s)**.
- **On clear perform action** — Identifies a user-defined action or a compound action. For example, OpenEdge Management might run a job that was defined as an action. The default is **None**.

Trending performance data for CPU resources

Once you enable trending by selecting the **Trend Performance Data** field, you can also override the default value of **1** in the **Trend Performance Data every poll(s)** field. Any value that you enter in this field that is greater than the default value of 1 causes a CPU resource monitor to trend an average of the data that it gathers between trends. This average is calculated as the sum of the values divided by the number of polls.

Editing the CPU resource defaults

If you want to change any of the default CPU resource monitor values, click **Edit** from the **Monitoring Plan** page. The **Edit Default_Schedule Monitoring Plan** page appears.

Make the modifications you want, and click **Save**. The CPU details page appears and reflects the changes.

Reviewing the memory resource monitoring plan

You do not have to set up monitoring plans and rules for your memory resource. OpenEdge Management automatically provides a default resource named Memory when you initially start up OpenEdge Management. You cannot delete this OpenEdge Management-provided default resource, nor can you change its name. However, you can change the values in the memory resource's monitoring plans definition and rules definition.

To review the default memory resource monitoring plan:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the default Memory resource whose monitoring plan you want to review.

   The monitoring plan and a tabular and graphical representation of the memory status appear.

3. To review the values set in the memory resource's monitoring plan definition and rule definition, click the monitoring plan in the **Name** column. The monitoring plan for the memory resource appears.

   The default monitoring plan consists of the following monitoring plan definition values:

   - **Schedule** — Identifies the system-defined 24/7 default schedule used when the plan is active. The default is **Default_Schedule**.
   - **Polling Interval** — Identifies the polling cycle, which is the frequency at which the resource's rules are checked. The default is **5 minutes**.
   - **Alerts Enabled** — Indicates whether alerts are active and will be generated when the plan is active. The default is **true**.
   - **Trend Performance Data** — Indicates whether you want to store performance data in the OpenEdge Management Trend Database. The default is **true**.
   - **Trend every** — Indicates how often performance data is trended. The default is **1 poll(s)**.

The default monitoring plan consists of the following rule definition values:

- The **Alert if system memory usage exceeds** field indicates the threshold percent value that generates an alert. This alert indicates that the percentage of physical memory in use by the system has been exceeded.
- The **OR if virtual memory usage exceeds** field indicates the threshold value that generates an alert. This alert indicates that the percentage of virtual memory in use by the system has been exceeded.

  These memory-related statistics are displayed as single snapshots when OpenEdge Management polls the system to determine the amount of free virtual memory. For example, OpenEdge Management displays the amount of free virtual memory at a single instant. Therefore, these statistics are more likely to reflect random spikes in performance that might not accurately reflect overall performance. You should consider configuring these monitors so that an alert is not generated the first time the threshold is crossed, but only after it has been crossed enough times to validate concern about performance.
- **Alert Severity** — Identifies the severity level at which OpenEdge Management triggers an alert. The default is **Error**.
- **Throw alert after** — Identifies at which point OpenEdge Management triggers an alert. The default is **2 failed poll(s)**.
- **Throw additional alerts** — Indicates if OpenEdge Management will throw additional alerts. Possible values are **true** or **false**.
- **On alert perform action** — Identifies the action that OpenEdge Management takes when an alert is generated. The default is **Default Action**.
- **Clear alert after** — Identifies the number of successful polls after which OpenEdge Management clears the alert. The default is **0 successful poll(s)**.
- **On clear perform action** — Identifies a user-defined action or a compound action. For example, OpenEdge Management might run a job that was defined as an action. The default is **None**.

Trending performance data for memory resources

Once you enable trending by selecting the **Trend Performance Data** field, you can also override the default value of **1** in the **Trend Performance Data every poll(s)** field. Any value that you enter in this field that is greater than the default value **1** causes a memory resource monitor to trend data that it gathered from the last poll, with the exception of page in and page out fields. Page in and page out data can only be queried through the OpenEdge Management Trend Database.

Note: The page in and page out values represent the delta values from the previous trend.

Editing the Memory resource defaults

If you want to change any of the default memory resource monitor values, click **Edit** from the **Monitoring Plan** page. The **Edit Monitoring Plan**page appears.

Make the modifications you want and click **Save**. The Memory details page opens and reflects the changes.

Reviewing a disk resource monitoring plan

OpenEdge Management provides a default disk resource monitor. The resource monitor name depends on the operating system and the drives that are found on the machine.

You can also add more disk resource monitors. Depending on the initial configuration options you selected, it is possible that all disks found on your system will have default monitors created for them. For more information about disk options available during the OpenEdge Management setup and installation process, see the appropriate section of *OpenEdge Management and OpenEdge Explorer: Getting Started*.

Reviewing disk monitoring data

OpenEdge Management provides a complete list of available system disk devices for which you can also set up and maintain individual disk resource monitors.

To access a list of available disk devices:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the disk resource.

   The **Disk: <disk name>** page appears.

3. From the **Disk: <disk name>** page you can:

   - Define monitors for each of the disks.

     If, when you were initially configuring OpenEdge Management, you chose the option for OpenEdge Management to define monitors for all disks, resource monitors will already be in place for each disk. For more information about this configuration setting, see the section about initial configuration in *OpenEdge Management and OpenEdge Explorer: Getting Started*.

   - Access detailed monitoring plans and rules for available disks you want to set up or whose values you want to edit.

Disk device statistics

OpenEdge Management monitors the value in the **Alert if disk activity exceeds** field. This value indicates the percentage of elapsed time that the selected disk drive is busy servicing read or write requests. You can set an alert to trigger when this percentage is exceeded.

Trending performance data for disk resources

Once you enable trending by selecting the **Trend Performance Data** field, you can also override the default value of 1 in the **Trend every poll(s)** field. Any value that you enter in this field that is greater than the default value of 1 causes a disk resource monitor to trend data that it gathered primarily from the last poll.

Reviewing the file system resource monitoring plan

OpenEdge Management provides a default file system resource monitor. The resource monitor name depends on the operating system and the file systems that are found on the machine.

You can also add more file system monitors. Depending on the initial configuration options you selected, it is possible that all file systems found on your machine will have default monitors created for them. For more information about file system options available during the OpenEdge Management setup and installation process, see the appropriate section of *OpenEdge Management and OpenEdge Explorer: Getting Started*.

Note: Keep in mind that some of the file systems that are available to be monitored might not, in fact, be suitable for monitoring. UNIX-related examples that are considered unsuitable for monitoring include pseudo file systems used by the Kernel such as /proc, file systems that are actually mount points for remote or unmounted file systems. File system resource monitors set up for these types of file systems would be ineffective.

File system statistics

By establishing values for statistics, you define the rules for a file system resource monitor. OpenEdge Management collects the following statistics related to the utilization of file systems:

- **File System Name** — The name of the file system as it is known to the operating system
- **File system usage** — The percentage of file system capacity used

The **File system usage exceeds** statistic is gathered primarily for trend analysis. You can review these statistics and trend details to predict, to some extent, when it is time to extend or redesign your file system. Provided nothing other than the database is located on the file system, this detail indicates the database growth over time. You can set an alert to trigger when the percentage of file system capacity is exceeded.

The file system information can also help with mapping device-busy metrics to a file system. You can see which files are involved with performance issues.

Reviewing file system monitoring data

OpenEdge Management provides a complete list of available file systems for which you can set up and maintain individual file system resource monitors. OpenEdge Management also provides detailed data for all available file systems, whether monitored or not.

To access a list of available file systems:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the File system resource.

   The File System details page appears.

3. Click **Edit** to find the details you want to view.

Trending performance data for file systems

Once you enable trending by selecting the **Trend Performance Data** field (of the monitoring plan definition), you can also override the default value of **1** in the **Trend Performance Data every poll(s)** field. Any value that you enter in this field that is greater than the default value **1** causes a file system resource monitor to trend data that is taken from the last poll. No calculations are performed on this data.

5

Monitoring Network Resources

For details, see the following topics:

- Network resource monitoring overview
- Reviewing network resource default values
- Discovering and monitoring TCP and UDP ports
- Creating a TCP or UDP resource monitor
- Creating a PING (ICMP) resource monitor
- Using the HTTP resource monitor to determine Web site status

Network resource monitoring overview

You can use OpenEdge Management to set up resource monitors for non-OpenEdge resources. To determine the availability of some non-OpenEdge network resources, you follow network discovery procedures that allow you to search for ports that are present on your network.

This chapter describes network discovery procedures and explains how to create specific port connections for the following network resources:

- **TCP port** — You can set up resource monitors for ports that use Transmission Control Protocol (TCP) to monitor machine availability and enable a data connection between two machines.
- **UDP port** — You can set up resource monitors for ports that use User Datagram Protocol (UDP) to determine the availability of these ports enabling a given machine to broadcast messages over a network.

- **PING (ICMP)** — Packet Internet Groper (PING) is a utility used to determine whether a specific IP address is accessible. It works by sending a packet to the specified address and waiting for a reply. ICMP supports packets that can contain error, control, and informational messages. OpenEdge Management allows you to set up resource monitors on a PING to monitor the availability of specific machines and IP addresses.
- **HTTP Monitor** — Hypertext Transfer Protocol (HTTP) is an OpenEdge Management-supported resource monitoring feature that allows you to monitor various aspects of a Universal Resource Locator (URL). Businesses that rely heavily on the Internet need to ensure that their company's Web site is reliable, accurate, and accessible. The HTTP Monitor supports an authentication feature. You cannot use the discovery process with the HTTP Monitor resource process.

Caution: Give serious consideration to the types of devices for which you are creating network resources. OpenEdge Management does not allow you to PING broadcast or multi-cast addresses. However, it is possible to PING other devices, such as hubs, switches, and routers. Polling devices such as these can lead to network problems.

Trending considerations for network resources

Trend information for network resources is limited to the response times received by the resource and status details generated up to and including that time frame. With regard to network resources, information trended to the OpenEdge Management Trend Database refers to the worst-case poll information received between trends.

Reviewing network resource default values

OpenEdge Management provides various default values for system, network, file, database, and OpenEdge server resources. You see these default values when you create a resource monitor and define its monitoring plan.

If the default values are acceptable to you, you can use them with the resource monitors you create. If, however, you want to change the default values, you can do so as follows:

- For each resource monitor you create, adjust the values on a per-resource basis, so that you get the feedback data you want on the resource. Each new resource monitor is still initially created with the OpenEdge Management defaults.
- Change the actual defaults provided by OpenEdge Management so that each new resource monitor you create already has the most useful default values.

Changing default values does not affect any resource monitor already created.

To display default values for network resources:

1. Select **Options** > **Resource Monitor Defaults** in the OpenEdge Management console.

   The **OpenEdge Management Resource Monitoring Defaults** page appears.

2. Click **Network Resource Defaults**. The **Network Resource Defaults** page appears.

   Review the default values defined for TCP ports, UDP ports, ICMP, and HTTP monitors. Values set at this resource level appear as default values when you create an individual resource's monitoring plan and rule definitions.

3. Change the default values, as necessary, and click **Submit**.

The new defaults will be in place when you create a new network resource. If, when you create the new resource, the default values are still not what you want, you can change the values of the individual resource at that time.

Note that you can revert back to the original OpenEdge Management-supplied default values at any time by selecting **Restore Defaults**.

Discovering and monitoring TCP and UDP ports

OpenEdge Management supports network discovery procedures, which allow you to review a list of well-known ports or discover network resources. Once you obtain discovery search results, you can then add resource monitors for the active ports or networked resources, populating the resource-monitoring plans with default values that you can later change.

You can also qualify the discovery procedures, seeing only IPv4 addresses within a specific range or IPv6 network addresses in the results. This latter option is useful because rather than initiating a search across a range of network addresses (which can be time-consuming if there are not a lot of IPv6 addresses), you can enter a specific host name or address in the field provided.

Reviewing the list of well-known ports

OpenEdge Management can provide you with a list of well-known ports.

To review the list of available ports:

1. Select **Resources** > **Discover** > **Network Resources** in the OpenEdge Management console menu.

   The **Discover Network Resources** page appears and displays all the resources managed by the console.

2. Click **Well-known Ports** in the **IPv4 address range** section. The **Well-known Ports** page appears.

   This page displays a list of predefined ports that are available to be discovered. Each port is identified by the port number, the port type, and an editable default description.

   You can delete any port from the list by selecting the port and clicking **Delete**.

   You can also add ports to this list. Type a description in the available field, and then click **Add**. (Note that descriptions cannot contain non-ASCII characters.) Repeat this procedure for each port you want to add.

   Besides adding and deleting ports, you can also perform the following actions related to port discovery:

   - Update a port description in the space provided, and click **Submit** to save the update.
   - Click **Discover Network Resources** to discover which machines are available.
   - Click **Resource Home Page** to return to the **OpenEdge Management Resources** page.

Naming conventions in port descriptions

If a port number is one of the registered well-known numbers, a short description of the port appears as the port's name. (Once you create a resource monitor for a port, OpenEdge Management displays this descriptive information on the resource monitoring page, under the page title.)

The following table identifies names associated with several well-known TCP and UDP port numbers.

Table 9: Well-known ports and their associated names

Well-known port	Description field acronym
TCP Port 21	FTP
TCP Port 23	Telnet
TCP Port 25	SMTP
TCP Port 80	HTTP
TCP Port 110	POP3
TCP Port 119	NNTP
TCP Port 143	IMAP
TCP Port 389	LDAP
TCP Port 547	DHCPv6
UDP Port 53	DNS
UDP 161	SNMP

If a port is not considered well known, OpenEdge Management displays a port name in the **Description** field that is based on a combination of the port number and the protocol type you specify. For example, if you specify the port type as `TCP` and the port number as `3510`, OpenEdge Management displays the port name as `tcp3510`.

Discovering network resources

You can use OpenEdge Management to determine if a particular machine is available.

To discover network resources:

1. Select **Resources** > **Discover** > **Network Resources** in the OpenEdge Management console menu.

   The **Discover Network Resources** page appears and displays all the resources managed by the console.

2. Do either of the following:

   - To search for IPv4 resources, enter a range of IPv4 Internet Protocol (IP) addresses in the **Starting Address** and **Ending Address** fields. This activity enables OpenEdge Management to discover machine and port information.

     To simplify your entries in the **Starting Address** and **Ending Address** fields, OpenEdge Management provides the first two fields of your local host in each of these entries. To enter a valid starting address, you must append these partially provided values with the remaining two fields.

     Both the **Starting Address** and the **Ending Address** must begin with the same values.

   - To search for IPv6 resources, click **Discover active resources** under **IPv6 Network Host**. The **Discover IPv6 Port Resources** page appears.

Type the host name or address.

3. Click **Discover**. The **Discovery Status** page appears.

Note: How quickly the page appears depends on the number of resources OpenEdge Management must PING and the time it takes to perform this activity.

The **Discovery Status** page displays and updates the status of the discovery process as resources and ports associated with the IP addresses are scanned. PING information obtained for active resources, and UDP and TCP ports that are discovered to be active, are reported and updated on this page approximately every three seconds.

Once the discovery process ends, the **Discovery Results** page appears.

The following details about discovered resources are returned and presented if a machine is available:

- TCP and/or UDP port data
- Machine name

The icon that displays next to a discovered resource indicates whether the resource is a TCP port, a UDP port, or an ICMP resource.

Adding resource monitors for discovered machines and ports

Once you have discovered machines and ports, you can set up resource monitors for them.

To set up the resource monitors:

1. From the **Discovery Results** page, select the individual check box associated with each resource you want to add, or click **Select All** to select all resources.
2. Click **Add Checked**. An informational message appears, stating that the network resources you selected have been successfully added to OpenEdge Management.
3. Click **OK** to close the informational message.

   Discovered network resources also follow a naming convention. These resource names are derived from the hostname associated with the discovered resource and a character string derived from the resource's type and port number. Each of these elements is separated by a colon, using the following format:

   ```
   resource typeresource port number:discovered resource
   ```

   For example, if an activity on port 80 on a machine with a host name of **DELL1** is discovered, the resource name that OpenEdge Management generates is **HTTP:DELL1**. In contrast, if an activity on port 1234 on a machine with a host name of **DELL2** is discovered, the resource name that OpenEdge Management generates is **tcp1234:DELL2**.

   When you add the discovered network resource to OpenEdge Management, the colon is replaced by an underscore; for example, **tcp23:nbaspauldixp** becomes **tcp23_nbaspauldixp**.

4. Select a newly added network resource to display its default monitoring plans and rules.

   You can now work with the monitor for the newly discovered resource just as you would any other monitor.

Accessing a previously discovered resource

You can later reaccess any discovered resource that you do not add to OpenEdge Management at the time of discovery. The results are retained from only one previous discovery process and are replaced once a subsequent session concludes.

To reaccess a previously discovered resource:

1. Select **Resources** > **Discover** > **Network Resources** in the OpenEdge Management console menu.

   The **Discover Network Resources** page appears and displays all the resources managed by the console.
2. Click **Previous Results**.

Creating a TCP or UDP resource monitor

OpenEdge Management supports setting up individual TCP and/or UDP monitors. The TCP monitor attempts to connect to a TCP port and reports status and response time. The TCP monitor can be useful in determining the status and response times for servers that listen on a TCP port. This includes servers such as mail, FTP, and Web servers.

Note: OpenEdge Management supports a default operator e-mail TCP monitor identified as a Simple Mail Transfer Protocol (SMTP) SMTP_Mail. You can set host, port, and default e-mail recipient values in the **Default alert recipient** section of the **OpenEdge Management Configuration** page. However, you can edit values as you would any user-defined TCP resource monitor.

The UDP monitor attempts to communicate with a UDP port and reports status and response time.

Note: Monitoring a NameServer using UDP resource monitors is not recommended. If you set up a UDP monitor to monitor a NameServer UDP listener port, the NameServer will detect the UDP messages. However, the NameServer will inspect the UDP messages and reject them because they do not follow the convention expected by the NameServer. Consequently, the NameServer will log error messages to indicate that invalid message requests are being received.

To define a TCP resource monitor:

1. Select **Resources** > **New** > **TCP** in the OpenEdge Management console menu.

   The **Edit TCP Monitor** page appears.
2. Enter values in the **Name** and **Description** fields.
3. In the **Host** field, enter either the name of the host or the IP address where the port is located.
4. In the **Port** field, enter the number of the port.
5. Ensure that the **Enabled** field is checked to begin monitoring.
6. Click **Save**.

   The **Create Monitoring Plan** page appears.

7. Update as necessary the default values, which provide for monitoring plan and rule definition fields. See Reviewing network resource default values on page 74 for details about the **Alert if response time greater than** and **OR no response within** fields.
8. Click **Save**.

   The TCP resource summary page appears displaying the monitoring plans.

Creating a PING (ICMP) resource monitor

OpenEdge Management supports setting up a large number of network resources simultaneously for which you can define PING activities individually. OpenEdge Management can PING a network resource to determine its availability. OpenEdge Management sends an ICMP echo request to an IP host and returns a response time and status.

To define an ICMP resource monitor:

1. Select **Resources** > **New** > **ICMP** in the OpenEdge Management console menu.

   The **Create new ICMP** page appears.
2. Enter values in the **Name** and **Description** fields.
3. In the **Host** field, enter the name of the host or IP address where the port is located.
4. Ensure that the **Enabled** field is checked to begin monitoring.
5. Click **Save**.

   The **Create Monitoring Plan** page appears.
6. Update as necessary the default values, which are provided for monitoring plan and rule definition fields. See Reviewing network resource default values on page 74 for details about the Alert if response time greater than and OR no response within fields.
7. Click **Save**.

   The ICMP resource summary page appears displaying the monitoring plans.

Using the HTTP resource monitor to determine Web site status

OpenEdge Management supports a Hypertext Transfer Protocol (HTTP) resource monitoring feature that allows you to monitor various aspects of a Web page. Businesses that rely heavily on the Internet to convey information to customers and conduct business transactions need to ensure that their company's Web site is reliable, accurate, and accessible. Establishing HTTP resource monitors for these Web sites can help fulfill this need.

Like other resources you can set up in OpenEdge Management, characteristics of an HTTP resource monitor are defined in properties that you establish. You can also establish optional authentication properties for an HTTP resource monitor. The purpose of this process is to provide basic authentication capabilities. However, this authentication scheme is not considered to be a secure method of authentication unless it is used with another external security system because username and password data are passed over the network as clear text.

General considerations

Review the following information about HTTP rules and Web page redirection before you set up HTTP resource monitors.

HTTP rules

HTTP monitors use a unique set of resource rules to monitor the status and content details of a Universal Resource Locator (URL). When enabled, the *status rule* can indicate if:

- The URL you are monitoring was redirected
- The GET method was successful and was accomplished within the time range you specified

Similarly, the *content rule* can also be enabled to indicate content details such as whether:

- A hash value comparison should be performed
- A search should be performed on a given page for the search string you specify

For each schedule you define for a URL that you intend to monitor, you also define status and content rules to suit your specific URL monitoring needs. See Editing rules for the HTTP resource monitor on page 82 for details about editing status rules and content rules.

Web page redirection

When attempting to retrieve a web page, your client might be directed to a page other than the one initially requested.

This situation can occur for the following reasons:

- **Redirection** — Occurs when the server informs the client that it is being redirected and sends the client to the new page. When redirection occurs within the context of using an HTTP Monitor, a redirection alert occurs if you have selected the Alert if URL was redirected option on the **Edit status rule** page.
- **Aliasing** — Occurs when the server sends the client to a different page without informing the client. The client believes that it has retrieved the Web page that was originally requested. Because the server does not inform the client that it has been sent to a different page, the HTTP Monitor cannot trigger redirection alerts, even if the Alert if URL was redirected option is selected on the **Edit status rule** page.

Creating an HTTP resource monitor

You can create an HTTP resource monitor to determine the availability of a Web site.

To define an HTTP monitor:

Select **Resources** > **New** > **HTTP** in the OpenEdge Management console menu.

The **Create new HTTP** page appears.

You can now set up the resource monitor's properties.

Setting up the properties of an HTTP resource

You must set properties for an HTTP resource you create.

To set the HTTP resource monitor's properties:

1. Enter the name of the HTTP monitor in the **Name** field. A maximum of 32 alphanumeric characters and underscores is allowed.
2. Enter the HTTP monitor description in the **Description** field.
3. Enter the name of the URL to monitor in the **URL** field. HTTP is the only supported protocol. If the URL does not begin with `http://`, OpenEdge Management prepends it.
4. Enter the name and port number of the proxy server in the **Proxy server** field that identifies this URL (optional). The correct format to enter this information is <*`proxy`*>:<*`port`*>.
5. Enter a set of name-value pairs to be appended to the URL in the **Form data** field.
6. Enter HTTP headers, one per line, in the **HTTP headers** field. The acceptable header format is <*`header name:header value`*>. Field names are case insensitive.

   The following example shows HTTP headers using this format:

   ```
   Accept:img/*;q=0.0,img/gif;q=0.0,text/*;q=0.0
   ```

7. Decide whether you want OpenEdge Management to allow the server to set cookies. The **Accept cookies** option is selected by default. Clear the option if you do not want the server to set cookies. Cookies are not retained as persistent to the disk.
8. Decide whether you want to enable the HTTP monitor. The **Polling Enabled** option is selected by default. Clear the option if you want to turn polling off for the monitor.

   You can now set the HTTP monitor's authentication properties.

Setting the Authentication properties of an HTTP monitor

Setting HTTP resource authentication properties is optional. However, depending on the Web site you are accessing, values might be required in these fields. For example, if you intend to link to a site that is typically secured, such as a bank's Web site, you might need to supply values in all the authentication fields. In contrast, public Web sites that are interested in providing information to all users who want access to information generally do not require an authenticated connection.

To set the HTTP resource monitor's authentication properties:

1. Enter the name of the URL realm that you want to access in the **URL realm** field.
2. Enter the name of the authenticated user in the **URL user name** field. The value you provide in this field is not authenticated when you enter this information.
3. Enter the password associated with the user name in the **URL password** field.
4. If accessing your URL involves a proxy server, you must also enter the following values:
   - The name of the realm associated with the proxy server in the **Proxy server realm** field
   - The name of the user in the **Proxy server user name** field
   - The password associated with the user name in the **Proxy server password** field
5. Click **Save**.

   The **Create Monitoring Plan** page for the HTTP Monitor appears.

   You can now set up the HTTP resource monitor's monitoring plans.

Setting up the HTTP resource monitoring plans

You determine the monitoring plan and rules for an HTTP resource monitor you create.

establish HTTP resource monitoring plans and rules:

1. Click on the **Available Schedules** drop-down list to select a schedule.
2. As necessary, change the default number that appears in the **Polling Interval** field.
3. To disable the alerts option, clear the **Alerts Enabled** option.

   If you leave this option selected, you should review and possibly change the additional alert-related fields presented in the Rule definition section. See Editing rules for the HTTP resource monitor on page 82 for more information.

4. Select the **Trend Performance Data** option to store data to the OpenEdge Management Trend Database.
5. As necessary, change the default number that appears in the **Trend Performance Data every** field. See Trending considerations for network resources on page 74 for more information about the **Trend Performance Data option** and the **Trend Performance Data every** field.
6. Click **Save**.

   The **Edit status rule** page appears.

   You can now edit the HTTP resource monitor's rules.

Editing rules for the HTTP resource monitor

You can edit all of the rules-related definitions for an HTTP resource monitor.

The following status rule definition fields are unique to the HTTP resource monitor:

- The **Enable status rule** option allows you to either enable or disable all of the criteria specified in the rule definition.
- The **Alert if URL was redirected** option triggers an alert to inform you that the URL was redirected to a page other than the one you expected to access. This alert is generated even if the response time and status code results are acceptable. You can also set the boolean expression **Alert if response time greater than or no response within xx milliseconds** for each of these phrases. When you finish, click **Save**. The **Edit content rule** page opens.

  The **Enable content rule** option is unique to the HTTP resource monitor. The option allows you to either enable or disable all of the criteria specified in the rule definition. This page defines two means of determining the validity of the page returned, based on your URL definition. You can set one or both of the following fields:

  - **Page is static, use hashcode comparison** — Select this option if the page associated with the URL you have defined is static. If the data displayed on the page does not change, OpenEdge Management performs a compare based on hashcode details.
  - **Alert if search string found in page** — Select this option when you know a particular string should appear on the page you want to access.

Reviewing the HTTP resource monitor settings

When you finish setting up an HTTP monitor and click **Save**, you can view a summary of the various values you entered or accepted.

Click **Status rule** to display status rule definition data. Click **Content rule** to display content rule definition data. You can choose to edit values on each of the pages that appear.

6

Monitoring File Resources

You can monitor a log file or other file by creating a resource monitor. This chapter describes how to create log file monitors (which monitor file contents) and file monitors (which monitor file characteristics).

For details, see the following topics:

- File resource monitoring overview
- Verifying mapped network drives
- Reviewing log file and file resource default values
- Creating a log file monitor
- Enabling the log file monitor
- Editing, copying, or deleting a log file monitor
- Creating a file resource monitor
- Creating a file monitor

File resource monitoring overview

When you create a *log file monitor*, you choose the specific file content you want to monitor. Your ability to focus on specific content is valuable because log files typically contain considerable data. For example, you might be especially concerned about tracking system errors or monitoring any warnings that are recorded in the log.

You create a *file monitor* to check on the characteristics, rather than the content, of a file. You can create a file monitor that will gather specific information about a file, such as its size or age.

Because the log file monitor and the file monitor are concerned with different aspects of a file, you can create one of each for the same file. For example, you might create a log file monitor for the AdminServer log (`admserv.log`) to keep you informed of specific areas of concern in the file's contents. You can also create a file monitor for the AdminServer log that will check one or more of the file's characteristics, such as its size or growth rate. If you find that the AdminServer log is growing too quickly, you can lower the logging level so that less data is written to the log.

Log file monitor features

Log file resource monitors have characteristics similar to other OpenEdge Management resource monitors. You define monitoring plans and rules for log file monitors just as you can for all resource monitors. You can also elect to use OpenEdge Management-supplied default values to create log file monitoring resources.

Unlike other OpenEdge Management resource monitors, however, log file resource monitors have several unique qualities that you can use to define their monitoring capabilities, such as:

- **A bookmark feature** — The bookmark feature helps ensure that log files are monitored unobtrusively and according to search criteria you define. See Using a bookmark on page 89 for information about setting and working with bookmarks.
- **A search criteria feature** — You specify search criteria in the log file monitor's rule properties and identify expressions that you can use to obtain specific information about log file contents. See Specifying search criteria on page 95 for information about search criteria.
- **A library feature** — The library feature allows you to define a collection of independently stored search criteria and rule sets that you can share (including by importing and exporting across different machines) among the log file monitors you create and maintain. See Associating rule sets with log file monitors on page 94 for information about library features. For details about importing and exporting search criteria and log file rule sets, see Exporting and Importing in OpenEdge Management on page 133.

File monitor features

You can monitor the following file attributes by creating a file monitor:

- The file's existence
- The file's size
- The file's age
- The file's growth rate
- If the file has been modified

The file monitor does not support rule sets.

Verifying mapped network drives

If you want to monitor a log file or other file and OpenEdge Management is installed in a supported Windows system, you must make sure that OpenEdge Management can see your mapped network drives.

To verify the mapped network drives in a Windows XP system:

1. Choose **Start** > **Settings** > **Control Panel**. Choose **Administrative Tools**, and then click **Services**.
2. Select **AdminService for OpenEdge**, and then click the link to start the service.
3. Double-click **AdminService for OpenEdge**, and then choose the **Log On** tab.
4. Select **This Account** and specify a user with administrative privileges for the account. (You must override the default value `LocalSystem` with this user information.)
5. Enter the password, and confirm the password entry.
6. Click **OK**.
7. If required, restart the AdminServer.

   Now you can define file-specific details to create and edit log file and file monitors.

Reviewing log file and file resource default values

OpenEdge Management supplies default values for the log file monitor and the file monitor. You can keep the default values established by OpenEdge Management, or you can change the values. Whatever default values exist are automatically used when you create a log file monitor or file monitor.

If you choose not to change the default values, you still have the option of altering the values when you create a new monitor. You simply remove the default values and replace them with values you choose.

Before you begin to create a log file monitor or a file monitor, review the default settings. Performing this review can:

- Help standardize resource rule definitions according to your company's needs
- Minimize the duplication of effort involved in unnecessarily creating new or editing existing file resource monitors
- Produce useful resource monitoring data as quickly as possible

To display the log file monitor and file monitor default values:

1. Select **Options** > **Resource Monitor Defaults** in the OpenEdge Management console menu.

   The **OpenEdge Management Resource Monitoring Defaults page appears.**
2. Click **File Resource Defaults**. The **File Resource Defaults** page appears.

   If the default values are still not what you want, you can change the values of the individual log file or file monitor when you create a new resource monitor. Changing default values does not affect any resource monitor already created.

   You can revert back to the original OpenEdge Management-supplied default values at any time by selecting **Restore Defaults**.
3. Change the default values as necessary, and click **Submit**.

Creating a log file monitor

Creating a log file monitor involves the following tasks:

1. Defining properties on page 88
2. Using a bookmark on page 89
3. Creating a monitoring plan on page 90
4. Adding log file rules on page 91
5. Adding log file rule sets on page 92
6. Associating rule sets with log file monitors on page 94
7. Specifying search criteria on page 95
8. Creating search criteria on page 96
9. Enabling the log file monitor on page 100

Typically, you set up one log file resource monitor at a time on a log file. However, you can set up a variety of monitoring plans, consisting of a number of different schedules that are all defined and enabled to run at different times, for a specific log file resource. It is most beneficial to have only one *active* log file monitor enabled at a time for any given log file.

Once you create a log file monitor, you can edit it, copy it as the basis for a new log file monitor, or delete it.

Defining properties

The log file monitor properties you define are the log file monitor name and description and the name of the file you are monitoring.

To define the log file properties:

1. Select **Resources** > **New** > **Log File** on the OpenEdge Management menu bar.

   The **Log File Monitor Resource** page appears.
2. Enter the log file monitor name in the **Name** field.
3. Select the **AdminServer** name from the drop-down list.
4. Enter the log file description in the **Description** field.
5. Enter the file name of the log file you are monitoring in the **Filename** field. You can enter the path to the log file and the log file's file name using either UNIX or Windows notation conventions.

   It is possible that you might need to follow additional steps to see your mapped network drives and to let the log file monitor you create recognize mapped network drives. See Verifying mapped network drives on page 86 for general instructions about how to perform this task.
6. Set the value for the **On First Poll** field.

   On the first poll of a log file, you must identify how you want the log file monitor to interact with existing log file data. You select one of two options:

   - **Search entire file** tells the log file monitor to search all preexisting log file data. If your log file is a relatively small file, consider this option.
   - **Search only new data** tells the log file monitor to **ignore** any preexisting log file data and perform the setup needed for subsequent polls based on the bookmark properties set at the current end of file.
7. Click **Save**.

Using a bookmark

You create a *bookmark* and insert it into your log file to indicate the exact point from which the log file monitor is reading data. To accurately monitor and manage a log file, the log file monitor must know its exact position within the file and have a dependable means to reset its file position if it loses its place within the file due to *truncation*. If you opt not to set a bookmark, you choose which truncation action OpenEdge Management follows: searching only new data, or searching the entire file.

You are urged to use bookmarks with monitors you create for OpenEdge log files. However, using the bookmark feature in non-OpenEdge log files might cause problems using the log. If you have an application that parses out its own log file, it could get tripped up by information it considers foreign.

You might detect a small performance gain in running the log file resource monitor when the log file is truncated. However, you will sacrifice accuracy within your log file data, exposing the data and monitoring operations to an unnecessarily high level of risk and unpredictable accuracy in your log file reads.

Bookmark and truncation considerations

Due to the volume of data stored, some log files have the potential to grow very large. A database administrator (DBA) will typically truncate these files periodically to minimize the amount of space they occupy on a system. The existence of bookmarks within a log file ensures that regardless of what happens to your file when it is truncated, your log file monitor always has the ability to reposition itself accurately within the file.

Data can be truncated at any point within a file. Consequently, when bookmarks are not used and a file has been truncated, the log file monitor has only a limited number of options to determine its new position. Bookmarks represent your most reliable means to reposition your log file monitor accurately so that the data you use is accurate.

For example, if a DBA truncated 500 lines of data at the beginning of a log and, in the meantime, 1000 more data lines were added at the end of the file, the log file will reposition the new reference mark somewhere within the last 500 data lines; having "lost" 500 lines, the log file monitor only knows that it has to read ahead 500 lines into the new data and reposition itself. In this situation, the fact that 500 lines of data have been truncated has been obscured by the existence of the new data. Not having a bookmark to reference as a point of origin, the log file monitor repositions itself 500 lines into the "newer" data lines that have been written. Unfortunately, this approach incorrectly allows the data to be misread, since the balance of the 500 newer data lines are left unread. Had bookmarks been used in this circumstance, this error would have been avoided.

Note: Under certain circumstances, you might not be concerned with unread data. For example, if you are performing software testing activities, log files might be deleted frequently. Therefore, you might not need to set bookmarks to ensure accurate repositioning occurs after a file has been truncated.

If the monitored log file has changed since the last poll and the bookmark feature is enabled, a new bookmark is inserted.

Setting a bookmark

You can either set a custom bookmark or use the last line as the bookmark.

To set a bookmark from the **Log File Monitor Resource** page:

1. In the **Bookmark** section, select **Use bookmark**.
2. Select one of the following ways to set bookmark attributes:
   - **Use custom bookmark (Specify bookmark)**

You can choose to either use the default bookmark named **Fathom_Bookmark** or define a new bookmark. If you define a new bookmark, the value you enter is not restricted to a certain length. However, for practical purposes, you might want to consider keeping it a reasonable length.

Additionally, you can choose to prepend a time stamp to your bookmark name selection. You should **always** select this option. The date and time value uniquely qualifies each bookmark that has this data prepended to its bookmark name. In situations where the log file monitor must reposition itself within the log file, this type of bookmark will always be found.

- **Use last line as bookmark**

  This option records the last line of a given poll as a bookmark internally (Unlike the custom bookmark, this information is not written in the log file being monitored.)

  Choosing this option indicates that you want to use the last line as the bookmark. The last line is referenced as it is randomly defined in the file, and it is likely that the line will not be unique. If your log file has a date and time stamp on each line, or other data that makes each line unique, select the Each line unique option. This refinement will increase the log file monitor's accuracy.

  If you do not use the **Each line unique** option, you must specify one of the possible two **Truncate Action** selections from the drop-down list.

3. Click **Save**.

   The **Create Monitoring Plan** definition page appears.

Truncating a log file

If you choose to truncate a log file, the file does not use a defined bookmark to reference its position within the log file. Therefore, you must indicate how you want the log file monitor to interact with the log file when the file has been truncated.

You might detect a small performance gain in running the log file resource monitor when the log file is truncated. However, you will sacrifice accuracy within your log file data, exposing the data and monitoring operations to an unnecessarily high level of risk and unpredictable accuracy in your log file reads.

You can choose either of the following truncation actions:

- **Search only new data**

  OpenEdge Management resets the log file monitor's internal place marker to the end of the file on its first poll after the file has been truncated.

- **Search entire file**

  OpenEdge Management searches the log file from beginning to end on its first poll after the file has been truncated. Remember that this technique will ensure that no data is missed when a file is truncated, but much of the file could be searched a second time. If the file is very large, this can be a slow operation.

Creating a monitoring plan

Once you make your decisions about the bookmark, create the log file monitor's monitoring plan. Begin by viewing the monitoring plan's default values.

The following default values are identified:

- **Available Schedules** — A list of all schedules available for this monitoring plan
- **Polling Interval** — The interval at which the resource is polled

- **Alerts Enabled** — An indication of whether alerts defined for this resource are enabled

You can elect to accept or change any of the default values.

To continue creating the monitoring plan from the **Create Monitoring Plan** page:

1. Click **Save** to save the monitoring plan definition.

   The **Monitoring Plan** page appears.

   You can now add rules or rule sets to this monitoring plan.

2. Under **Rules selected for this plan**, click either **Add Rule** or **Select Rule Sets**.

   You can now continue with the Adding log file rules on page 91 section for details about including rules or the Adding log file rule sets on page 92 section for details about including rule sets in the plan.

Adding log file rules

You can include one or more rules and/or rule sets in a log file monitor. When you add a rule, you must also select search criteria and then associate it with severity and alert information to fully define the rule.

Log file monitor rule components

A log file monitor rule consists of the following components:

- **Search criteria** — The core elements of log file rules. A search criterion allows you to create and store exactly what data a log file resource monitor will look for when it searches a log file. OpenEdge Management- and user-defined search criteria are stored in the library so they can be shared and reused in various log file monitor rules.

  See Specifying search criteria on page 95 for details about the search criteria.

- **Severity** — One of the four standard, user-selected severity levels: **Information**, **Warning**, **Error**, or **Severe**. This information is stored in the **Severity** field.

- **Alert** — Identifies which action OpenEdge Management performs when an alert occurs. This user-selected information is stored in the **On Alert Perform Action** field.

You can include a rule in the log file resource monitor in three ways:

- By adding it to a single log file resource under the resource's monitoring plan, in which case the rule definition **cannot** be shared by other log file resource monitors.
- By adding it to one or more rule sets, in which case the rule **can** be shared among several log file resource monitors.

  Note: Edits to any user-defined search criteria immediately affect all current instances in which the search criteria is used in individual rules or rule sets. These edits will also be used in instances created in the future. You can edit or delete OpenEdge Management-supplied search criteria.

- Associating it with a log file resource monitor. An identical rule definition, defined in a rule set, can also be applied to this same log file resource monitor. In this situation, the individual, or locally defined rule, takes precedence. The rule definition associated with the rule set will be listed as part of the rule set, but the status icon for this duplicate rule will be gray and the status will be listed as **Inactive**.

Adding log file rule sets

A log file rule set is composed of one or more rule definitions. The purpose of a rule set is to help you manage a large number of log file monitors consistently and efficiently by sharing rule definitions. By creating and selecting appropriate log file rule sets to be added to log file resource monitors, you eliminate the requirement to add log file rules on an individual basis to each log file monitor.

Once a rule set exists, you can access it from the list frame and then associate it with the monitoring plan of a log file monitor. Then, when the monitoring plan is active and the log file monitor is polled, OpenEdge Management can evaluate each rule in the set. Keep in mind that a log file monitor's monitoring plan can have one or more rule sets associated with the plan.

Accessing existing log file rule sets

The OpenEdge Management Library provides access to existing log file rule sets.

To access existing log file rule sets:

1. Select **Library** > **View** > **Rule Sets** in the OpenEdge Management console menu. The **Library.Rule Sets** page appears.

   The following predefined categories display:

   - **AppServer**
   - **Database**
   - **LogFile**
   - **MSS DataServer**
   - **NameServer**
   - **Oracle DataServer**
   - **PAS for OpenEdge**
   - **SonicMQ Adapter**
   - **WebSpeed**
   - **Web Services Adapter**
   - **OE Web Server**

2. Click **LogFile** to display a list of all available log file rule sets.
3. To display details about a specific rule set, click the individual rule set name in the list frame. You can also edit the rule set.

Creating a new log file rule set

A log file rule set allows you to centrally define rules at the library level and associate rule sets with one or more log file resource monitors.

Rules within rule sets are composed of three rules-required components: search criterion, severity level, and alert details. Library-based log file rule sets are comprised of rules that have search criteria you define, or you can select from the search criteria stored in the library.

To set up library-based log file rule sets:

1. In the OpenEdge Management console menu, choose one:
   - Select **Library** > **New** > **Log File Rule Set**.
   - Select **Library** > **Go to Library** > **Create Log File Rule Set**.

   The **Create Log File Rule Set** page appears.
2. Enter the name of the rule set in the **Name** field.
3. Enter a description of the rule set in the **Description** field.
4. Click **Save**.

   The **Log File Rule Set: <rule set name>** page appears.
5. To add a rule to this rule set, click **Add Rule**.

   The **Create new Log File Rule** page appears.
6. Choose one of the following options:
   - Use existing search criteria. Select values available from the drop-down list fields on the **Create Log File Rule** page. See Selecting search criteria to define a local rule on page 99 to complete the steps for this option. Then proceed with Step 7 in this procedure.
   - Use new search criteria. Continue with Step 7 in this procedure.
7. Click **Create Criterion**.

   The **Create new Search Criterion** page, from which you will create a new rule to add to the rule set, appears.
8. Enter a value in the **Name** field.

   The search name can be a maximum of 32 characters, and it must not contain spaces between words or use special characters such as an asterisk (*), an ampersand (&), or a period (.).
9. Enter a description in the **Description** field.
10. In the **Search Text** field, enter the search string you want this rule to look for in the log file.

    The entry must be consistent with the type of search you are performing.
11. In the **Search Type** field, select either **Literal Search** or **Regular Expression**.

    See Specifying search criteria on page 95 for a detailed description of these search types.
12. Store this search criterion in a **Search Criteria** category. For example:
    - **Use Existing Category** (default value) — From **Category**, scroll through the list of OpenEdge Management predefined categories and select one.
    - **Use New Category** — In the **Category** field, use the predefined **misc** category, or enter the name of a new category.
13. Click **Save**.

    The **Create Log File Rule** page reappears, displaying the newly identified search criterion values in the **Choose Criteria Category** and **Choose Search Criteria** fields.
14. Identify the remaining properties for this rule:
    - In the **Severity** field, select the level of severity you want to set for this rule.

- Select the **Always throw new alert** option to generate an alert each time the rule is violated.
- In the **On Alert Perform Action** field, select the action you want OpenEdge Management to perform when an alert triggers.

15. Click **Save**.

    The **Log File Rule Set** page reappears and displays the newly created rule in the **Rules selected for this Rule Set** section.

16. If required, repeat Step 5 through Step 15 to add additional rules to this rule set.

    Your rule set appears in the list frame under the **LogFile** category. See Associating rule sets with log file monitors on page 94 for details about associating the rule set with a log file monitor.

Associating rule sets with log file monitors

Once you have created at least one rule set, you can associate it with a log file monitor.

To associate a rule set with an existing log file monitor:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select a Log File resource.

   The summary information about the log file monitor appears.

3. In the **Monitoring Plans** section, click **Edit** against the monitoring plan for which you want to add a rule set.

   The current monitoring plan and rules appear in edit mode.

4. Click **Select Rule Sets**.

   The rule sets that currently exist appear.

5. Select the check box associated with each rule set you want to add to the log file monitor.

6. Click **Save**. The **Monitoring Plan** page reappears and includes the new rule sets.

   Note the following points relative to the rule set addition:

   - All rule sets are identified by a rule set icon. The icon appears to the left of each rule set name in the **Rules selected for this plan** list.
   - Rule sets appear after individual rules in the list.
   - Any edits you make to individual rules within a rule set, or to the rule set itself, immediately affect all instances of the rule sets—current and future.
   - Individual rule definitions take precedence over the same rule that is defined within a rule set.

7. To review the specific rules contained in any rule set, click the rule set link in the **Rules selected for this plan** section of the **Monitoring Plan** page.

   The **Log File Rule Set** page associated with the specific rule set appears, displaying each rule that it contains.

8. You can perform the following actions:

- Click **Remove From Plan**. The **Monitoring Plan** page appears. You can elect to remove any of the rule sets defined for the current plan. Any changes you make to rule sets affect **only** this plan's use of the rule set.
- Click **View Library Definition**. A new page appears, allowing you to link to and display the detail properties that compose each rule. You can edit, copy, or delete individual rules from the library-stored definition at this time. Any changes you make will also affect the current and future uses of this rule.

9. Repeat these steps to associate rule sets with additional monitoring plans for this log file monitor.

Specifying search criteria

Search criteria identify expressions that OpenEdge Management can use to obtain specific information about the contents of log files. You can use either a text literal or a Perl 5 regular expression, which you select or create as one of the rule properties, as an individual search criterion. When you specify search criteria, OpenEdge Management searches the log file for strings and tries to match the search criteria you specify.

OpenEdge Management provides some default search criteria; you can also create your own search criteria. This section describes each type of search criteria.

Accessing search criteria

You access existing search criteria from the OpenEdge Management Component Library.

To access either OpenEdge Management- or user-defined search criteria:

1. Select **Library** > **View** > **Search Criteria** from the OpenEdge Management console menu bar. The **Library.Search Criterion** page appears.

   The following predefined categories display:

   - **AppServer Internet Adapter**
   - **AppServer Broker**
   - **AppServer Server**
   - **Database**
   - **Messengers**
   - **Miscellaneous**
   - **MSS DataServer Broker**
   - **MSS DataServer Server**
   - **NameServer**
   - **OE_Replication**
   - **Oracle DataServer Broker**
   - **Oracle DataServer Server**
   - **Progress Application Server**
   - **SonicMQ Adapter Broker**
   - **SonicMQ Adapter Server**

- **WebSpeed Broker**
- **WebSpeed Server**
- **Web Services Adapter**
- **OE Web Server**

2. Click the category from which you want to select a predefined search.

   An expanded sub-list of the search criterion page appears.

3. Select the specific search you want to access.

   The Search Crietrion properties summary page appears.

4. Review, edit, copy, or delete the search, as needed. Remember the following points if you elect to change or delete any existing search criteria:

   - If you change any values associated with a search criterion that is currently associated with log file monitors, those changes will propagate through all existing instances of the log file monitors that use that specific search criterion.
   - If you attempt to delete a search criterion that is the last one currently defined in a given category, a message displays asking if you would like to delete the category at the same time.
   - If you attempt to delete a search criterion that is associated with any log file monitor, OpenEdge Management displays a warning message.

     This message appears regardless of the log file monitor's status as enabled or disabled.

Creating search criteria

You create the search criteria phrases to be accessed by and shared among several log file monitors from the OpenEdge Management library. These search criteria phrases, along with the severity level and alert detail you define, are the rule properties you use to create log file resource monitor rules and rule sets.

You can use the search criteria that you create and store at the library level in multiple log file monitors simultaneously. You can also share search criteria by exporting it from or importing it to other machines. For details about importing and exporting search criteria, see Exporting and Importing in OpenEdge Management on page 133.

Within your work environment, it is important for you to know when a database is unexpectedly down. Perhaps the effectiveness of your company's operations have suffered lately due to unplanned database down time. Obtaining information from the log file about these events can help you better research the issue.

To help you assess what might be happening, you can set up a literal search for the words **database down**. Once you enable the search against the log file, OpenEdge Management checks the file for an occurrence of the phrase, based on your polling cycle, and reports on its findings.

You must provide values for the **Search Type** and the **Search Text** fields on the **Create Search Criterion** page. The **Search Type** is a drop-down list. Depending on the type of expression you enter in the **Search Type** field, select either the literal search string or Perl 5 regular expression option. In the **Search Text** field, you enter the specific value for the type of expression you selected.

Another example of a literal text string you might search for in a log file is **abnormal shutdown**. A search based on this value yields a match only if the exact expression **abnormal shutdown** is found in the file. In contrast, a Perl 5 regular expression search supports a variety of notations that you can use for pattern matching, potentially yielding broader search results.

Note: The **Search Text** field does not impose any length limits on your entry. For practical reasons, however, you might want to limit the length. Also, keep in mind that the **Search Text** field is case-sensitive.

To create a search criterion:

1. In the OpenEdge Management console menu, choose one:
   - Select **Library** > **New** > **Search Criteria**.
   - Select **Library** > **Go to Library** > **Create Search Criteria**

   The **Create new Search Criterion** page appears.
2. In the **Criterion Properties** section, enter a value in the **Name** field.

   The search name can be a maximum of 32 characters, and it must not contain spaces between words, or use special characters such as an asterisk (*), an ampersand (&), or a period.
3. Enter a description in the **Description** field.
4. In the **Search Text** field, enter the search string you want the log file resource monitor to look for in the log file.

   The entry must be consistent with the type of search you are performing.
5. In the **Search Type** field, select either **Literal Search** or **Regular Expression**.

   See Specifying search criteria on page 95 for a detailed description of these search types.
6. Store this search criterion in a **Search Criteria** category by choosing either:
   - **Use Existing Category** — From **Category**, scroll through the list of OpenEdge Management predefined categories and select one.
   - **Use New Category** — In the field, use the predefined **misc** category, or enter the name of a new category.
7. Click **Save**.

   The search criterion is now available. You can display it from the category to which you assigned it.
8. Repeat these steps for each additional criterion you want to create.

Using Perl 5 expressions to create search criteria

The following table identifies some Perl 5 regular expressions you can use. Similarly, the Examples using the Perl 5 regular expressions table suggests specific examples based on this notation.

Table 10: Perl 5 regular expressions

Perl 5 expression	Description
.	Matches exactly one character, regardless of what the character is.
?	The preceding item is optional and matched at most once (error if no preceding item).
*	The preceding item will be matched zero or more times (error if no preceding item).
+	The preceding item will be matched one or more times (error if no preceding item).

Perl 5 expression	Description
^	Match at beginning of a line.
$	Match at end of a line.
{n}	The preceding item is matched exactly n times (error if no preceding item).
{n, }	The preceding item is matched n or more times (error if no preceding item).
{,m}	The preceding item is optional and is matched at most m times (error if no preceding item).
{n,m}	The preceding item is matched at least ntimes, but not more than m times (error if no preceding item).
[abc]	Matches the characters a OR b OR c.
[a-z]	Matches any character from a to z.
[^abc]	Matches any character EXCEPT a, b, or c.
\d	Matches exactly one digit.
\D	Matches any character EXCEPT a digit.
\w	Matches exactly one letter, number, or the underscore character(_).
\W	Matches any one character EXCEPT a letter, number, or the underscore character.
\s	Matches exactly one character of white space (for example, spaces, tabs, newlines, or any character that would not use ink if printed on a printer).
\S	Matches any character that is NOT a white space.
\	Dereferences metacharacters (called "quoting").
|	Separates two or more choices such as either|or behavior.

The following table provides some examples based on the Perl 5 regular expressions identified in the above table.

Table 11: Examples using the Perl 5 regular expressions

Example	Description
^error	Matches the exact word `error` only when it appears at the beginning of a line.
\(9239\)$	Matches the exact entry `(9239)` only when it appears at the end of a line.
da.*e	Matches the exact words `date`, `daze`, `database`, and `dattape`. This Perl 5 regular expression, .*, is similar to the wild card * on UNIX.

Example	Description
abc|abd|abe	Matches `abc`, `abd`, and `abe`.
b.d	Matches `bad`, `bud`, and `bid`, but not `bald`.
da.....e	Matches `database` and `dattape`, but not `date` and `daze`.
3.14	Matches `3.14`, `3f14`, and `3814`.
3\.14	Matches `3.14`, but not `3f14` and `3814`.
ab?c	Matches `ac` and `abc`.
ab*c	Matches `ac`, `abc`, `abbc`, `abbbc`, and so forth.
ab+c	Matches `abc`, `abbc`, and so forth, but not `ac`.
d\.*z	Matches `dz`, `d.z`, `d..z`, `d...z`, and so forth.
d.*z	Matches `da*z`, `db*z`, `dc*z`, and so forth.
1\.\d\d	Matches any three-digit floating point number from 1.00 to 1.99******.
a\Dc	Matches `abc`, `a&c`, and `aFc`, but not `a2c` or `a8c`.
a\wc	Matches `abc`, `aGc`, and `a_c`, but not `a%c`.
a\Wc	Matches `a%c`, `a?c`, and `a c`, but not `abc`, `aGc`, or `a_c`.
a\sc	Matches any three-character string starting with `a` and ending with `c` whose second character is a space, tab, or newline.
a\Sc	Matches any three-character string starting with `a` and ending with `c` whose second character is not a space, tab, or newline.
ab{3,5}c	Matches `abbbc`, `abbbbc`, `abbbbbc`, only.
.{3,5} pentane	Matches `cyclopentane`, `neopentane`, and `isopentane`, but not `n-pentane`.
a[bc]d	Matches `abd` and `acd`, only.
a[a-z]c	Matches any three-character string starting with `a` and ending with `c`, and whose second character is any letter from `a` to `z`, inclusive.

Selecting search criteria to define a local rule

Adding and defining a rule requires you to select search criteria and to associate severity and alert information with it.

To add a rule to a log file monitor:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the required Log File resource.

   The summary information about the log file monitor appears.

3. In the **Monitoring Plans** section, click **Edit** against the monitoring plan for which you want to add a rule set.

   The current monitoring plan and rules appear in edit mode.

4. Click **Add Rule**.

   The **Create new Log File Rule** page appears.

5. From **Choose Criteria Category**, select the category in which the search criteria you want to use is defined.

6. From **Choose Search Criteria**, display all currently defined search criteria for the category previously selected and select a search criteria.

   The search text automatically appears in the **Criterion Search Text** field, and a description automatically appears in the **Search Criterion Description** field.

7. From **Severity**, select the level of severity you want to set.

8. Select the **Always throw new alert** option if you want OpenEdge Management to generate an alert each time the rule is violated.

9. From **On Alert Perform Action**, select the action you want OpenEdge Management to perform when an alert is triggered.

10. Click **Save**.

    The **Monitoring Plan** page reappears and includes the new rule.

11. As necessary, repeat Step 1 through Step 7 to add more individual rules.

    Note the following points concerning individual rules that you add to a log file monitor resource:

    - Individual rules display in alphabetical order and appear in the list before rule sets.
    - Each rule appears with an associated resource status indicator.
    - You can click any rule to display summarized details about it.
    - You must explicitly create an individual rule to add it to the rule set; an individual rule is not directly available for use in any rule set.

12. Choose one of the following to continue:

    - If you want to add rule sets, see Associating rule sets with log file monitors on page 94 for details.
    - If you are done adding rules, click **Done Adding Rules**. The **Log File Monitor** summary page appears, displaying any rules and rule sets selected for a log file monitor.

Enabling the log file monitor

A log file monitor cannot begin polling until you do one of the following tasks:

- Click **Done Adding Rules** on the **Log File Monitor** summary page.

- Edit the log file properties, select the **Enabled** option, and click **Save**.

Polling will begin based on the polling criteria you have set, and polling data will begin to display in the upper right-hand corner of the **Log File Resource** page.

Editing, copying, or deleting a log file monitor

All log file monitors have a summary page. The **Log File Monitor** summary page displays all the values entered to define the bookmark feature, the monitoring plans, and the rules and rule sets for this log file monitor. From the summary page you can:

- Click **Edit**, **Copy**, or **Delete** at the top of this page to access the basic properties to update, copy the values for the entire log file monitor to create a new one, or delete the log file resource monitor.
- Click **Log File Viewer** to see the contents of the log file.
- Click **Default Schedule Plan** (or any plan you have established) to display the monitoring plans and rules for this log file resource monitor for review.
- Click **Edit** in the monitoring plans section to display the monitoring plans and rules and edit them.
- Select **Add Plan** in the monitoring plans section to create additional, unique monitoring plans and rules for this log file resource monitor. Once you create a new plan, the details of the new plan are appended to the plans that already exist.

Note: OpenEdge Management provides a database log file monitor designed to be used specifically with OpenEdge databases. Its functionality is similar to that of the log file monitor. For detailed information about the database log file monitor and how to use it, see the appropriate section in *OpenEdge Management: Database Management*.

Creating a file resource monitor

You can set up OpenEdge Management to monitor a file resource in the following ways:

- **File size** — Examines the file size and triggers an alert if the file size is less than, greater than, equal to, or not equal to a size you specify. You can specify the size in bytes, kilobytes, megabytes, gigabytes, or terabytes.
- **File existence** — Checks to see if the file exists or not. You can choose the condition under which an alert will be generated—for example, if the file does not exist after two polls. You also choose the specific action you want performed if the file is older than the number of seconds, minutes, hours, or days you specify.
- **File growth rate** — Examines the file size and generate an alert based on a size change for a particular period of time. For example, you might specify that OpenEdge Management generate an alert if the file growth exceeds 100 megabytes per day.
- **File modification** — Examines the file data and determines if the file has been modified. You can specify that OpenEdge Management generate an alert if there is evidence that the file has been modified, because either the timestamp has changed or the timestamp and size of the file have both changed.

Before you can set up a file resource monitor, you must ensure that OpenEdge Management can see your mapped network drives.

Trending performance data for file monitor resources

From the monitoring plan, you can trend performance data for file monitors.

Once you enable trending by selecting the **Trend Performance Data** field, you can also override the default value of `1` in the **Trend Performance Dataevery poll(s)** field. Any value that you enter in this field that is greater than the default value `1` causes a file system resource to trend data that is taken from the last poll. No calculations are performed on this data.

Creating a file monitor

To monitor a file, you must create a monitoring plan definition, select rules for the plan, and set the rules' properties.

To define a file monitor:

1. Select **Resources** > **New** > **File** in the OpenEdge Management menu.

   The **Create new File** page appears.

2. Provide the following information:

   - **Name** — Enter the name of the file monitor. You must not use spaces or special characters such as an asterisk (*), ampersand (&), or period (.).
   - **AdminServer** — Select the AdminServer name from the drop-down list.
   - **Description** — Provide an optional description of the monitor.
   - **File Name** — Provide the name of the file you want to monitor.

   The **Enabled** option is selected by default. To disable the monitor, clear the option.

3. Click **Save**.

   The **Create Monitoring Plan** for the file monitor appears.

4. Choose the monitoring plan specifics. When you finish, click **Save**.

   The file monitor's monitoring plan definition and a list of rules selected for this plan appear.

5. To add one or more rules to the plan, click **Add Rule**.

   The **Available File Monitor Rules** page appears.

6. Click the first rule you want to add. If you choose file age, the **File Rule: File Age** properties page appears.

7. Provide values for the following rule properties (which might differ depending on the rule you are adding):

   - **Alert if file older than** OR **Alert severity**
   - **Throw alert after**
   - **Throw additional alerts**
   - **On alert perform action**
   - **Clear alert after**
   - **On clear perform action**

8. Click **Save**.
9. Repeat Step 7, Step 8, and Step 9 for each rule you want to add.
10. When you finish adding rules, click **Done Adding Rules**.

    The file monitor's monitoring definition and rules appear. You can now edit or delete the plan.

7

Monitoring Batch Program resources

You can setup monitoring plans and schedules for a Batch Program resource using the **Batch Program Monitoring Plans** page in the OpenEdge Management. This allows you to process specific conditions when to trigger an alert.

For details, see the following topics:

- Viewing and modifying the default monitoring plan
- Viewing and modifying the default schedule

Viewing and modifying the default monitoring plan

To view and modify the default monitoring plan for a Batch Program resource:

1. From the OpenEdge Management console men, click **Resources** > **Go to Resources**.

   All the resources managed by your console appear in a grid frame.

2. Filter or search for, and select the required Batch Program resource.

   The **Batch Program: <instance name>** page appears.

3. In the **Command and Control** section, click **Monitoring Plans**.

   The **Batch Program Monitoring Plans** page appears with the default monitoring plan for the Batch Program resource.

4. To modify the default monitoring plan, click **EDIT** in the **Monitoring Plans** section.

   The **Edit Default_Schedule Monitoring Plan** page appears.

5. In the **Monitoring plan definition** section, you can:
   a) Click the **Available Schedules** drop-down to select a schedule.
   b) Modify the default **Polling interval** values as required.
   c) Clear the **Alerts Enabled** option to disable the alerts option.
6. Click **SAVE** to save the monitoring plan.

Viewing and modifying the default schedule

To view and modify the default monitoring plan for a Batch Program resource:

1. From the OpenEdge Management console men, click **Resources** > **Go to Resources**.

   All the resources managed by your console appear in a grid frame.
2. Filter or search for, and select the required Batch Program resource.

   The **Batch Program: <instance name>** page appears.
3. In the **Command and Control** section, click **Monitoring Plans**.

   The **Batch Program Monitoring Plans** page appears with the default monitoring plan for the Batch Program resource.
4. To modify the default schedule, click **Default_Schedule_Plan** link in the **Monitoring Plans** section.
5. In the **Monitoring plan definition** section, click the **Default_Schedule** link and then click **EDIT**.
6. In the **Edit Schedule** page, you can:
   a) Modify the **Name** and **Description** of the Schedule in the **Properties** section.
   b) Modify the days and times you want the schedule to include in the **Schedule period** section.
7. Click **SAVE** to save the schedule.

8

Working with Jobs

OpenEdge Management provides several different ways for you to work with jobs:

- You can create a custom job to run once or on a recurring basis.
- You can create a job template and run job instances based on the template's properties.
- You can use one of the predefined database maintenance job templates.

You can work with local jobs and also remote jobs that are running on remote machines. The remote host where the job is to be run is known as the *remote AdminServer*; the local host is also referred to as the *local AdminServer*.

Remote job support requires that you set up remote configuration, in which you install OpenEdge Management on the local host. The remote machine requires an installation of the AdminServer; you must disable OpenEdge Management or OpenEdge Explorer on the remote machine if either one is installed. For more information about remote configuration, including the setup procedures, see the relevant section in *OpenEdge Management and OpenEdge Explorer: Getting Started*.

In addition to setting up remote configuration, you need to ensure that remote job commands and scripts are available on the remote AdminServer machine. The OpenEdge job scripts are automatically included with an installation of OpenEdge Management or OpenEdge Explorer. If you are using scripts or commands from an operating system, be sure they are also available on the remote machine.

OpenEdge Management also provides several predefined database-maintenance job templates that you can use to keep your OpenEdge database performing efficiently. For details about these maintenance templates, see *OpenEdge Management: Data Management*.

For details, see the following topics:

- Creating a custom job
- Creating a job template

- Using a predefined database maintenance job template
- Creating a job instance
- Editing a job
- Copying a job
- Killing a job
- Deleting a job
- Running a job
- Scheduling a job
- Viewing scheduled jobs
- Viewing running jobs
- Viewing job history
- Viewing database maintenance job instances for an individual database
- Working with job templates
- Viewing debug details about jobs

Creating a custom job

A *custom job*, also known as a *job instance* is an individual job that OpenEdge Management executes once or at a regularly scheduled interval. A custom job is generally a batch (noninteractive) task that runs as a background process. When you create a new job, be sure to assign it a unique and distinctive name, so that there is no confusion when you later view job history.

Examples of custom jobs include:

- Hourly database after-image roll over
- Weekly disk defragmentation
- Monthly database index rebuild
- Monthly scan disk

Once you create a job instance, you can schedule it to execute at a regular interval, such as hourly, daily, and weekly, and also at a specific time, such as at startup. Job schedules differ from resource monitor schedules in that job schedules define a frequency of occurrence (every fifteen minutes, for example) while monitor schedules define a block of time (9:00 AM – 5:00 PM, for example). You can create schedules through either the repeat interval or a cron expression.

You can view a list of scheduled jobs as well as check on the status of currently running jobs. Once jobs have completed, you can also query job history data. See Creating a job instance on page 110 and Scheduling a job on page 122 for details about creating and scheduling a custom job.

Creating a job template

OpenEdge Management allows you to create and maintain job templates. The advantage of using a job template is that you can use it more than once. This saves time in recreating the job's specifics and also makes sure that the job based on the job template runs with the same specifications each time.

You do not run the job template itself; instead, you create a unique job instance from the common values in a job template and then run the job instance. See Creating a job template on page 129 for details about creating a job template.

Using a predefined database maintenance job template

OpenEdge Management supplies several predefined database maintenance job templates to help you get started performing some routine database maintenance activities. The following table lists the templates and also indicates which templates are available when you run a job on a local AdminServer or a remote AdminServer.

Table 12: Predefined database maintenance job templates

Job template	Availability on local AdminServer	Availability on remote AdminServer
Configuration Backup	Yes	–
Database Analysis	Yes	Yes
Database Restore	Yes	Yes
Data Compaction	Yes	–
Grow BI	Yes	Yes
Index Compaction	Yes	Yes
Offline Backup	Yes	Yes
Online Backup	Yes	Yes
Truncate BI	Yes	Yes
Unlock Compaction	Yes	–

The Data Compaction and Unlock Compaction job templates allow you to obtain details about the OpenEdge Management Trend Database. For details, see *OpenEdge Management: Trend Database Guide and Reference*.

The information in this Web paper focuses primarily on creating custom jobs and job templates. For definitions and detailed information about using the predefined database maintenance templates, see *OpenEdge Management: Database Management*.

Creating a job instance

When you want to create and run an OpenEdge Management job, you have several options. You can create a job that runs only once or one that runs on a recurring basis. You can also create a template for a job and then run separate job instances based on the specifics you provide only once in the template.

When you want to create a job for one-time or recurring use, or if you want to create a job based on an existing template, you create a job instance.

To create a job instance:

1. Choose the type of job to create:

   - To create a unique job not based on any template, you can choose one in the OpenEdge Management console menu:

     - Select **Jobs** > **New** > **Job**.

     - Select **Jobs** > **Go to Jobs** > **Create Job**.

     The **Job** page appears.

   - To create a job from an existing job template, you can choose one in the OpenEdge Management console menu:

     - Choose **Jobs** > **New** > **Job from a Template**.

     - Click **Jobs** > **Go to Jobs** > **Create Job from a Template**.

     The **Create Custom Job** page appears. Choose one of the existing templates. The **Job** page for that template appears.

2. Complete the following fields:

   - **Name** — Type the job's name. You use the job name when you want to edit the job's properties or check on the job's status. Job names cannot contain spaces or special characters. Be sure to give the job a unique name to avoid confusion when you review job history.

   - **Description** — Provide a sentence or phrase that describes the job's purpose. (Optional, but recommended.)

   - **Resources** — Identify database resources to associate with a job or job action. When you create a job instance, you can associate zero or more databases with it.

     If you have configured remote AdminServers, the list of resources includes both local databases and databases on the remote AdminServers.

     If you are creating a unique job, highlight the database you want to add from the list of available database resources, and click the **Add** arrow icon. The selected database displays in the **Selected** list. To remove a database that you do not want to select, highlight the database name in the **Selected** list and click the Remove arrow icon.

     If you are creating a job from a template and want to include a database, select the database from the list of **Available** resources.

Note:

- If you are using one of the predefined database maintenance job templates, you can choose only one database. If you are using a job template that you have defined, you can choose multiple databases.
- In OpenEdge 12.0, there is no support for the **Job from a Template** functionality to create a job instance in OpenEdge Management that has a 12.0 or higher versioned remote adminserver.

- **User name** — Provide the operating system user account. The operating system user account might not be the same as the OpenEdge Management user account name. If specified, the name must be a valid account on the server machine (or server domain) where the AdminServer and OpenEdge Management are running. In Windows platforms, the name can also include a domain.

  If you are creating a job for a remote AdminServer and you provide a user name and password, these are the credentials used when the job is run. If you do not provide a user name and password, the job is run by the same user running the AdminServer.

- **Group** — On supported UNIX platforms, identify the name of the group. A *group* is a collection of users who share the same privileges.
- **Password** — Provide your password. If the value you enter in the **User name** field does not have an associated password, you must leave this field blank; otherwise, an error message is generated. (See related information about **User name** earlier in this list.)
- **AdminServer** — Choose from the local AdminServer or any remote AdminServers you have configured.
- **Command** — Identify the command you want the job to perform. If you are creating a job instance from a template, it is possible that the command will be predefined for you. If you are creating a unique job, you can define any command that would typically execute from an OS shell, and you can include a full or relative place-name. You can also use environment variables such as `%DLC%` or `$DLC`.
- **Command parameters** — Define the input parameter to the command. The list of parameters takes the same format as from an OS shell. Additionally, you can use environment variables (for example, `$SHELL or %WINDIR%`) and Windows registry values.
- **Working directory** — Identify the command's current working directory. If you specify a directory, the directory must exist. This property defaults to the working directory defined at installation.
- **Input file** — Identify the file used for read redirection with a job's command. This is typically used for any keyboard input the command might require.
- **Output file (stdout)** — Identify the file to which any output written to `stdout` is routed. OpenEdge Management creates a default `stdout` filename based on the name of the job instance. However, you can change it. If you select the Append option, any new information will be added to the existing `stdout` file.
- **Output file (stderr)** — Identify the file to which any output written to `stderr` is routed. OpenEdge Management creates a default `stderr` filename based on the name of the job instance. However, you can change it. If you select the Append option, any new information will be added to the existing `stderr` file.

Note: When you create or edit a job, OpenEdge Management will prefill the `stdout` and `stderr` filenames for you if any of these fields are left blank. The filenames will be the same as the job name, with a different suffix. You can change the filename or remove it if you do not want to create the output files. If you edit the job name, OpenEdge Management will not change the filenames to reflect the new job name. You should review these filenames **before** saving the page to ensure the job creates the expected output files.

- **Environment name=value pairs** — Define application-specific variables to be set in the process context of the task that runs to execute the specified job. If you are creating a job instance from a template, a **name=value pair** will be predefined for you.
- **Debug log file** — Select this option to obtain diagnostic details that can help debug a job. For example, if you set this option and use the **Run Now** feature to run the job, a debug log file will be generated in the default working directory. The debug log file will be named based on the job name with a `.log` extension.

  Once you have submitted the job, OpenEdge Management makes debug trace file data available. See Viewing job history on page 126 for details.
- **Indicate if the job can be used as an action** — Select this option if you want the job to be available as a job action. As with other OpenEdge Management actions, such as an e-mail or log file action, you can set up job actions to automatically trigger in response to alerts.

  You can execute a job in response to the following events:

  - An alert trigger
  - An alert that has been cleared
  - A schedule that is set up so that the job action occurs on a predetermined basis
- **Completion Actions and Alerts** — If you intend to set up actions and alerts for this job, click **Edit**. For more information about this page, see Step 3 in this procedure.

3. After you specify the job instance's properties (or accept the predefined properties if you are using a template), you can do either of the following:
   - Define action and alert occurrences for the job as described in Setting up job completion actions and alerts on page 112.
   - Save the job. You can then run the job immediately, as described in Running a job on page 121, or hold off until another time.

Setting up job completion actions and alerts

When you select the option to monitor and manage information about jobs that have executed, you can:

- Associate job status details with actions to automatically notify the appropriate personnel about the status of the job's execution.
- Set up alerts to trigger if a job execution exceeds the time frame in which the job was scheduled to run.
- Initiate a series of job activities based on exit code values. This activity, known as *job chaining*, allows you to plan job actions to occur for more than one exit code condition. See Job chaining on page 114 for more information.

Job completion actions

You can associate exit codes with existing actions to define which actions occur based on the status of a job's execution. Once a job execution completes, an *exit code* is generated. The exit code indicates whether or not the process succeeded. Based on the exit code and the action you associate with the code, OpenEdge Management can ensure that the specified action occurs.

Typically, an exit code of zero indicates success, while a nonzero code indicates an error. For more information on nonzero exit codes, search the log file. If a job running in Windows returns a positive, nonzero code, use the `net helpmsg` command for information.

For example, consider a job that returns an exit code of 1326. Entering `net helpmsg 1326` at the command line returns the following message:

```
Logon failure: unknown user name or bad password
```

You can also use exit codes to chain jobs together using flow control. See Job chaining on page 114 for more information.

Job completion alerts

A job alert definition instructs OpenEdge Management to trigger a warning alert if the job execution time frame exceeds the time frame during which the process is scheduled to run. You can also associate an action with the alert.

To specify job completion actions and alerts:

1. Click the **Edit** button at the bottom of the **Job** page.

   The **Job Completion Action and Alerts** page appears.

2. Enter an exit code number in the **Exit Code** field.

   The number must be between 0 and 255. You can type the word `default`. From the **Action** drop-down list, select a predefined action to associate with the exit code. Click **Add/Update**. The exit code and action pair created display in parallel columns under their respective fields.

3. Repeat step 2 in this section to define additional exit code and action pairs. For example, you might have several exit code and action pairs.
4. In the **Alert if execution time exceeds** field of the **Job Completion Action and Alerts** page, enter the threshold of time after which you want a warning alert to trigger.

   You can specify this time in minutes or hours, and you must use whole numbers. For example, if you want to be notified at the 30-minute mark that the job execution has not occurred, enter the value `30` in the first field and click Minutes from the drop-down list.

5. From the **On alert perform** action drop-down list, select an action from the list of available, predefined actions that you want to associate with this alert.
6. Click **Save** to submit your selections.

   The **Job** summary page appears.

Updating job completion actions and alerts

You can update exit code and action pairs.

To update exit code and action pairs:

1. Enter the exit code you want to update in the **Exit Code** field.
2. From the **Action** drop-down list, select a different action to associate with the exit code.
3. Click **Add/Update**.

   The exit code and its new action are displayed in parallel columns under their respective fields.
4. Repeat steps 1 to 3 to update additional exit code and action pairs.
5. Once you have updated all necessary exit codes and action pairs, click **Save**.

   This will update pairs available for use the next time the job is executed. The **Job** summary page appears.

Deleting job completion actions and alerts

To delete exit code and action pairs:

1. Enter the exit code you want to delete in the **Exit Code** field.
2. Click **Delete**.

   The exit code and its action will be deleted from their respective columns.
3. Repeat steps 1 and 2 to delete additional exit code and action pairs.
4. Once you have deleted all the necessary exit code and action pairs, click **Save**.

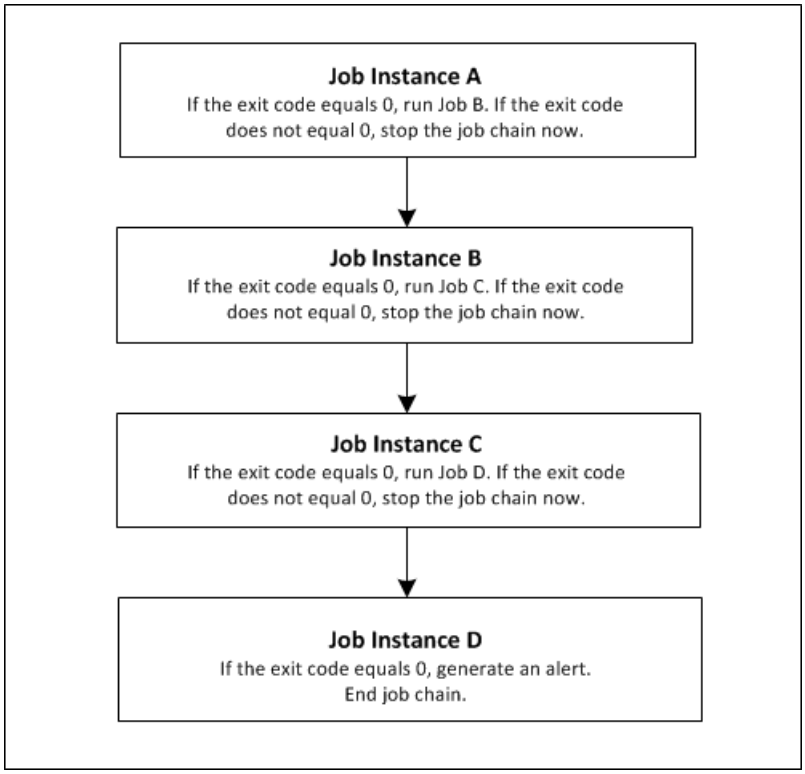
   The **Job** summary page appears.

Job chaining

When you link individual job instances together in a sequence based on the presence of a specific exit code, you are performing job chaining. Using the value of an exit code that occurs in response to a job having been run, you can determine the processes, or control the flow of processes, that occur once one job ends and the next one begins.

The following figure shows a simple job chaining scenario involving the following four job instances: Job Instance A, Job Instance B, Job Instance C, and Job Instance D. Each job instance has an exit code of 0 already defined. The job chaining in this scenario occurs only if the exit code of 0 is recognized at the end of job instances A, B, and C. If any of the jobs concludes with an exit code other than 0, the job chain ends at that point in the process.

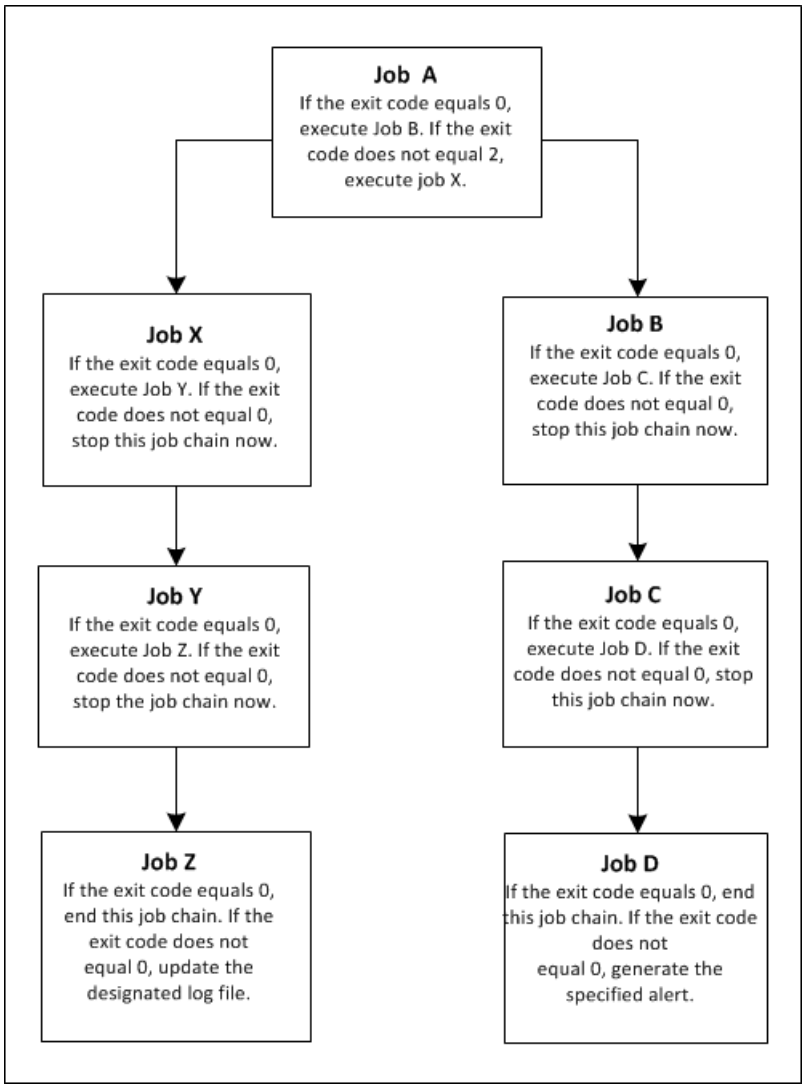
Figure 1: Sample job chaining scenario



If you want to account for more than one exit code condition, you can use a more sophisticated implementation of job chaining known as *job chaining with flow control*. You can set up one job chain to implement different job chains, depending on the exit code generated by a given job in the chain.

The following figure illustrates a sample job chaining with flow control scenario. One of two possible flows is determined based on the value of the exit codes.

Figure 2: Sample job chaining with flow control scenario



A word about recursion

You cannot create a job exit action that is self-referencing. OpenEdge Management prevents you from creating this type of situation in a job chain. However, OpenEdge Management will not detect the situation in which you create a self-reference further down a job chain.

It is possible for recursion to occur in an instance in a job chain when the job called is any other number of steps removed from the original job in the chain. For example, OpenEdge Management will not be able to detect a self-referencing job in a situation in which Job A references Job B and concludes with a reference back to Job A.

Recursion is not a potential issue in the example described in the Sample job chaining scenario figure above. Each of the job instances in the job chain is uniquely defined. Similarly, recursion is not an issue in the Sample job chaining with flow control scenario figure above because only one of the two possible chains in the job flow will be followed.

Using environment variables

Environment variables define the context in which a job or job action runs. OpenEdge Management sets up an environment variable context to associate jobs with resources and alerts and provides a set of predefined environment variables that are common to all jobs you can run.

There are additional, specific environment variables that will vary from one job to another, depending on the job to be executed. For example, jobs run in response to an alert have an alert context and jobs for which associated database resources have been defined have a resource context. OpenEdge Management provides a unique context definition for a job run against the OpenEdge Management Trend Database.

The following table identifies environment variables that are common to all jobs.

Table 13: Environment variables common to all jobs

Environment variable	Description
FM_HTTPPORT	OpenEdge Management's Web server port
ADMSRVRPORT	The AdminServer's listening port
JOBNAME	The name of the job as defined in OpenEdge Management
FATHOMINSTALLDIR	OpenEdge Management's installation directory as defined in `fathom.init.params`
FM_CONTAINERNAME	The resources's AdminServer name and the resource name (for example, Dev01.sports2004)
FATHOMCONFIGDIF	OpenEdge Management's directory for configuration data as defined in `fathom.init.params`
FATHOMLOGSDIR	OpenEdge Management's log directory as defined in `fathom.init.params`
FATHOMWORKDIR	OpenEdge Management's work directory as defined in `fathom.init.params`
FATHOMREPORTDIR	OpenEdge Management's directory for report data as defined in `fathom.init.params`
PROGRESSINSTALLDIR	Installation directory that OpenEdge Management is associated with as defined in `fathom.init.params`
PERL5BIN	Perl v5 interpreter bin directory as defined in `fathom.init.params`
PERL5LIB	Perl v5 library directory as defined in `fathom.init.params`

The following table identifies environment variables for jobs associated with database resources. If a job is associated with more than one database, the value of each of the variables will be identified in a comma-separated list. For example, if there are three databases associated with a job, and the databases are named `db1`, `db2`, and `db3`, the database names will be defined as follows:

```
RESRC_DBNAME=db1, db2, db3
```

Table 14: Environment variables for jobs with associated database resources

Environment variable	Description
RESRC_NAME	The name of the associated resource. This will be a database name because only a database resource can be associated with a job.
RESRC_DBNAME	The name of the database.
RESRC_DBPATH	The full pathname of the database. For example, `/usr1/myapp/sports.db`.
RESRC_DBCONFIG	The database configuration name.
RESRC_DBDISPLAY	The database display name.
RESRC_DBPORT	The TCP port number of the database broker.
RESRC_DBHOST	The TCP host name of the database DataServer.
RESRC_DBADMSRVRPORT	The AdminServer TCP port number on the database server host.

The **Action** option that appears at the bottom of the **Job** page. Select this option to make these environment variables available to the job's alert context.

The following tables identify the environment variables that can also be run in response to alerts. These environment variables are available when you identify jobs as job actions.

Table 15: Alert environment variables for e-mail actions

Environment variable	Description
ALERTID	The ID of the current alert.
ALERTNAME	The name of the current alert.
ALERTSEVERITY	The alert severity (`Severe`, `Error`, `Warning`, or `Informational`).
HOST	The host name on which OpenEdge Management is running.
CONTAINERNAME	The name of the container that is used to transport the resource.
RESOURCENAME	The name of the resource that has caused the alert to generate.
ALERTLASTOCCURRENCE	The date and time of the previous occurrence of the alert.

Environment variable	Description
ALERTREASON	A text description of the alert.
ALERTCOUNT	The number of times the alert has occurred.
ALERTRIGVALUE	The value that has triggered the alert.
HOSTLINK	The HTTP path of the host name.
RESOURCELINK	The full path of the resource that has generated the alert causing the e-mail to be sent.
ALERTLINK	The link to the description of the alert that has caused the e-mail to be sent.

Table 16: Alert environment variables for task actions

Environment variable	Description
TASKTYPE	The type of the current task (`Job` or `Alert`).
TASKNAME	The name of the current task.
TASKCOMMAND	The command required to complete the task.
TASKCOMMANDARGS	The command arguments required to complete the task.
TASKSTARTTIME	The start date and time of the current task.
TASKENDTIME	The end date and time of the current task.
TASKEXITCODE	The exit code defined for task completion.
TASKRESOURCES	The database resource that is used for the task to be performed.
TASKENVVARS	A list of environment variables set for the task type.
TASKCONTAINER	The name of the container that is used to transport the resource.
TASKHOSTLINK	The HTTP path of the host name.
TASKFAILMSG	If an error occurs while running the task action, an error message is displayed.

Note: If you set the **Default_Mail_Action** action for job completion alerts, the values of the job alert variables are not passed in the alert. Use the **Default_Task_Action** action to receive valid values of the job alert variables in the alert.

The following table describes the environment variables that you can identify for jobs with database resource associations.

Table 17: Environment variables for jobs associated with the OpenEdge Management Trend Database

Variable	Description
FM_DBNAME	This variable is set to the full name of the database. In this situation, it is OpenEdge Management Trend Database.
FM_DBPATH	This variable is set to the absolute pathname of the OpenEdge Management Trend Database.
FM_DBCONFIG	This variable is set to the OpenEdge Management Trend Database defined configuration.
FM_DBDISPLAY	This variable is set to the OpenEdge Management Trend Database defined display name.
FM_DBPORT	This variable is set to the OpenEdge Management Trend Database TCP port number.

Editing a job

You can edit a created job. If you edit the job name, be sure to change the output filenames (`stout` and `stderr`) and any other fields that retained the name of the original job. Review these file names **before** saving the edits to ensure that the job creates the expected output files.

You can also edit a job by changing the specified resources and AdminServer.

To edit a created job:

1. Select **Jobs** > **Job Scheduling** > **Defined Jobs** in the OpenEdge Management console menu.

   The **Jobs.Defined Jobs** page appears.
2. Select the job you want to edit.

   The job summary page appears.
3. Click **Edit**.
4. Make the changes you want, and click **Save**.

   The summary of the job reappears, reflecting the changes you made.

Copying a job

You can copy a created job and use it as the basis for another job. When you make a copy of an existing job, you must rename the newly created job. You should also change the updated `stout` and `stderr` filenames, and any other fields that retained the value of the original job from which you created this copy. Review these new file names **before** saving the page to ensure that the job creates the expected output files.

To copy a created job:

1. Select **Jobs** > **Job Scheduling** > **Defined Jobs** in the OpenEdge Management console menu.

   The **Jobs.Defined Jobs** page appears.
2. Select the job you want to copy.

   The job summary page appears.
3. Click **Copy**.
4. Rename the copy, and click **Save**.

   The summary of the job reappears.

Killing a job

You can kill a running job.

To initiate the kill process:

1. Select **Jobs** > **Job Scheduling** > **Running Jobs** in the OpenEdge Management console menu.

   The **Running Jobs as of: <date time>** page appears.
2. Click the **PID** associated with the job you want to terminate.
3. Click **Kill** to terminate the job. OpenEdge Management will prompt you once again to verify you want to terminate the job.
4. Click **OK**.

Deleting a job

You can delete a job you no longer need.

To delete a job:

1. Select **Jobs** > **Job Scheduling** > **Defined Jobs** in the OpenEdge Management console menu.

   The **Jobs.Defined Jobs** page appears.
2. Select the job you want to delete.

   The job summary page appears.
3. Click **Delete** and then click **Yes** in the confirmation dialog box to confirm deletion.

Running a job

You can run a new created job immediately.

Note: For existing jobs created prior to OpenEdge 11.1, you must update the start date and time to the current date and time else the management console might return an error when running the job. For updating the start date and time, see Scheduling a job.

To run a created job immediately:

1. Select **Jobs** > **Job Scheduling** > **Defined Jobs** in the OpenEdge Management console menu.

   The **Jobs.Defined Jobs** page appears.

2. Select the job you want to run.

   The job summary page appears.

3. Click **Run Now**.

   A window appears showing that the job has been scheduled to run now.

Scheduling a job

When you create a job instance, you also specify how often it will run. Jobs adhere to a number of execution policies: **Short Duration**, **Long Duration**, and **Single Time**. A full understanding of these policies allows you to schedule jobs more effectively.

Note: For existing jobs created prior to OpenEdge 11.1, you must update the start date and time to the current date and time before running the job.

Using cron-based scheduling

You can set up more complex scheduling by incorporating cron expressions in your job schedules. You can still specify that a job execute at a repeating calendar interval, such as every sixty minutes or once each day, for example. With cron-based scheduling, you can be even more specific about when you want a job to run, according to a business period interval, such as on the last Friday of the month, or every ten minutes from 4 PM to 6 PM daily.

If you choose to use a cron expression in scheduling a job, you select it from the **Repeat interval** drop down list.

You can include from one to five cron expressions (separated by semi-colons) in the **Cron expression** field of a job schedule. Each cron expression consists of the following five fields:

- **Minutes**
- **Hours**
- **Day of the month**
- **Month**
- **Day of the week**

To help you set up a cron expression, you can click **Assist** (once you select **Cron expression** from the **Repeat Interval** drop-down list).

A window with several different calendars appears. You can use this window to choose the month(s), date(s), day(s), etc., when you want the job to run.

When you finish making your choices, click **Apply**. Cron expressions that correspond to the choices you made appear in the **Cron expression** field.

Each field describes a different part of the scheduling definition. For example, the `15 15 ? * MON-FRI`cron expression specifies that a job should be run at 3:15 PM every weekday of every month.

If you also want the job to run once every Saturday and Sunday but at a different time from when it runs each weekday, you can use two cron expressions. For example, the `15 15 ? * MON-FRI; 30 19 ? * SAT,SUN`cron expression specifies that a job will run at 3:15 PM each weekday of every month, and the job will also run at 7:30 PM every Saturday and Sunday of every month.

For more details about using cron expressions, click the **Cron expression** field Help button.

Scheduling a job instance or a job template

You schedule a job according to whether you are scheduling a job instance or a job template, as follows:

- To associate a schedule with a job instance, complete all the steps in the following procedure.
- To associate a schedule with a job template, complete only step 1 through step 7 in the following procedure.

To schedule a job instance or a job template:

1. Select **Jobs** > **Job Scheduling** > **Defined Jobs** in the OpenEdge Management console menu.

   The **Jobs.Defined Jobs** page appears.
2. Select the job you want to schedule.

   The job summary page appears.
3. Click **Schedule**.

   The **Job Schedule** page appears.
4. Indicate the date on which you want the job to start, using the drop-down lists.

   Note that the **Start Date** is entered in a dd/mm/yyyy format.
5. Indicate the time you want your job to start, using the drop-down lists.
6. Choose the job's repeat interval from the drop-down list. The repeat interval options are:
   - **One time** — The job runs once on the specified start date and time.
   - **At startup** — The job runs once when the AdminServer and OpenEdge Management start. After running, the job is disabled until the next time the AdminServer restarts.
   - **Weekly from date** — The job runs once a week, beginning on the specified date.
   - **Monthly from date** — The job runs once a month, beginning on the specified date.
   - **Every 5, 15, 30, 60 Minutes** — The job runs every *n* minutes, beginning at the specified start date and time on the days chosen in the **Include days** list.

     Keep in mind that these increments of 5, 15, 30, and 60 offer a relatively short run time for a job. Since the length of time it takes a job to run can vary, a set time in your schedule might be missed.
   - **Daily** — The job runs once a day, beginning on the specified date.
   - **Cron Expression** — The job runs as specified by the cron expression.
7. Determine which days to select from the **Include days** list.

   The choices in the **Include days** list are valid and enabled only for repeat intervals of every 5, 15, 30, and 60 minutes and Daily. For all other repeat intervals, all days are active.

8. If you want, enter a cron expression in the **Cron expression** field. For assistance in creating the cron expression, click **Assist** and choose the job specifics from the available calendars. For help specific to cron expressions, click **Help**.
9. Select the **Enabled** option to enable the job's schedule.
10. Click **Save**.

Duration Execution Policy

When you define a job to run at intervals of Daily, Weekly, or Monthly, the Duration Execution Policy determines the time between runs. With the Duration Execution Policy, the Start Time is used to set the next run Date and Time. The Duration Execution Policy allows the time between runs to vary.

Example: Daily interval

Job execution begins at its scheduled time, 3:00 AM.

Job execution ends at 8:00 AM.

The next scheduled job execution begins at 3:00 AM on the next day this job is scheduled to run.

Single Execution Policy

When you define a job to run **At Startup** or **One Time**, the Single Execution Policy determines the time between runs. If a job runs at startup, it executes when the AdminServer starts (through the proadsv -start command on UNIX, or Services in Windows). With the Single Execution Policy, the job does not run again until the next time the AdminServer starts. If the job is defined as **One Time**, the Single Execution Policy ensures that the job, after completion, is permanently disabled. **One Time** jobs are reenabled only if you define new start times and **Repeat Intervals** for them on the **Schedule** page.

At startup

Job execution begins with the starting of the AdminServer at 7:30 AM.

Job execution ends at 7:45 AM.

The job does not begin again until the AdminServer is restarted.

One time

Job execution begins at its scheduled time, 6:15 PM.

Job execution ends at 6:53 PM.

The job is disabled. It will not run again unless its starting date and time are redefined on the **Schedule** page.

Job Execution Policy over time

Under typical conditions a job's **Start Date** and **Time** will reflect a point in the past. If the AdminServer and OpenEdge Management are shut down and restarted, OpenEdge Management determines the next execution date of any enabled jobs from the original **Start Date** and the time the AdminServer restarted (which would be the present time). OpenEdge Management then uses the **Repeat Interval** and the original **Start Time** to calculate the job's next occurrence.

Job execution with a restarted AdminServer

Original Job's Start Date and Start Time: July 21, 2004, at 10:00 AM.

Original Job's Repeat Interval: Every 60 Minutes as determined from the original Start Time indicated as 10:00 AM.

Date of AdminServer shutdown and restart: July 28, 2004.

Time of AdminServer restart: 8:25 PM.

The next scheduled execution of this job begins at 9:00 PM on July 28, 2004.

If the job in the above example had a Repeat Interval of Every 15 Minutes, it would next execute at 8:30 PM on July 28, 2004.

Monthly By Date Execution Policy

Tasks scheduled to run monthly on the thirtieth or thirty-first of the month run on the last day of the month, regardless of the month.

Tasks in April

A report of the Inventory database is scheduled to run on the thirty-first of the month. Since there are only thirty days in April, OpenEdge Management executes the report on April 30.

Tasks in February

The same report used in the above example would execute on February 28 in nonleap years and on February 29 in leap years.

Viewing scheduled jobs

To view jobs scheduled to run, choose one in the OpenEdge Management console menu:

- Select **Jobs** > **Job Scheduling** > **Scheduled Jobs**.
- Select **Jobs** > **Go to Jobs** > **View Scheduled Jobs**

The **Scheduled Jobs** page appears and the following information is given for each scheduled job:

- The name of the job
- Whether the scheduler is enabled

- The next scheduled time the job will run
- The job AdminServer, if the AdminServer is remote. If the AdminServer is local, the AdminServer name is blank.
- The command to be executed when the job runs
- Any parameters defined for the command
- The working directory

To see a job's summary page, click the specific job name.

Viewing running jobs

To view a list of jobs currently running in the OpenEdge Management, choose one:

- Select **Jobs** > **Job Scheduling** > **Running Jobs**.
- Select **Jobs** > **Go to Jobs** > **View Running Jobs**.

The **Running Jobs as of: <data time>** page appears.

OpenEdge Management displays the following information about each running job:

- The **Name** of the currently running job
- The **PID** (process ID) of the running job
- The **Start Time** of the job
- The **Command** executed when the job runs
- The job **AdminServer**
- Any **Parameters** defined for the command
- The **Working Directory**

To see a job's summary, click the specific job name.

Viewing job history

The **Job History** section of the **Job** summary page allows you to:

- Display historical details about individual instances that are run for a specific job within a date range that you define
- Display output, error, and debug trace details for the specific job instance

To view data for all jobs that have completed running in a specified time frame:

1. Choose one:
   - Select **Jobs** > **Job Scheduling** > **Jobs History**.
   - Select **Jobs** > **Go to Jobs** > **View Job History**.

The **Job History** page appears.

2. Specify the date range for the job history query you intend to submit, using the drop-down lists. Note that the **From** and **To** values of the date range are entered in a dd/mm/yyyy format.

   The date range default value displays the past seven days. The seven days are determined as inclusive of the current date through the seventh day as you count forward in a given week.

3. Click **Submit** to submit your query. If no job history exists for the time frame you submitted, a **No history records available** statement displays.

   However, if a job history does exist, the history appears.

4. Click an underlined name in the **Name** column to see a job status that includes:

   - The name of the completed job
   - The start time (includes the start date)
   - The end time (includes the end date)
   - The command executed when the job ran
   - The exit code

   When a job completes, OpenEdge Management stores its data in the OpenEdge Management Trend Database. OpenEdge Management pulls the information about a completed job from the database table `Cf_Task_Detail`. For more information about OpenEdge Management Trend Database tables, see *OpenEdge Management: Trend Database Guide and Reference*.

Viewing database maintenance job instances for an individual database

OpenEdge Management supports functionality to view job instances currently associated with a specific database.

To view job instances associated with databases:

1. Click **Resources** > **Go to Resources** in the OpenEdge Management console menu.

   All resources managed by your console appear in the grid frame.

2. Filter or search for, and select the database whose job instances you want to view.

   The **Database: <database name>** page appears.

3. Click **Maintenance Jobs** in the **Command and control** section.

   The **Database Maintenance Jobs** page appears and displays the list of jobs that are associated with the particular database and each job's next scheduled run time. Jobs that do not display information in the **Next Scheduled Run** column do not have schedules that are enabled.

4. From this page you can perform the following tasks:

   - Click a job's name to display its **Job** summary page.
   - Click **Create** to go to the **OpenEdge Management Jobs** page.

Note: Complete functionality associated with the **Database Control** page is discussed in *OpenEdge Management: Database Management*.

Working with job templates

OpenEdge Management makes it possible to create and reuse job templates. A *template* is a sample outline for a resource category that you create. It usually includes some instance data. Values you define in a job template can be reused to help create individual job instances.Job templates are also a time-saving tool to help you standardize the creation of job instances and ensure that they are exportable.

In situations where job instances require just a few unique values, you can define a job template from which you can create individual jobs. To create a job template, you define common values for the template. Then, when you create a job instance from the template, you supply the unique values for the job, accepting appropriate default values from the template. This approach minimizes your data entry tasks and data entry errors.

These common values typically include:

- Commands such as executables or scripts
- Command arguments
- Job properties
- A default schedule

When you are ready to create instances from the template, you can access and reuse these values in job instances derived from the template.

For example, in analyzing your current OpenEdge database maintenance practices, your company determines that productivity in your IT department could be greatly improved if database restructuring activities were performed on a regularly scheduled basis for all databases. To help standardize these processes, and ensure that they are performed routinely for all databases, you can create a user-defined job template and associated schedule for each of these unique database activities.

Using OpenEdge Management-supplied generic tools such as environment variables, you can tailor your job template to your specific needs. When you are ready to use this template, you associate the appropriate database resources with a job instance derived from the specific template.

Note: When reviewing your current database maintenance practices, keep in mind the availability of the OpenEdge Management-supplied, specialized Database maintenance job templates from which you can create job instances to address fundamental OpenEdge database maintenance activities. Designing user-defined job templates can be extremely effective, but they can also require additional analysis and design time to effectively produce. In contrast, the Database maintenance job templates are already tailored to address specific database tasks. Consider all your template options as you analyze your database needs. It is possible that a combination of user-defined and OpenEdge Management-supplied Database maintenance job templates might be optimum for your company.

For more information about these database maintenance jobs, see *OpenEdge Management: Database Management*.

Keep in mind that templates can be imported or exported. For example, users at your company's satellite offices can use the identical templates that are employed at your corporate office, ensuring that your practices are standardized throughout your company. For details about import and export activities, see Exporting and Importing in OpenEdge Management on page 133.

Creating a job template

When you create a job template, you can choose the job's characteristics. You need not recreate the job template criteria each time you want to create another job instance based on the job template.

Once you create a job instance from a job template, the two are no longer tied together. You can modify the characteristics of a job instance without affecting the template on which the job is based, and you can update the template without affecting the characteristics of a job instance you have already created.

You choose the following characteristics for the job template: **Menu group**, **Menu entry**, and **Menu description**. Once you create the template, these characteristics enable you to identify it from among all other existing templates. For example, consider a scenario in which a user creates a Database Restore template for use on company databases related to inventory. The user makes the following choices regarding menus:

Menu group	Database Maintenance
Menu entry	DB Restore Inventory
Menu description	DBrestore for all inventory dbs

Once the job template is created, the template appears.

The new job template **New** menu entry appears in the **Database Maintenance** menu group with the **New job template** menu description.

You specify each OpenEdge Management job template's properties and characteristics on two pages. The first template page defines the properties of the job. On the second page, you provide default values for the job instance; however, you can change these values, such as the schedule, for each instance.

To create a job template:

1. In the OpenEdge Management console, choose one:
   - Select **Jobs** > **New** > **Job Template**.
   - Select **Jobs** > **Go to Jobs** > **Create Job Template**.

   The first of the two **Job Template** pages appears.
2. In the **Name** field, enter the job template name. Note that the name must be unique among the job templates and that spaces are not allowed in the name fields on this page.
3. In the **Menu group** field, enter the name of the new menu group, or choose an existing menu group in the **Existing** field. (When you choose an existing group, the name automatically appears in the **Menu group** field.)

   The menu group is the heading under which the new template will appear (for example, **Database Maintenance**).
4. In the **Menu entry** field, enter the template name you want to appear on the **Custom Job** page.
5. In the **Menu description** field, enter a brief description of the job template. This description appears in the list of available job templates on the **Create Custom Job** page, under the template name you provide in step 4.
6. Under the **Limit how many databases can be associated with the job** field, select **one**, **two**, or **no limit of databases** from the drop-down menu. **No limit** is the default value.
7. Click **Save**.

   The second of the two **Job Template** pages appears with the **Name** field predefined.

8. Provide the properties.

   - Entering the **User name** in the **Account information** section is optional; if you do not supply a user name, any job instance created based on this template is run under the account OpenEdge Management is running under. Entering a password is required only if you entered a user name. Otherwise, a password is optional.
   - You have three AdminServer options from which to choose. The choice affects whether you will need to select an AdminServer when you create a job using the template. Choose one of these options:
     - **Prompt for AdminServer** — You will need to provide the name of the AdminServer when you create a job from the template.
     - **Use AdminServer of resource** — You will not need to provide the name of the AdminServer when you create a job from the template.
     - **Use specified AdminServer** — You will not need to provide the name of the AdminServer when you create a job from the template.

   For details about any other field, see Creating a job instance on page 110.

9. Click **Save**. The job template summary appears.

   You can now edit or copy the template's characteristics, delete the template, or schedule a job instance to run based on the template.

10. After you specify the properties that you want to default to each instance created from this template, you can:

    - Click **Edit** at the bottom of the page directly to the right of the **Completion Actions and Alerts** field to continue entering values. The **Job Completion Actions and Alerts** page appears. This page allows you to define default action and alert occurrences for each job instance created from this template. See Setting up job completion actions and alerts on page 112 for details.
    - Click **Save** at the top of the page to conclude entering values for the job template. The **Job** summary page appears.
    - If you want to set additional values for job actions and/or alerts, you can access the **Job Completion Actions and Alerts** page and perform these tasks at a later time by editing the **Job** page.
    - You can also initiate the job immediately by clicking **Run Now** on the **Job** summary page.

Additional characteristics of job templates

Note the following points as you create and work with custom job templates:

- You do not run a job template. You run a job instance derived from a job template.
- To access user-defined templates and the OpenEdge Management-supplied database maintenance templates, select **Jobs** from the menu bar. Review the contents of the **Defined Jobs** and **Job Templates** categories that appear in the list frame.
- When you create a job template, you can define schedule information, but you cannot enable the schedule to run with the template. This schedule definition can be enabled once you create an instance from the template.

  You can associate one or more command actions with job templates.

- You cannot define specific database resources for a job template. The job template is an outline of common values that you can copy and apply to various instances derived from the template. You associate specific database resources when you create a job instance.
- Once you have created a job instance from a job template, these items are completely separate from each other. Edits to the job instance do not affect the job template, and edits to the job template do not affect instances previously derived from the template.
- Unlike job instances, which can only be used local to the machine on which they are created, job templates can be exported and imported. However, the definition of job templates that you export or import does not include resource- or device-specific properties.
- You can also access existing job templates from the **Jobs** home page. See Accessing existing job templates on page 131 for more details.

Choosing the job template schedule

When you create a job template, you also determine the schedule you want job instances created from the template to use.

Setting up a schedule for a template is very similar to setting up a schedule for a job instance. However, you cannot enable a schedule at the template level. To create schedule data, complete the **Job Schedule** page. Then see Scheduling a job on page 122 for more information.

Accessing existing job templates

To access existing job templates, choose one:

- Select **Jobs** > **New** > **Job from a Template**.
- Select **Jobs** > **Go to Jobs** > **Create Job from a Template** > **.**

The **Create Custom Job** page appears and displays the following OpenEdge Management-supplied menu groups:

- **Database Maintenance**
- **OpenEdge Management Configuration**
- **OpenEdge Management Trend Database Maintenance**

Any templates that you create are also listed here, according to the menu groups, descriptions, and entities that you define for these templates.

Viewing debug details about jobs

OpenEdge Management provides a diagnostic report that includes debug information about jobs and reports. For details about the **Task Detail** report, see *OpenEdge Management: Reporting*.

9

Exporting and Importing in OpenEdge Management

When you create a resource monitoring plan, you can use OpenEdge Management to export and import selected plan components from one machine running OpenEdge Management to another.

For details, see the following topics:

- Exporting and importing overview
- Exporting OpenEdge Management Library components
- Importing OpenEdge Management Library components

Exporting and importing overview

As you work with OpenEdge Management, you may find that you want to share certain related components that you created on one machine with one or more other machines as well. You can export or import any of the sharable OpenEdge Management components, as follows:

- Actions
- Report templates
- Job templates
- Rule sets for databases, log files, or OpenEdge server components
- Search criteria for log files

General export and import considerations

Keep the following points in mind before you perform either an export or an import activity:

- It is advisable to back up the OpenEdge Management configuration database **before** you proceed with the import process.
- You will be informed of any potential conflicts through error messages in the OpenEdge Management console and in the log file. For example, if a component you are exporting from Machine A to Machine B has a schedule that conflicts with the schedule of an existing component on Machine B, you will be notified of this conflict.

Exporting OpenEdge Management Library components

When you *export* a sharable OpenEdge Management component from one machine, you place a copy of the OpenEdge Management component's definition into a file that you can then import onto another machine and use. When you export sharable Library components, the export process automatically determines any related resource dependencies and exports these components, too. For example, if you select a job template to export and the job template references an action to be performed when the job completes, OpenEdge Management will automatically include that action in the exported file.

Note: You can perform only one export or import operation at any one time during a given HTTP session.

The OpenEdge Management components you can export are contained in the OpenEdge Management Library.

To perform the export process:

1. Click **Library** > **Go to Library** in the OpenEdge Management console menu.

   The **OpenEdge Management Component Library** page appears.

2. In the **Import and Export Components** section, click **Export Components**.

   The **OpenEdge Management Export Step 1: Select Resources for Export** page appears.

3. Select the component you intend to export.

   The resources available for the export list appears.

   You can click **Select All** to choose every one listed. To deselect all search criteria, click **Unselect All**.

4. Click **OK** to return to the **OpenEdge Management Export Step 1: Select Resources for Export** page.

   The list of components selected for export appears.

5. Click **Export** when you have selected all components you want to export at this time. The **OpenEdge Management Export Step 2: Specify Export File** page appears.

   In the **Export to file** field, OpenEdge Management provides a default path and filename for the export file to be created. You can change either default value in that field.

   You can also indicate what action you want taken if the file already exists. The options available are:

   - **Display an error** (default option).
   - **Overwrite existing file** — Removes any data currently in the file.

- **Update/add components to existing file** — Allows you to add the components you just selected to the list of components already in the file. If you have selected for export a component that already exists in the `.xml` file, the old definition will be removed and the selected component added.

6. Choose one of the following:

   - **Back** — Returns you to the **OpenEdge Management Export Step 1: Select Resources for Export** page.
   - **Finish** — Initiates the export operation. The **Exporting File** page appears, noting the path and filename of the exported file. The **Exporting File** page is updated every three seconds. In general, the larger the number of components you are exporting in the `export.xml` file, the longer the operation can take.
   - **Cancel** — Ends the export setup process.

   Once the export operation is completed, the **Export Summary Report** page appears.

   The header information on this page identifies the path and filename to which the exported components were sent. The date and time information identifies when the export operation occurred. The report details on the **Export Summary Report** page provide the list of components that were exported, updated, or determined to be dependent because they are referenced by an exported resource.

   This process also determines if any errors occur that prevent the export from succeeding. If an error occurs, the error will appear in a separate window, and the report page will provide detailed information about the error.

7. Click **Clear Report** once you have reviewed and/or printed your exported resources list. You must select **Clear Report** before you can initiate another export operation.

Exporting referenced items

Components that you export might be dependent on other components to function as intended. For example, if you want to export a job resource and the job resource references an action to be performed when the job completes, OpenEdge Management automatically includes that action in the exported file.

An *explicit* component is one that you select. An *implicit* component is one that is also exported because a selected component has a dependency on it. On the **OpenEdge Management Step 1: Select Resources for Export** page, the implicit components under **Components selected for export** look disabled to identify their dependency. The **Exported components** list differentiates the explicit and implicit components in a similar manner.

Importing OpenEdge Management Library components

When you import a component, you add the component definition from an import file to your project.

To perform the import process:

1. Click **Library** > **Go to Library** in the OpenEdge Management console menu.

   The **OpenEdge Management Component Library** page appears.

2. In the **Import and Export Components** section, click **Import Components**.

   The **OpenEdge Management Import** page appears.

3. If you want, you can change the default value in the **Import from file** field. (The default path and filename are the same as the default values used for the most recent export operation.)

4. Indicate the action you want taken if any of the components that you are importing already exist. You can choose to:

   - **Display an error** — Indicates that none of the components is imported if any component in the import file conflicts with an existing component. (This is the default.)
   - **Keep existing resource** — Indicates that the existing definition of a given component is maintained and the imported definition discarded.
   - **Replace existing resource** — Indicates that the existing definition is deleted and the imported definition is added.

5. Choose one:

   - **Import** — Initiates the import operation. The **Importing File** page displays and is updated every three seconds. In general, the larger the number of components you are importing from the `export.xml` file, the longer the operation can take.
   - **Cancel** — Terminates the current import selections.

   Once the import operation is completed, the **Import Summary Report** page appears.

   The report data on the **Import Summary Report** page indicates which components have been updated. This update occurs when you select the **Replace existing resource** option on the **OpenEdge Management Import** page.

   Note: Unlike with the export operation, there are no explicit and implicit components associated with the import operation. No components appear dimmed in the import operation's report page.

   If you select the **Display an error** option on the **OpenEdge Management Import** page and then a conflict is encountered, you are notified that the import failed.

Index

A

Actions
  creating
    compound *57*
    e-mail *54*
    log *56*
    SNMP Trap *57*
  default mail *54*
  definition *23*, *54*
  updating and deleting *59*
AdminServer
  definition *22*
Alerts
  definition *23*, *53*

C

Collection views
  creating *38*
  customizing layout for *38*
  definition *36*
  deleting *39*
  differences from collections *30*
  editing *39*
  working with *37*
Collections
  and administrators or operators *33*
  creating *33*
  definition *22*, *31*
  deleting *34*
  differences from collection views *30*
  editing *34*
  My Collections *31*
  referenceable collections *31*
  sample scenario *32*
  shared collections *31*
  terminology and hierarchy *31*
  working with *33*
Configuration Advisor
  comparing and selecting threshold settings *52*
  determining effectiveness of selections *53*
  guidelines *50*
  overview *50*
  reviewing calculation details *52*
  reviewing threshold settings recommended by *52*
  submitting threshold setting selections *53*
Copying resource monitors *59*
Creating resource monitors
  Monitoring Plan and Rule Definition page *44*
  resource naming conventions *43*
Cron-based scheduling
  for jobs *122*
Custom views
  customizing the layout of *38*
  Dashboard Home page *30*
  displaying custom viewlets in *37*

D

Dashboard page *30*
Database maintenance job templates *109*
Debug file
  reviewing for job details *131*
Default action (Default_Action) *54*
Default mail action (Default_Mail_Action) *54*
Default schedule *106*
Default values
  changing at the individual resource level *62*
  CPU and memory resources *43*
  definition *60*
  file monitor resources *87*
  global settings *60*
  resource type *61*
Default_Action *54*
Deleting resource monitors *60*
Displaying custom viewlets
  in custom views *37*

E

Edit default monitoring plan *105*
Edit default schedule *106*
Editing resource monitors *59*
Environment variables
  job *117*
Exit code
  for jobs *113*
Exporting sharable OpenEdge Management components
  considerations before *134*
  details *134*
  explicit *135*
  implicit *135*
  that include referenced items *135*

F

File monitors
  defining *101*
  rules *101*

G

Global settings
reviewing and changing 60

H

HTTP resource monitor
considerations 80
creating 80
using to determine Web site status 79

I

Importing sharable OpenEdge Management components
considerations before 134

J

Job Action field 117
Job chaining
recursion 116
Job instances
associated with an individual database 127
common environment variables in 117
definition 108
Job Action field in 117
job chaining in 114
viewing when associated with an individual database 127
Job templates
accessing existing 131
creating 129
definition 62
for database maintenance 109
job chaining 114
menu group, menu entry, menu description 129
setting up
job completion actions and alerts for 112
schedules for 131
Jobs
and job chaining 114
creating
as an instance 110
templates for 129
custom 108
database maintenance templates for 109
definition 27
environment variables in 117
execution policies
monthly by date 125
over time 125
single 124
Job Action field and 120
Jobs *(continued)*
job completion actions and alerts for 112
overview 110
recursion 116
reviewing debug file for details about 131
scheduling
using cron expressions 122
templates
accessing existing 131
additional characteristics 130
schedule data in 131
setting up 128
viewing
history 126
running 126

L

Library functions
defining search criteria 63
definition 63
schedules and actions 63
Log file monitors
accessing library-defined search criteria for 95
creating library-based search criteria for 96
definition 86
enabling 100
library features
accessing rule sets 92
associating rule sets with log file monitors 94
overview 85
rule components 91
setting up
library-based log file rule sets for 92
monitoring plans 90
properties 88
rules 91
Truncate Action 90
viewing summary details 101

M

Management console
default user name and password 24
Menu description
job template 129
Menu entry
job template 129
Menu group
job template 129
Monitor
definition 23
Monitoring plan 105
Monitoring remote resources 42
My Collections 31

N

Naming conventions for resources *42*
Network resources
  adding *77*
  determining machine and port availability for *76*
  discovery *75*
  monitoring *73*

O

OpenEdge Management Trend Database
  definition *24*

P

Perl 5 regular expressions
  as search criteria *97*
PING/ICMP monitors
  creating *79*
  defining *79*

R

Realtime reports *26*
Redirection
  Web page *80*
Referenceable collections *31*
Remote Configuration Utility *42*
Report templates
  definition *62*
Resource monitoring
  overview *43*
Resource monitoring plans
  definition *23*
Resource monitors
  deleting *60*
  editing and copying *59*
Resource viewlets, in custom views *37*
Resources
  definition *21*
  naming conventions *42*
Rules
  definition *23*
Rules definitions
  default *48*
  for CPU usage *67*
  for file systems *70*
  for files *101*
  for memory usage *68*
  for system disk devices *69*
  resource-specific *48*

S

Schedules
  considerations *46*
  default *46*
  definition *23*, *45*
  predefined by OpenEdge Management *46*
  updating and deleting *47*
Search criteria
  definition *95*
  for log file monitors *95*
  literal strings as *95*
  Perl 5 regular expressions in *95*, *97*
Shared collections *31*
SNMP Trap Action
  creating *57*
Standard viewlets
  in custom views *37*
System disk devices
  monitoring *69*

T

TCP monitors
  creating *78*
Terminology, OpenEdge Management
  actions *54*
  AdminServer *22*
  alerts *53*
  alerts and actions *23*
  collections *31*
  common *21*, *24*
  monitor *23*
  resource-monitoring plan *23*
  rule set *23*
  rules *23*
  rules or rules definitions *47*
  schedule *23*
  schedules for jobs *122*
  Trend Database *24*
Trending considerations for
  CPU resources *68*
  disk resources *70*
  file systems *71*
  memory resources *69*
  network resources *74*
Truncate Action *90*

U

UDP monitors
  creating *78*

V

Viewing
job history *126*
running jobs *126*
scheduled jobs *125*

Viewlets
characteristics and functionality *37*
customizing *39*
definition *36*
resource *37*
standard *37*
view panel *37*