webMethods Monitor User's Guide

Version 10.15

October 2022

WEBMETHODS

Document ID: MON-UG-1015-20221015

Table of Contents

About this Guide

This guide includes information about how to setup and use webMethods Monitor to monitor business processes, services, and documents; how to work with process models that are in your production environment; and how to archive audit data from the IS Core Audit Log and Process Audit Log database components. Access the webMethods Monitor functions described in this guide using the My webMethods user interface.

Deprecation of webMethods Broker

webMethods Broker is deprecated for use with webMethods 10.2. If you are starting development using webMethods 10.2, you should use webMethods Universal Messaging instead of webMethods Broker. If you are upgrading to webMethods 10.2, you should consider migrating to Universal Messaging. If you choose to continue to use webMethods Broker, you will still be fully supported, but only until the announced end-of-life dates for webMethods Broker. For details, see https://empower.softwareag.com/brokerendoflife/

Document Conventions

Convention	Description
Bold	Identifies elements on a screen.
Narrowfont	Identifies service names and locations in the format *folder.subfolder.service*, APIs, Java classes, methods, properties.
Italic	Identifies: Variables for which you must supply values specific to your own situation or environment. New terms the first time they occur in the text. References to other documentation sources.
`Monospace font`	Identifies: Text you must type in. Messages displayed by the system. Program code.
{ }	Indicates a set of choices from which you must choose one. Type only the information inside the curly braces. Do not type the { } symbols.
|	Separates two mutually exclusive choices in a syntax line. Type one of these choices. Do not type the | symbol.
[ ]	Indicates one or more options. Type only the information inside the square brackets. Do not type the [ ] symbols.
...	Indicates that you can type multiple options of the same type. Type only the information. Do not type the ellipsis (...).

Online Information and Support

Product Documentation

You can find the product documentation on our documentation website at https://documentation.softwareag.com.

In addition, you can also access the cloud product documentation via https://www.softwareag.cloud. Navigate to the desired product and then, depending on your solution, go to “Developer Center”, “User Center” or “Documentation”.

Product Training

You can find helpful product training material on our Learning Portal at https://knowledge.softwareag.com.

Tech Community

You can collaborate with Software AG experts on our Tech Community website at https://techcommunity.softwareag.com. From here you can, for example:

- Browse through our vast knowledge base.
- Ask questions and find answers in our discussion forums.
- Get the latest Software AG news and announcements.
- Explore our communities.
- Go to our public GitHub and Docker repositories at https://github.com/softwareag and https://hub.docker.com/publishers/softwareag and discover additional Software AG resources.

Product Support

Support for Software AG products is provided to licensed customers via our Empower Portal at https://empower.softwareag.com. Many services on this portal require that you have an account. If you do not yet have one, you can request it at https://empower.softwareag.com/register. Once you have an account, you can, for example:

- Download products, updates and fixes.
- Search the Knowledge Center for technical information and tips.
- Subscribe to early warnings and critical alerts.
- Open and update support incidents.
- Add product feature requests.

Data Protection

Software AG products provide functionality with respect to processing of personal data according to the EU General Data Protection Regulation (GDPR). Where applicable, appropriate steps are documented in the respective administration documentation.

1 Concepts

About webMethods Monitor

webMethods Monitor consists of two primary components:

- The WmMonitor package that is installed on webMethods Integration Server.
- The Monitor User Interface that is installed in My webMethods Server.

Both of these components must be installed to enable webMethods Monitor operation. To work with Monitor, you log in to My webMethods Server and work with the Monitor user interface available there. For more information, see "About the Monitor User Interface" on page 15.

webMethods Monitor enables you to view a wide variety of information logged by webMethods Integration Server and webMethods Optimize for Process for business processes, services, and documents in your webMethods environment.

webMethods Monitor operates with the following process types:

- *webMethods -executed processes*. These are business processes orchestrated by a Process Engine and executed on a an Integration Server
- *Externally executed processes*. These are processes that were executed by an external application and not orchestrated by a Process Engine or executed on an Integration Server.
- *Integration processes*. These are services that run on Integration Server and invoke each other in a sequence.

In addition to displaying logged information, Monitor enables you to:

- Resubmit documents, and edit them prior to resubmission.
- Resubmit services and processes (if you log pipeline data for those services and processes).
- Create and log audit data for integration processes.
- Archive your logged data. You can delete data instead of archiving it, and you can delete archived data as well.

Architecture and Workflow

The diagram below shows the workflow and architecture for webMethods Monitor:

- webMethods Broker (deprecated) and Software AG Universal Messaging pass documents to Integration Server using the Logging Utility package.
- Integration Servers log service data and documents to the IS Core Audit Log database.
- Process Engines log data for webMethods-executed processes to the Process Audit Log database.
- Optimize Analytic Engines log data for externally executed processes to the Process Tracker database.
- Monitor retrieves the logged data and documents from all three databases.

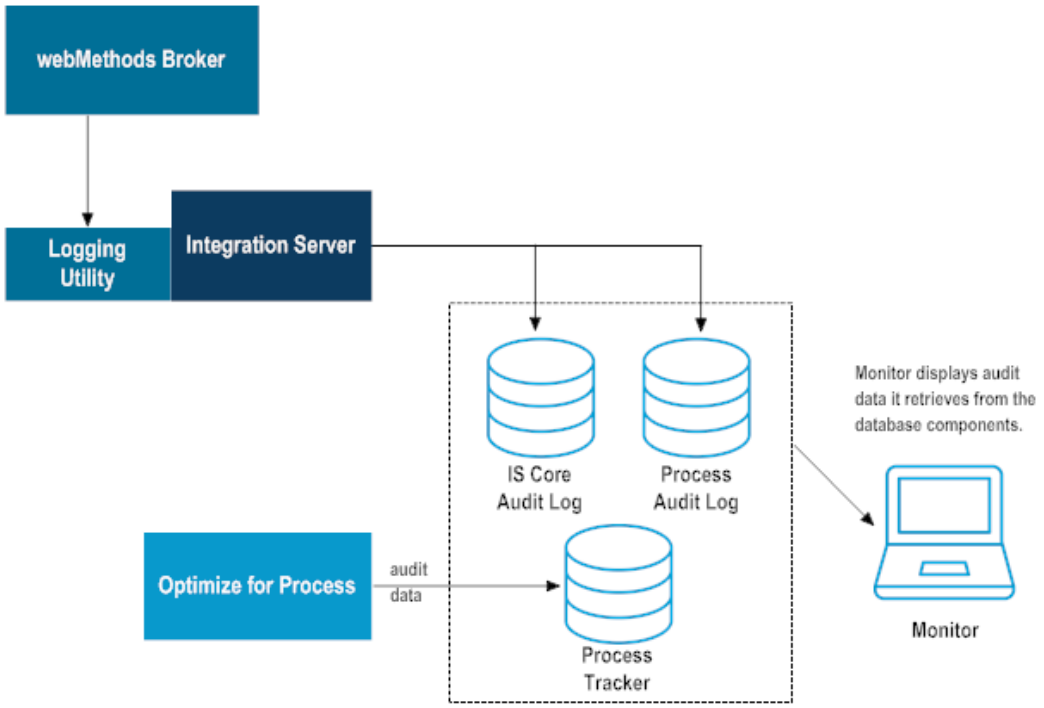

Data You Can Monitor

webMethods products log the following types of data for monitoring with webMethods Monitor:

- Audit data for flow and coded (for example, Java) services.
- Audit data for Integration Server and webMethods Broker (deprecated) documents.
- Audit data for webMethods-executed process instances.
- Audit data for externally executed processes.

For detailed information about Integration Server logging, see the *webMethods Audit Logging Guide*.

Optimize Analytic Engines log business and process audit data for externally executed business processes. For detailed information, see *Administering webMethods Optimize*.

Monitoring Flow and Coded Services

In Monitor, you can view the following data logged by Integration Server:

- When a service starts.
- Service status and duration.
- Whether the service completed successfully or failed.
- The client that called the service.
- The pipeline data from the service.

- The Integration Server port on which the client connected.

Monitoring Documents

In Monitor, you can view data logged by Integration Server for these types of documents:

- Integration Server documents that are in doubt, that have failed, or that have exhausted trigger retries (see the *Publish-Subscribe Developer's Guide*).
- Documents that webMethods Broker (deprecated) clients publish or subscribe to.

Monitoring webMethods-Executed Process Instances

In Monitor, you can perform the following tasks for webMethods-executed process instances.

- Identify process instances.
- See the path process instances took at run time.
- See the estimated time of completion of a process instance.
- Track when process instances and process steps started and when they ended.
- Track changes in the status of process instances and steps.
- Track whether process instances and steps completed successfully or failed.
- Track stages and milestones defined in a business process instance.
- See values of fields and custom data that were logged for steps.
- See error messages for a process instance.
- See control actions (such as resubmit) taken for a process instance.

Monitoring Externally Executed Processes

In Monitor, you can perform the same tasks for externally executed processes that you can for webMethods-executed processes, except you cannot perform control actions (such as resubmit).

Monitoring Integration Processes

In Monitor, you can perform the following tasks using data logged for integration processes.

- Identify process instances.
- Track the status of process instances and their steps.
- See values of fields and custom data logged for steps.
- See error messages for a process instance.

- If you are analyzing process instances in Optimize for Process, you can use Monitor to view metrics that relate to your business processes, such as the average time to complete a process or the number of times that a step was executed.

Monitoring Dynamically-Invoked Processes

Monitor supports monitoring of processes that are dynamically-invoked by implementing a call activity step.

For more information on dynamically-invoked processes and call activity concepts, see *webMethods BPM Process Development Help*.

About the Monitor User Interface

You access Monitor functionality through the Monitor user interface in My webMethods. The user privileges assigned to your My webMethods account control what you can access in Monitor. If a procedure in this guide instructs you to use an item or page that is not available to you, see your system administrator about acquiring the necessary privileges. For more information about My webMethods, see the PDF publication, *Working with My webMethods*.

Note:
To enable Monitor to perform actions on Integration Server using existing My webMethods user accounts, you must implement the central user management feature. For more information, see "Configuring Central User Management" on page 20.

Monitor Administration Tasks

In My webMethods, browse to **Navigate > Applications > Administration** to access the administration tasks. Unless noted otherwise, instructions for the tasks are provided in this guide.

The following table lists these tasks and the administration navigation path to each of them.

Task	Administration Navigation Path
Define users, groups, and roles, and configure access to Monitor functionality.	**System-Wide > User Management**
Enable process models for execution and define settings used by running process instances.	**Business > Business Processes**
Delete or archive and delete stored data.	**Business > Data Management > Archive Audit Data**
Identify the Integration Server that hosts the WmMonitor package to the My webMethods Server that hosts the Monitor user interface in My webMethods.	**My webMethods > System Settings**

Monitoring Tasks

In My webMethods, browse to **Navigate > Applications > Monitoring** to access the monitoring tasks.

The following table lists these tasks and the monitoring navigation path to each of them.

Task	Monitoring Navigation Path
■ View details about process instances and steps. ■ Suspend, resume, and stop process instances.	**Business > Process Instances**
■ View service audit data. ■ Resubmit services.	**Integration > Services**
■ View logged documents. ■ Resubmit documents.	**Integration > Documents**

Archiving Data

To keep logging at peak performance, Software AG recommends that you remove data from the IS Core Audit Log and the Process Audit Log databases on a regular basis. Monitor enables you to:

- Delete data directly from the databases without archiving it.
- Archive the data from the database to a separate archive database. You can then delete data from the archive database as needed.

Before you can archive data, you must configure the archiving procedure. For more information, see "Configuring Archive Settings" on page 30.

2 Configuring Monitor

Overview of the Monitor Configuration

This chapter covers mandatory configuration tasks for Monitor. Most configuration tasks are described in this chapter. Additional configuration tasks, as listed in the following table, are described in other chapters and other guides.

Task	See
Create the databases that store the data to be monitored.	*Installing Software AG Products*
Configure logging for services and documents.	*webMethods Audit Logging Guide*
Configure logging for custom fields in services.	*webMethods Service Development Help*
For webMethods-executed processes:	"About Process Model Data Logging" on page 93
■ Configure process instance audit logging.	"About Process Model Data Logging" on page 93
■ Specify process step input and output document fields to log as run-time values.	*webMethods BPM Process Development Help*
■ If you log process transitions so you can see the path the process took at runtime, set the model image format.	*Working with My webMethods*
Configure audit logging for integration processes.	"Logging an Integration Process" on page 110.
Configure analysis for webMethods-executed and externally executed processes using Optimize for Process.	*Administering webMethods Optimize*

Identifying the Integration Servers to Be Monitored

Use the following procedure to define which instances of Integration Server to monitor. You must specify at least one Integration Server. If you specify multiple servers, you identify one as the default. The default server or server pair cannot be deleted.

To identify the Integration Servers to be monitored

1. In My webMethods: **Navigate > Applications > Administration > My webMethods > System Settings > Servers**.

2. Click **Add Server** to add a new Integration Server (or to add an Integration Server/Analytic Engine host pair, if you are using Monitor with Optimize for Process).

3. Enter a name for your new server in the **Name** column, and then do one of the following:

   - To add an Integration Server for a BPM-only host, enter the host name or network address and port number in the **Integration Server (Monitor) Host** and **Port** columns. Select the check box in the **Use SSL** column if the server uses an SSL connection.
   - To add an Integration Server/Analytic Engine host pair, enter the host names or network addresses and the appropriate port numbers in *both* the **Integration Server (Monitor) Host** and **Port** columns and the **Analytic Engine Host** and **Port** columns. Select the check box in the appropriate **Use SSL** column if either server uses an SSL connection.

   Note:
   Depending on your installation environment, the Analytic Engine fields may be missing from the **System Settings > Servers** page.

4. Click **Save**.

5. Repeat the steps above to identify all Integration Servers to be monitored.

6. The default server or server pair appears in the **Default** column. To choose a new default, select the new default server and click **Save**.

   The selected server is the default for any new My webMethods user. After a user selects a server or server pair on a Monitor page, that server or server pair becomes the user's default.

   Important: Monitor uses the default Integration Server remote server alias “local” to resubmit a process instance or service when the node on which they were initially submitted is down. The default Integration Server remote server alias is required and must not be altered.

   Tip:
   Click **Check Server Status** to verify that a specified Integration Server or Analytic Engine is accessible.

Identifying the My webMethods Server that Hosts the Monitor User Interface

The Integration Server that hosts the WmMonitor package must know which My webMethods Server hosts the Monitor user interface to enable the user interface and package to communicate.

To identify the My webMethods Server that hosts the Monitor user interface

1. In Integration Server Administrator for the host Integration Server: **Packages > Management**.

2. In the WmMonitor row, click the **Home** icon.

3. Complete the first five fields in the **Configuration Settings**.

> **Note:**
> By default, the My webMethods Server port number is 8585. Enter a different port number in the **MWS Port** field only if a non-default port was specified during installation of My webMethods Server. If no value is entered, the **MWS Port** value is set to 8585.

4. Change any of the remaining configuration fields as necessary.

5. Click **Submit**.

Configuring Database Connection Retries

You can configure the number of times that Monitor attempts to connect to a database (such as the Process Audit Log database) from which it reads data. If Monitor cannot connect in the specified number of tries, it logs the error to the host Integration Server's error log.

To configure Monitor connection attempts

1. In Integration Server Administrator for the host Integration Server: **Packages > Management**.

2. Click the **Home** icon for the WmMonitor package.

3. In the **Database Retries** field, specify the number of tries.

4. Click **Submit**.

Configuring Central User Management

If you want My webMethods users to perform Monitor tasks using their My webMethods user name and password, you must enable and configure central user management. With central user management, when a My webMethods user issues a Monitor request, My webMethods Server invokes a service in the WmMonitor package on Integration Server to handle the request.

The service is invoked using the user name and password of the requesting user, and Integration Server authenticates the user. If the user name and password do not match an Integration Server user, Integration Server uses central user management to authenticate the user.

For complete information about enabling and configuring central user management, see the PDF publication *webMethods Integration Server Administrator's Guide*. Central user management may already be configured in your environment. If not, follow the instructions in *webMethods Integration Server Administrator's Guide* to enable and configure it. After central user management is working, complete the following tasks:

- "Verifying the Configuration of Central User Management in Integration Server " on page 21.
- "Adding My webMethods Users to the Monitor ACLs" on page 21.
- "Customizing How Monitor Sets Up ACLs When Using Central User Management" on page 22.

Note:
If you do not use central user management, you must ensure that each Monitor user defined in My webMethods has a corresponding user account defined in Integration Server.

Verifying the Configuration of Central User Management in Integration Server

To verify the configurations of central user management in Integration Server

1. In Integration Server Administrator for the Integration Server that hosts the WmMonitor package: **Security > User Management**.

2. Verify that the **Central User Management** field is set to **Configured**. If it is not, ask the administrator for that Integration Server to configure central user management.

3. In Integration Server Administrator for the Integration Server that hosts the WmMonitor package: **Settings > Resources**.

4. Under **Single Sign On with My webMethods Server**, verify that **MWS SAML Resolver URL** field is set to `https://mws-host:mws-port/services/SAML`. If it is not, ask the administrator for that Integration Server to configure single sign on.

5. In Integration Server Administrator for the Integration Server that hosts the WmMonitor package: **Settings > Extended**. Next, click **Edit Extended Settings** and verify that the following key/value pair is included in the extended settings:

   ```
   watt.server.auth.samlResolver=http://mws-host:mws-port/services/SAML
   ```

   If the setting is not defined, ask the administrator for that Integration Server to configure the setting.

Adding My webMethods Users to the Monitor ACLs

To add My webMethods users to the Monitor ACLs

1. In Integration Server Administrator for the Integration Server that hosts the WmMonitor package: **Security > ACLs**.

2. In the **Select ACL** field, click **MonitorAdministrators ACL**.

3. Click **Add** under the **Allowed** list to view the current groups in the Select Role/Group dialog box.

4. In the **Provider** field, click **Central**.

5. Type an asterisk (*) in the **Search** field and then click **Go** to populate the list of roles and groups.

6. Click **My webMethods Users** to add that role to the **Allowed** list.

7. In the **Select ACL** field, select **MonitorUsers ACL**.

8. Repeat steps 3 - 6 to add the **My webMethods Users** role to the MonitorUsers ACL.

9. Click **Save Changes**.

Customizing How Monitor Sets Up ACLs When Using Central User Management

By default, Monitor sets the ACLs for the WmMonitor services based on My webMethods functional privileges. This enables users to perform all actions for which they have functional privileges. However, you can configure Monitor so that it does not automatically set the ACLs; if you do so, *you must* set the ACLs for the WmMonitor services.

If a user has the functional privilege to perform an action in My webMethods and you fail to assign the corresponding ACLs to WmMonitor services, the user will receive errors in the My webMethods user interface.

To customize how Monitor sets up ACLs when using central user management

1. In the Integration Server Administrator for the Integration Server that hosts the WmMonitor package: **Packages > Management**.

2. Click the **Home** icon for the WmMonitor package.

3. To enable Monitor to automatically set the ACLs based on My webMethods functional privileges, select the **Add 'My webMethods Users' role to 'MonitorUsers' ACL** check box. To prevent Monitor from doing so, clear the check box.

4. Click **Submit** to save your changes.

Granting Users Access to Monitor

Access to Monitor pages and the functions available to you on those pages is subject to the My webMethods Server access and functional privileges feature. Although this product guide describes all pages and functions, some pages or functions may not be available to every user. If you require additional privileges, contact your My webMethods Server administrator.

Configuring Access to Monitor Pages, Actions, and Data

My webMethods Server administrators determine which pages in the Monitor user interface a user can access by assigning access privileges. For example, you can configure My webMethods so that a user can view pages related to monitoring process instances, but not allow the user to view pages related to monitoring services.

My webMethods Server administrators also determine which Monitor actions a user can perform by assigning functional privileges. For example, you can allow a user to view documents, but not to resubmit documents.

A My webMethods Server administrator can assign access and functional privileges to a user, group, or role.

Finally, My webMethods Server administrators determine the audit data (that is, specific business processes, services, or documents) upon which a user can act. This type of user privilege is called *content-based access* or *data-level security*. You assign these privileges to a role. For example, you can allow the Service Administrator role to act on service audit data.

- For more information about permissions management and the pages discussed below, see *Administering My webMethods Server*.
- For more information about the My webMethods user interface and its administrative functions, see *Working with My webMethods* and *Administering My webMethods Server*.

Granting Users Access to Monitor Pages

You must be a member of the My webMethods Server Administrator role to grant privileges. In My webMethods, use the **Navigate > Applications > Administration > System-Wide > Permissions Management** page to assign access privileges.

The following table describes the access privileges you can assign for Monitor pages.

To allow users to	In the Access Privileges section, select the check box
View process models that are available for monitoring.	**Administration > Business > Business Processes**
Archive data from the IS Core Audit Log and Process Audit Log databases.	**Administration > Business > Data Management**
View data about process instances.	**Monitoring > Business > Process Instances**
View data about services.	**Monitoring > Integration > Services**
View data about documents.	**Monitoring > Integration > Documents**

Granting Users the Ability to Perform Monitor Actions

You must be a member of the My webMethods Server Administrator role to grant privileges. In My webMethods use the **Navigate > Applications > Administration > System-Wide > Permissions Management** page to assign functional privileges.

The following table describes the functional privileges you can assign for Monitor pages.

To allow users to	In the Functional Privileges section, select the check box
Stop, suspend, and resume process instances.	**Business Monitoring > Processes > Stop, Suspend, Resume**
Resubmit process instances.	**Business Monitoring > Resubmit**
Modify the pipeline for a process instance and resubmit the process instance.	**Business Monitoring > Modify and Resubmit**
Resubmit services.	**Integration Monitoring > Services > Resubmit**
Modify the pipeline for a service and resubmit that service.	**Integration Monitoring > Services > Modify and Resubmit**
Resubmit documents.	**Integration Monitoring > Documents > Resubmit**
Modify and resubmit documents.	**Integration Monitoring > Documents > Modify and Resubmit**
Archive data or archive and delete data from the IS Core Audit Log and Process Audit Log databases.	**Data Management > Archiving**

Granting Users Permissions for Non-DBO Schemas

webMethods Monitor archiving services support non-dbo schemas in an SQL Server database.

To assign permissions for non-dbo schemas

1. Create a login with default database as Archive Database.
2. Create Archive schema in the same database.
3. Create Archive user for Archive login with default schema as Archive schema.
4. Add required roles to grant access to Archive user on Process/ISCore schema.
5. Create Archive user for archive login in Process/ISCore database.

Example

Example on assigning permissions for non-dbo schemas.

The following data is used in the example:

- Database Name forISCORE/PROCESS SCHEMA : *MYBPMDB*
- Database name for archival : *MYBPMDB_ARCHIVE*
- Integration Server core login name : *MYBPMDB_USER1*
- Integration Server core username : *MYBPMDB_USER1*
- Integration Server core schema name : *dbo*
- Archive login name : *MYBPMDB_ARCHIVE_LOGIN1*
- Archive username : *MYBPMDB_ARCHIVE_USER1*
- Archive schema name : *MYBPMDB_ARCHIVE_SCHEMA1*

1. Create login named `MYBPMDB_ARCHIVE_LOGIN1` with default database as `MYBPMDB_ARCHIVE` using MS SQL Studio.

2. Create Schema and add user in database `MYBPMDB_ARCHIVE` for login `MYBPMDB_ARCHIVE_LOGIN1`

```
USE MYBPMDB_ARCHIVE
 GO
CREATE SCHEMA MYBPMDB_ARCHIVE_SCHEMA1
 GO
CREATE USER [MYBPMDB_ARCHIVE_USER1] FOR LOGIN [MYBPMDB_ARCHIVE_LOGIN1] WITH
DEFAULT_SCHEMA=
[MYBPMDB_ARCHIVE_SCHEMA1]
 GO
ALTER AUTHORIZATION ON SCHEMA::MYBPMDB_ARCHIVE_SCHEMA1 TO MYBPMDB_ARCHIVE_USER1
 GO
ALTER USER MYBPMDB_ARCHIVE_USER1 WITH DEFAULT_SCHEMA = MYBPMDB_ARCHIVE_SCHEMA1
 GO
GRANT CREATE TABLE TO MYBPMDB_ARCHIVE_USER1
 GO
GRANT ALTER TO MYBPMDB_ARCHIVE_USER1
 GO
GRANT CREATE VIEW TO MYBPMDB_ARCHIVE_USER1
 GO
GRANT CREATE PROCEDURE TO MYBPMDB_ARCHIVE_USER1
 GO
GRANT EXECUTE TO MYBPMDB_ARCHIVE_USER1 WITH GRANT OPTION
 GO
GRANT REFERENCES TO MYBPMDB_ARCHIVE_USER1
 GO
```

3. Using the Database Component Configurator, create an Archive component using `MYBPMDB_ARCHIVE_LOGIN1` as User ID in the `MYBPMDB_ARCHIVE` database.

4. Add archive user in ISCore database so that it can access the main database and grant permission.

```
USE [MYBPMDB]
 GO
CREATE USER MYBPMDB_ARCHIVE_USER1 FOR LOGIN MYBPMDB_ARCHIVE_LOGIN1 WITH
DEFAULT_SCHEMA=dbo
 GO
GRANT CONTROL ON SCHEMA::dbo TO MYBPMDB_ARCHIVE_USER1
 GO
GRANT SELECT, INSERT, UPDATE, DELETE ON SCHEMA :: dbo TO MYBPMDB_ARCHIVE_USER1;
 GO
GRANT CREATE TABLE TO MYBPMDB_ARCHIVE_USER1
 GO
GRANT ALTER TO MYBPMDB_ARCHIVE_USER1
 GO
GRANT CREATE VIEW TO MYBPMDB_ARCHIVE_USER1
 GO
GRANT CREATE PROCEDURE TO MYBPMDB_ARCHIVE_USER1
 GO
GRANT EXECUTE TO MYBPMDB_ARCHIVE_USER1 WITH GRANT OPTION
 GO
GRANT REFERENCES TO MYBPMDB_ARCHIVE_USER1
 GO
```

5. Update schema names in `OPERATION_PARAMETER` as follows:

   Run set operation parameter service and set Integration Server core schema and Process schema, for example, *<databasename>.<schemaname>*,

Creating Archive Table and Assigning Permissions for Postgres Database

To create an archive table, assign proper permissions to archive users, and execute archiving services with postgres database configured in Integration Server, follow the examples below:

- Set the default archiving parameter in the `OPERATION_PARAMETER` table.

```
PostgreSQL: Process Audit Log schema name
```

- BPM Archival for PostgreSQL database only supports multiple schemas with a single database only.

```
CREATE DATABASE bpmarchive(db name);
 \c bpmarchive;(connect to database)
CREATE SCHEMA processbpm(process schema name);
SET search_path TO processbpm;
-- to create user --
CREATE USER processbpmuser (process schema user);
ALTER USER processbpmuser WITH ENCRYPTED PASSWORD 'manage';
ALTER USER processbpmuser SET search_path = processbpm;
-- to grant permissions --
GRANT CONNECT ON DATABASE bpmarchive TO processbpmuser;
GRANT USAGE, CREATE ON SCHEMA processbpm TO processbpmuser;
GRANT ALL ON SCHEMA processbpm TO processbpmuser;
```

```
-- the below grants are needed and they are available for new objects created in
the future --
ALTER DEFAULT PRIVILEGES IN SCHEMA processbpm GRANT ALL ON TABLES TO processbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA processbpm GRANT ALL ON SEQUENCES TO
processbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA processbpm GRANT ALL ON FUNCTIONS TO
processbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA processbpm GRANT ALL ON TYPES TO processbpmuser;
-- to create archive schema --
CREATE SCHEMA archiveprocessbpm(archive schema name);
SET search_path TO archiveprocessbpm;
-- to create user --
CREATE USER archiveprocessbpmuser(archive schema user) ;
ALTER USER archiveprocessbpmuser WITH ENCRYPTED PASSWORD 'manage';
ALTER USER archiveprocessbpmuser SET search_path = archiveprocessbpm;
-- to grant permissions --
GRANT CONNECT ON DATABASE bpmarchive TO archiveprocessbpmuser;
GRANT USAGE, CREATE ON SCHEMA archiveprocessbpm TO archiveprocessbpmuser;
GRANT USAGE, CREATE ON SCHEMA processbpm TO archiveprocessbpmuser;
GRANT ALL ON SCHEMA archiveprocessbpm TO archiveprocessbpmuser;
GRANT ALL ON SCHEMA processbpm TO archiveprocessbpmuser;
-- the below grants are needed and they are available for new objects created in
the future --
ALTER DEFAULT PRIVILEGES IN SCHEMA archiveprocessbpm GRANT ALL ON TABLES TO
archiveprocessbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA archiveprocessbpm GRANT ALL ON SEQUENCES TO
archiveprocessbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA archiveprocessbpm GRANT ALL ON FUNCTIONS TO
archiveprocessbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA archiveprocessbpm GRANT ALL ON TYPES TO
archiveprocessbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA processbpm GRANT ALL ON TABLES TO
archiveprocessbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA processbpm GRANT ALL ON SEQUENCES TO
archiveprocessbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA processbpm GRANT ALL ON FUNCTIONS TO
archiveprocessbpmuser;
ALTER DEFAULT PRIVILEGES IN SCHEMA processbpm GRANT ALL ON TYPES TO
archiveprocessbpmuser;
GRANT ALL ON ALL TABLES IN SCHEMA processbpm TO archiveprocessbpmuser;
GRANT ALL ON ALL SEQUENCES IN SCHEMA processbpm TO archiveprocessbpmuser;
GRANT ALL ON ALL FUNCTIONS IN SCHEMA processbpm TO archiveprocessbpmuser;
```

Identifying the Audit Data on Which Users Can Perform Actions

My webMethods Server administrators can limit the types of data that a user can view or manage. This type of access control is referred to as data-level security. If a user belongs to more than one role, that user has access to all of the types of data and functions granted to all of the roles of which that user is a member.

To limit access to audit data on a role basis, you must:

- Enable data security as described in "Enabling Data-Level Security" on page 28.
- Configure role access to available process audit data, as described in "Identifying Processes, Services, and/or Documents on Which a Role Can Act" on page 29.

How Data-Level Security Works with Functional Privileges

Functional privileges are global across all of the data to which a user has been granted access. For example, assume the following two conditions:

- The role `HR` is granted the functional privileges to start and stop process instances and is granted data-level security access to the `newHire` process. As a result, users assigned to the `HR` role can view, start, and stop instances of the `newHire` process.
- The role `Interns` is granted data-level security access to the `ProblemReporting` process. As a result, users assigned to the `Interns` role can view instances of the `ProblemReporting` process.

If a user is assigned to *both* the `HR` and the `Interns` roles, because functional privileges are global and the `HR` role has the privilege to start and stop processes, the user assigned to both roles are able to start and stop not only instances of the `newHire` process, but also instances of the `ProblemReporting` process.

If you want to limit privileges, one straight-forward way to do so is to set up two user accounts. For example, assume that you want to give a user the ability to start and stop instances of the `newHire` process, but you also want that user to be able to only view instances of the `ProblemReporting` process. For this scenario, you could set up user account `joeHR` and assign the user account `joeHR` to the `HR` role, and then set up user account `joeIntern` and assign the user account `joeIntern` to the `Interns` role. When logged in as `joeHR`, the user can view, start, and stop `newHire` process instances. When logged in as `joeIntern`, the user can only view `ProblemReporting` instances.

Note:
Data-level security is currently only supported in a single server environment.

Enabling Data-Level Security

When data-level security is disabled, users have unrestricted data access and can access all audit data. If you want to limit the data to which users have access, enable data-level security and then specifically identify the data to which different user roles have access.

To enable data-level security for Monitor

1. In Integration Server Administrator for the Integration Server that hosts the WmMonitor package, click **Packages > Management**.
2. Click the **Home** icon for the **WmMonitor** package.
3. Select the **Enable Data Level Security** check box.
4. In the **Data Level Security Administrator** field, type the user name of a user who has access to all My webMethods data and all pages of the My webMethods user interface.
5. Click **Submit** to complete your settings.

Identifying Processes, Services, and/or Documents on Which a Role Can Act

The following table lists the pages that users with access privileges can view when data-level security is *disabled*.

Pages that display data for	User can view
Services	Audit data for *all* services.
Documents	All logged documents.
Process instances	Audit data for *all* process instances.

When you *enable* data-level security, by default, roles are blocked from accessing information about any processes, services, or documents. After you enable data-level security, you must configure data-level security for specific roles to identify the processes, services, and/or documents that each role can view and act on. After you have configured data-level security for roles, if a user belongs to multiple roles, that user will be able to work with all of the processes, services, and documents identified in all the roles to which the user belongs.

To identify the data on which a user role can act

1. In My webMethods, click **Navigate > Applications > Administration > System-Wide > User Management > Roles**.

2. Search for the role for which you want to configure data-level security, and edit it. For more detailed instructions, see *Administering My webMethods Server*.

3. To configure data-level security for processes:

   a. On the Edit Role page, click the **Data Level Security** tab, and then click the **Business Process** link. My webMethods displays the list of all processes the role can currently access. The list is empty if no processes have been added yet.

   b. To add processes you want to allow this role to access, click **Add Processes**, use the Add Processes page to identify the processes you want to allow this role to access, and click **OK**.

   c. Click **Apply** on the Edit Role page.

4. To configure data-level security for services, repeat step 3 but click the **Service** link.

5. To configure data-level security for documents, repeat step 3 but click the **Document** link.

Audit Data Archiving and Deletion in Monitor

The types of audit data that Monitor records for transactions can be categorized in two schemas:

- **IS Core Audit Log tables.** Store audit data for documents, processes, services and Integration Server data.
- **Process Audit Log tables**. Store audit data for document control, process control, and service control data (for example, resubmit actions).

To archive IS Core Audit Log data, you must configure an Archive database. To archive Process Audit Log data, you can use either an Archive database or partitions.

- **Archive database**. Use stored procedures or built-in services to move audit data from the IS Core Audit Log and Process Audit Log tables into an Archive database. In this configuration, you must regularly schedule archiving or delete audit data from IS Core Audit and Process Audit Log tables to maintain peak logging performance.
- **Database partitions of Process Audit Log data**. Use partitioning for systems that generate high volume transactions, to automatically manage audit data from the Process Audit Log tables. You must still use an Archive database to manage all other audit data.

For information about how to configure the archive database, see "Overview of Archiving or Deleting Data in an Archive Database" on page 116. For information about partitions for archiving Process Audit Log data, see "Overview of Using Partitions for Process Audit Log Data" on page 126.

Preliminary Requirements

Prior to archiving or deleting data, ensure that the following requirements are met:

- You have created an Archive database. For instructions, see "Configuring the Archive Database" on page 31.
- Identify the users who will archive or delete data and assign them the appropriate data management permissions in My webMethods Server. For more information, see "Granting Users Access to Monitor Pages" on page 23.

Configuring Archive Settings

You can configure how Monitor archives audit data to the Archive database using the following options:

- **Stored procedures.** This is the default. When Monitor executes a stored procedure to archive or delete audit data, the database performs the entire archive or delete without further interaction from Monitor. To use stored procedures to perform an archive, the audit data must be archived to the same database where the stored procedure is located. Archiving using stored procedures is an asynchronous operation and should have little system impact. Stored procedures are especially useful for preserving Integration Server resources in high-volume situations.
- **Partitioning**. For more information about archiving with partitioning, including configuration, see "Overview of Using Partitions for Process Audit Log Data" on page 126.

You can archive or archive and delete audit data. When Monitor archives audit data, it moves it to the Archive database and removes it from the source tables. When Monitor deletes data, it deletes it from the source tables and does not move it to any other location.

After you archive or delete audit data, you can no longer view that data in My webMethods. However, you can still execute queries on the data in the Archive database using SQL statements.

After you configure data archiving and deletion, see "Overview of Archiving or Deleting Data in an Archive Database" on page 116 for information about data archiving and deletion procedures.

If you use an Oracle database, you can define a recipient of email alerts when the Oracle Purge operation completes. For instructions, see *Administering webMethods Optimize*.

Configuring the Archive Database

To use non-partitioned archiving, you must define the Archive database.

The following instructions provide a high level overview of the steps for creating the Archive database. For complete instructions, see the chapter, "Creating and Dropping Database Components" in *Installing Software AG Products*. The section, "Product Database Component Descriptions and Installation Requirements, describes specific details for each database provider.

To configure data archiving

1. Using Database Component Configurator, create the Archive database for the Process Audit schema.

   a. In the **Action** fields, select the following values from the table below:

Field	Properties
Action Type	Create
Action Component	Archive
Version	Latest

   b. In the **Connection** fields, define the connection to your Archive database, as follows in the table below.

Field	Properties
RDBMS	Select the database provider. The Process Audit and Archive databases must be of the same type.
URL	*URL address*

Field	Properties
User ID	Username to access the database. This must be a new user and have sufficient privileges to access both the source and target Process Audit database.
Password	Password

c. In the **Create Database and Database User** fields, define the database Administrator, as follows in the table below.

Field	Properties
Admin ID	Add the Archive database administrator.
Admin Password	Password for the Archive database administrator
Database	Name of the Archive database, for example, wmProcessAuditArchive.

d. Click **Execute**.

For detailed instructions on creating the Archive database, see the chapter, "Creating and Dropping Database Components" in *Installing Software AG Products*.

2. In the Database Administration console, assign the user the appropriate permissions for the tables in the Archive and Process Audit database.

3. Connect the Archive database to an Integration Server. For complete instructions on connecting to a database, see the section on configuring databases in *webMethods Integration Server Administrator's Guide*.

4. Define a new JDBC connection pool alias settings.

   a. In Integration Server Administrator, click **Settings > JDBC Pools**.

   b. In Pool Alias Definitions, click **Edit**.

   c. Add the URL, user ID and password to match the Connection settings defined with the Database Component Configurator and click **Save Settings**.

5. Define the JDBC pools for the Archive database.

   a. In Integration Server Administrator, click **Settings > JDBC Pools**.

   b. In Functional Alias Definitions, click **Edit** for Archiving.

c. In **Associated Pool Alias**, select the alias and click **Save Settings**.

d. Click **Restart**.

6. Configure the default archiving parameter in the OPERATION_PARAMETER table.

   a. In Designer, run the pub.monitor.archive:setOperationParameters service.

   b. Specify the input parameters as follows. pub.monitor.archive:setOperationParameters sets the values you specify in the `OPERATION_PARAMETER` table of the Archive database.

      The following table lists these parameters.

Parameter	Entry
`PROCESS_SCHEMA`	To archive data from the Process Audit Log tables, specify the following information for your database provider: ■ **Oracle**: Process Audit Log database user ■ **SQL Server**: Process Audit Log database name ■ **DB2**: Process Audit Log schema name ■ **MySQL**: Process Audit Log database name ■ **PostgreSQL**: Process Audit Log schema name
`ISCORE_SCHEMA`	To archive data from the IS Core Audit Log database, specify the following: ■ **Oracle**: IS Core Audit Log database user ■ **SQL Server**: IS Core Audit Log database name ■ **DB2**: IS Core Audit Log schema name ■ **MySQL**: IS Core Audit Log database name ■ **PostgreSQL**: IS Core Audit Log schema name

Note:
You can set additional parameters not listed in the table. For example, you can specify how many days of audit information to keep in the IS Core Audit Log and Process Audit Log schemas. For more information, see the setOperationParameters service in *webMethods Monitor Built-In Services Reference*.

7. Set database permissions to allow the Archive database user permission to select and delete data from the IS Core Audit Log tables, the Process Audit Log tables, or both, depending on the data you want to archive. To do so, execute the following SQL:

```
GRANT SELECT ANY TABLE, UPDATE ANY TABLE, DELETE ANY TABLE, INSERT ANY TABLE
```

Verify that you set permission for the Archive tables listed in "Overview of Archive Tables" on page 136.

Configuring Monitor Properties Using an MSR Configuration Variables Template

You can configure Monitor properties by using a Microservices Runtime configuration variables template. The following table lists the Monitor properties that you can configure by adding them to the template file.

Property	Value
My webMethods Server Host	`monproperty.wm.monitor.myWebmethodsHost=localhost`
My webMethods Server Port	`monproperty.wm.monitor.myWebmethodsPort=8585`
My webMethods Server Transport	`monproperty.wm.monitor.myWebmethodsTransport=http`
My webMethods Server Username	`monproperty.wm.monitor.myWebmethodsUserName=Administrator`
My webMethods Server Password	`monproperty.wm.monitor.myWebmethodsPassword=`
Database Retries	`monproperty.wm.monitor.dbRetries=3`
Connection Timeout	`monproperty.wm.monitor.connTimeout=10000`
Enable Data Level Security	`monproperty.wm.monitor.dlsEnabled=false`
Data Level Security Administrator	`monproperty.wm.monitor.dlsSuperUser=Administrator`
Resubmit to local Integration Server	`monproperty.wm.monitor.resubmitOnLocalIS=true`
Add 'My webMethods Users' role to 'MonitorUsers' ACL	`monproperty.wm.monitor.addMwUsertoMonitorACL=false`
My webMethods Server Root Context Path	`monproperty.wm.monitor.myWebmethodsPath=`

For more information on creating and using Microservices Runtime configuration variables templates, see *Developing Microservices with webMethods Microservices Runtime*.

3 Service Monitoring

About the Services Search Page

You can search for logged service data using the Services page. You can configure the following search options for this page:

- Search tab to display when initially displaying the page.
- Search to execute, if any, when initially displaying the page.
- Search results display. You can sort the results, define the number of rows to display, and define the columns to display.

You can save searches for logged service data and you can re-execute saved searches for logged service data.

For instructions on all these tasks, see *Working with My webMethods*.

Important:
To monitor services, you have to configure the ProcessAudit JDBC Pool Alias. For more information, see *webMethods Integration Server Administrator's Guide*.

Finding Logged Service Data Using Keywords

You can search for logged service data by specifying keywords found in the names of services or in service context IDs.

Searching based on context IDs is useful if you have set custom context IDs using the pub.flow:setCustomContextID service. For information about this service, see *webMethods Integration Server Built-In Services Reference*.

Note:
Whether a search is case-sensitive or case-insensitive depends on the way the underlying database (for example, Oracle, DB2, or SQL Server) handles the queries that Monitor issues to obtain data.

To find logged service data using keywords

1. In My webMethods, click **Navigate > Applications > Monitoring > Integration > Services**.
2. Click the **Keyword** search tab.
3. If Monitor is configured to use multiple Integration Servers, use the **Server** selection box (above the search panel) to specify the server to search.

   The **Server** selection box defaults to the last server specified by the logged in user. If the current user has never selected a server, the default server configured by the My webMethods Server administrator on the System Settings page is used.

4. In the text box, type keywords that are contained in the names or context IDs of the services to find. For example, you might specify:

   - The fully-qualified name of a service (such as, OrderPartner.Services:processOrder).
   - A partial service name (such as, processOrder) to select all services that contain the specified keyword.

   For more information about how to specify keywords, see *Working with My webMethods*.

   To view data for all logged services for which you are authorized (up to the maximum rows setting), leave the text box blank.

5. Click **Search**.

Finding Logged Service Data Using an Advanced Search

Use an advanced search to specify detailed criteria to search for specific logged service audit data.

> **Note:**
> Whether a search is case-sensitive depends on how the underlying database (for example, Oracle, DB2, or SQL Server) handles the queries that Monitor issues to obtain data.

To find logged service data using an advanced search

1. In My webMethods: **Navigate > Applications > Monitoring > Integration > Services**.

2. Click the **Advanced** search tab.

3. If Monitor is configured to use multiple Integration Servers, use the **Server** selection box (above the search panel) to specify which server to search. The **Server** selection box defaults to the last server specified by the logged in user. If the current user has never selected a server, the default server set by the My webMethods Server administrator on the System Settings page is used.

4. Specify the search criteria using the fields below. To not restrict the search by a certain field, leave the field blank.

   > **Note:**
   > You can find context IDs for services by viewing the Service Detail page (see "Service Statuses" on page 40).

The following table lists and describes these fields.

Field	Description
Service Name	Fully qualified or partial name of services to find (such as OrderPartner.Services:processOrder) or processOrder).

Field	Description
	Note: The top, unlabeled text box on the **Advanced** search tab and the **Service Name** field are both for specifying the full or partial name of the services to find. Use only one of these two fields.
Server ID	Integration Server on which the services to find ran or are running. Type the Integration Server's DNS name and port (such as `integration.east.rubicon.com:5555`) or partial DNS name or port (such as `rubicon`).
Context ID	Full service context ID of services to find.
Root Context ID	Full root context ID, to find all services that were invoked one after another starting with the specified root service.
Parent Context ID	Full parent service context ID, to find all services invoked by the specified parent service.
Status	Status of services to find. To select multiple statuses, hold down the CTRL key while selecting each status. For information about statuses, see "Service Statuses" on page 40.
User	Full or partial user name of the client that invoked services to find.
Activity Message	Full or partial message entered in the **Full Message** field in the **Activity Messages** panel on the Service Detail page. The **Full Message** field is populated if a service logs user-defined messages by calling the pub.prt.log:logActivityMessages service.
Filter By	Predefined date option to narrow the search. ■ Choose **Date Last Updated** to search for services based on the most recent date and time service data was logged. ■ Choose **Start Date** to search for services based on when services were started.
Range area	Search for services based on the most recent date and time data was logged for the services. You can choose a predefined time period in the **Range** list, or you can use the calendar pickers to specify a **Start Date** and **End Date** and then select the numbers for the hours and minutes from the lists. **Note:** If you want to add the search results to a My webMethods workspace, using a predefined time period causes the search results on the workspace to be dynamic, showing data relative to the current date (for example, yesterday). Using exact start and end dates causes the search results to always contain data for the specific dates you use, regardless of the current date.

5. If you want to search for services based on custom logged fields, use the **Filter** section of the page.

The following table lists and describes these fields.

Field	Description
Log Field Name	Full name of a custom logged field to use for the search. The field name is case-insensitive. Wildcard characters are not supported.
Operator	Select the operator to use: Equal, Contain, or NotContain, != , < , >, <= , >=
Value	Specify the value to use for comparison. Click **Add Row** to specify additional fields.

6. In the **Search Condition** list, select **AND** to find services that match all search criteria. Select **OR** to find services that match any search criteria.

7. Click **Search**.

Viewing Detailed Information for a Service

You can view detailed information for the services that Monitor displays in the search results on the Services page. In the search results, locate the service for which to view details and click **View Detail**. Monitor displays the Service Detail page.

On the Service Information panel, Monitor displays information that identifies the service, as described in the table below:

Field	Description
Service Name	Fully-qualified name of the service.
Root Context ID	Context ID of the root-level service.
Parent Context ID	Context ID of the service that invoked the service, which is referred to as the parent service. The parent context ID can be the same as the root context ID.
Context ID	Context ID of the service.
Custom Context ID	Custom value set for the context ID of the service using the pub.flow:setCustomContextID service.
Server ID	DNS name and port number of the Integration Server on which the service ran or is running.
User	User name of the user who invoked the service.

Field	Description
Timestamp	Date and time on which the activity indicated by **Current Status** (for example, **Failed**) was logged.
Current Status	Current status of the service. For more information, see "Service Statuses" on page 40.
Error Message	Most recent error message associated with the service. Monitor displays the **Error Message** field only if the **Current Status** is **Failed**.
Root Service	Fully qualified name of the root service of the service whose details Monitor is displaying.
Parent Services	Fully qualified name of all the service that directly invoked the service whose details Monitor is displaying.

The History panel shows the statuses the service has gone through and the date and time each status occurred. For a list of statuses, see "Service Statuses" on page 40.

If a service logged user-defined messages by calling the pub.prt.log:logActivityMessages service, the Activity Messages panel shows the date and time a message was logged, the type of the message (that is error, warning, or message), and a brief and long version of the text of the message.

If a service has been resubmitted, the Control Actions panel shows information about the resubmission. The panel shows the date and time the service was resubmitted, the action taken, the user name of the user who resubmitted the service, and the Integration Server on which the service was resubmitted.

Note:
Each time a service is resubmitted, the Integration Server assigns that service a new context ID.

If the service logged run-time values for custom fields, the Logged Fields panel shows the date and time the custom field was logged, the input or output parameter of the service for which run-time values were logged, and the name and value of the custom logged field.

If errors occurred while a service was running, the Service Errors panel shows the date and time each error was logged and a description of the error.

Service Statuses

Monitor displays statuses for services on the Services page and the Service Details page using a status keyword (for example, Completed or Started) and a status icon.

The following table lists the possible service status icons along with their meanings.

Icon	Meaning
	Service completed successfully.
	Service is running, but has encountered errors.

Icon	Meaning
	Service completed, but encountered errors.

The following table lists the possible service statuses along with their meanings.

Status	Meaning
Completed	Service has finished processing.
Failed	Service stopped processing because it encountered an error.
Resubmitted	Service has been resubmitted.
Started	Service has started and is currently executing.

Resubmitting a Service

You can resubmit a root-level service whose input pipeline was logged. The service can have any status.

To resubmit a service, an Integration Server remote server alias is required and the default Integration Server alias must exist and be unaltered. The default alias is used to resubmit a service when the original node on which the service was submitted is down.

When you resubmit a service, Monitor changes the status of the service to **Resubmitted**. Monitor then starts a new instance of the service and sets its status to **Started**. Monitor uses the context ID of the original service as the parent context ID for the new instance of the service. All data about the resubmission is logged for the new instance of the service.

To resubmit a service without first editing the input pipeline, search for the service, select the check box next to the service in the search results, and then click **Resubmit**.

To edit the input pipeline and then resubmit the service, search for the service, click **View Detail** for the service in the search results, and then click **Edit Pipeline**. On the Edit Pipeline page, update the fields, click **OK**, and then click **Resubmit**.

Important:
When you leave the Service Detail page, your changes are lost, so you must resubmit from this page.

4 Document Monitoring

About the Documents Search Page

You search for logged documents using the Documents page. You can configure the following search options for this page:

- Search tab to display when initially displaying the page.
- Search to execute, if any, when initially displaying the page.
- Search results display. You can sort the results, define the number of rows to display, and define the columns to display.

You can save searches for logged documents and you can re-execute saved searches for logged documents.

For instructions on all these tasks, see *Working with My webMethods*.

Finding Logged Documents Using Keywords

You can search for logged documents by specifying keywords found in the names of the documents.

Note:
Whether the search is case-sensitive or case-insensitive depends on how the underlying database (for example, Oracle, SQL Server, or DB2) handles the queries that Monitor issues to obtain data.

To find logged documents using keywords

1. In My webMethods: **Navigate > Applications > Monitoring > Integration > Documents**.
2. Click the **Keyword** search tab.
3. If Monitor is configured to use multiple Integration Servers, use the **Server** selection box (above the search panel) to specify the server to search.

   The **Server** selection box defaults to the last server that was specified by the current user. If the current user has never selected a server, Monitor uses the default server set by the My webMethods Server administrator set on the System Settings page.
4. In the text box, type keywords that are contained in the names of the documents to find. For example, you might specify:
   - The full document name as it exists on the webMethods Broker (deprecated) (such as `wm::is::OrderProcess::Implementation::CanonicalOrder`).
   - The full document name as it exists on the Integration Server (such as `OrderProcess.Implementation:CanonicalOrder`).
   - A partial document name (such as `OrderProcess`) to select all documents that contain the specified keyword.

> **Note:**
> If a document was routed through Universal Messaging, search for the fully qualified name of the publishable document type as it exists on Integration Server. You cannot search for the Universal Messaging channel name associated with a publishable document type.

For more information about how to specify keywords, see *Working with My webMethods*.

To view all the logged documents you are authorized to view (up to the maximum rows setting), leave the text box blank.

5. Click **Search**.

Finding Logged Documents Using an Advanced Search

Use an advanced search to specify detailed criteria to search for specific logged documents.

> **Note:**
> Whether the search is case-sensitive or case-insensitive depends on how the underlying database (for example, Oracle, DB2, or SQL Server) handles the queries that Monitor issues to obtain data.

To find logged documents using an advanced search

1. In My webMethods: **Navigate > Applications > Monitoring > Integration > Documents**.

2. Click the **Advanced** search tab.

3. If Monitor is configured to use multiple Integration Servers, use the **Server** selection box (above the Search panel) to specify which server you want to search. The **Server** selection box defaults to the last server specified by the current user. If the current user has not selected a server, the default server set by the My webMethods Server administrator on the System Settings page is used.

4. Specify the search criteria using the table below. To not restrict the search by a certain field, leave the field blank.

Field	Description
Document Name	Fully qualified or partial name of documents to find (such as `wm::is::OrderProcess::Implementation::CanonicalOrder` or `OrderProcess`). **Note:** The top, unlabeled text box on the **Advanced** search tab and the **Document Name** field are both for specifying the full or partial name of the documents you want to find. Use only one of these two fields.

Field	Description
Document ID	Full or partial ID of documents to find, or blank to not restrict the search by document ID. The webMethods Broker (deprecated) or Integration Server that publishes the document generates the ID when it publishes the document.
Client ID	Full or partial client ID associated with documents to find. Use partial client ID to search for documents associated with multiple clients. The value you specify for **Client ID** depends on the types of documents you are searching for (see **Type** field, below). ■ The format for webMethods Broker (deprecated) IDs is *Broker@host:port* (for example, `CustOps@qatest07:6849`, or partial ID `CustOps`. ■ The format for IDs of webMethods Broker (deprecated) clients is *clientprefix_folder1.folder2.foldern_trigger* (for example, `smitha_documenthistory.history.triggers_MsgHistoryWithNoResServiceTrigger`, or partial ID `smitha`). **Note:** A webMethods messaging trigger that subscribes to document types routed through webMethods Broker (deprecated) has a corresponding client on the webMethods Broker (deprecated). ■ An In Doubt document received from Universal Messaging does not have a client ID. In the search results, My webMethods displays "NA" for the client ID.
Type	Type of documents to find (such as **Broker**(deprecated), **In Doubt**, or **Retries Exceeded**). To select multiple types, hold down the CTRL key while selecting each type.

To search for this type of document	Specify
Broker(deprecated)	IDs of the webMethods Brokers (deprecated) that logged the documents.
In Doubt	When searching for documents routed through webMethods Broker (deprecated), specify the webMethods Broker (deprecated) client IDs for the triggers that processed the documents originally. **Note:** A document routed through Universal Messaging does not have a client ID.
Failed and **Retries Exceeded** when failure/retries exceeded occurred during delivery	IDs of the original destination webMethods Broker (deprecated) clients. For webMethods messaging triggers the

Field	Description	
		client ID is the client prefix for the webMethods Broker (deprecated) connection alias plus the trigger name.
	Failed documents when failure occurred during retrieval	IDs of the webMethods Broker (deprecated) clients associated with the triggers for which Integration Server originally tried to retrieve the documents.
	Note: For **Failed** and **Retries Exceeded** documents when failure occurred during publication, there is no client ID on which to search.	
Range	Search for documents based on the most recent date and time data was logged for the documents. You can choose a predefined time period from the **Range** list, or you can use the calendar picker to specify a **Start Date** and **End Date** and then select the numbers for the hours and minutes from the lists.	
	Note: If you want to add the search results to a My webMethods workspace, using a predefined time period causes the search results on the workspace to be dynamic, showing data relative to the current date (for example, yesterday). Using exact start and end dates causes the search results to always contain data for the specific dates you use, regardless of the current date.	

5. In the **Search Condition** list, select **AND** to find documents that match all search criteria. Select **OR** to find documents that match any search criteria.

6. Click **Search**.

Viewing Detailed Information for a Document

You can view detailed information for the documents that Monitor displays in the search results on the Documents page. In the search results, locate the document for which to view details and click **View Detail**. Monitor displays the Document Detail page.

The Document Information panel displays information that identifies the document as follows:

- Fully qualified name and identifier of the document, and date and time the document was logged.
- Type of the document (that is, webMethods Broker (deprecated), Failed, In Doubt, or Retries Exceeded).
- For webMethods Broker (deprecated) documents, ID of the publishing webMethods Broker (deprecated).

- For In Doubt documents received from webMethods Broker (deprecated), Failed documents that failed during delivery or retrieval, and Retries Exceeded documents that could not be delivered, the client ID for the intended recipient.
- For Failed documents that failed during publishing and for Retries Exceeded documents that could not be published, no client ID is listed.
- For an In Doubt document received from Universal Messaging, the message ID (or UUID) assigned to the document by Integration Server and the document ID automatically generated by the auditing subsystem in Integration Server.

  Note:
  An In Doubt document received from Universal Messaging does not have a client ID. In the search results, the Document Details page displays "NA" for the client ID.

- If the document was logged by webMethods Broker (deprecated), date and time the webMethods Broker (deprecated) first enqueued the document (that is, added it to the first subscriber's queue).

The Control Actions panel displays information relating to document resubmission, as follows:

- User name of the user that resubmitted the document and date and time the document was resubmitted.
- For webMethods Broker (deprecated) documents, the webMethods Broker (deprecated) to which the document was delivered (always the webMethods Broker (deprecated) to which the Monitor-equipped Integration Server is connected).

Resubmitting a Document

When you resubmit a document, Monitor logs a new instance of the document and all data about the resubmission is logged for the new instance.

To resubmit a document without first editing the fields, search for the document, select the check box next to the documents in the search results, and then click **Resubmit**.

To edit the fields of a document and then resubmit, search for the document, click **View Detail** for the document in the search results, and then click **Edit Document**. On the Document Detail page, update the fields, click **Save**, and then click **Resubmit**.

Important:
When you leave the Document Detail page, your changes are lost, so you must resubmit from this page.

Monitor resubmits each type of document as follows in the table below:

Document	Monitor
webMethods Broker (deprecated)	Publishes the documents to the webMethods Broker (deprecated) to which the Monitor-equipped Integration Server is connected.

Document	Monitor
In Doubt	Delivers the documents to the triggers that originally processed the documents.
Failed	■ **Failed during delivery**: Delivers the documents to the original destination webMethods Broker (deprecated) clients. ■ **Failed during publication**: Publishes the documents to the webMethods Broker (deprecated) to which the Monitor-equipped Integration Server is connected. ■ **Failed during retrieval**: Delivers the documents to the triggers for which Integration Server originally tried to retrieve the documents.
Retries exceeded	■ **Exceeded during delivery**: Delivers the documents to the original destination webMethods Broker (deprecated) clients. ■ **Exceeded during publication**: Publishes the documents to the webMethods Broker (deprecated) to which the Monitor-equipped Integration Server is connected.

5 Process Monitoring

Working with the Process Instances Page

You can search for logged process instance data on the Process Instances page in My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**. You can search by keyword or by defining advanced search criteria, and you can configure the following search options for this page:

- Specify the search tab to be displayed when the page initially opens.
- Specify the search to execute, if any, when the page initially opens.
- Define the search results display. You can sort the results, define the maximum number of rows to display, and define the columns to display.

You can save the search criteria for any search you create, so you can re-execute those saved searches without having to reconstruct them. You can also add any search results to a workspace in My webMethods.

For instructions for all these tasks, see the PDF publication *Working with My webMethods*.

Finding Process Instances Using Keywords

You can search for process instance data logged *for the current date* by specifying a keyword search term that is compared to process instance model names and process instance IDs.

Important:
Whether the search is case-sensitive or case-insensitive depends on how the underlying database (for example, Oracle, DB2, or SQL Server) handles the queries that Monitor issues to obtain data.

To find process instance data using keyword search

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.
2. Click the **Keyword** search tab.
3. If Monitor is configured to use multiple Integration Servers, use the **Server** selection box to specify the server to search.

   Note:
   The **Server** selection box defaults to the last server specified by the current user. If the current user has never selected a server, the default server set by the My webMethods Server administrator on the System Settings page is used.

4. Type a keyword search term in the **Keywords** field. The keyword search compares the specified search term with the Process Name and Process Instance ID fields for the available process instances and returns all matching process instances you are authorized to view (up to the maximum rows setting).

- If you want to return all process instances your are authorized to view, leave the **Keywords** field empty.
- Use * as a wildcard character
- Use " " for an exact phrase
- Outside of an exact phrase, a space character is treated as a logical AND.

5. Click **Search**. Optional: Click **Save** to retain the search criteria on the **Saved** tab.

Finding Process Instances Using an Advanced Search

Use an advanced search to specify multiple criteria to find specific process instances.

Important:
Whether the search is case-sensitive or case-insensitive depends on how the underlying database (such as Oracle, DB2, or SQL Server) handles the queries that Monitor issues to obtain data.

To find logged process instance data using an advanced search

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.

2. Click the **Advanced** search tab.

3. If Monitor is configured to use multiple Integration Servers, use the **Server** selection box to specify the server to search.

   Note:
   The **Server** selection box defaults to the last server that was specified by the current user. If the current user has never selected a server, the default server set by the My webMethods Server administrator on the System Settings page is used.

4. The **Advanced** tab provides a **Keywords** field. For information on its operation, see "Finding Process Instances Using Keywords" on page 52.

5. Specify additional search criteria using the fields from the table below.

Field	Selection
Process	Select **All** (the default), or select one or more models from a scrollable list of all available process models. To select multiple entries, hold down the CTRL key while you select each name. For webMethods-executed process models, the entries in the **Process** list identify the model version in addition to the process model name using the format *process_model_name - process_model_version*.

Field	Selection
Status	Select **All** (the default), or select one or more status values. To search for multiple statuses, hold down the CTRL key while you select each status.
Filter By	Select either of the predefined options: ■ Click **Date Last Updated** to search for process instances based on the last date and time process instance data was logged. ■ Click **Start Date** to search for process instances based on when process instances were started.
Date Range	Specify a date range for your search with either of the following options: ■ In the **Date Range** list, select **All** (the default), or select one of the predefined date ranges from the drop-down list. ■ Use the calendar pickers to specify custom date range in the **Start Date** and **End Date** fields, and then specify start and end time values in the **Time** fields.

6. In the **Search Condition** list:
   - Select **AND** to find process instances that match all search criteria.
   - Select **OR** to find process instances that match any search criteria
7. Click **Search**. Optional: Click **Save** to retain the search criteria on the **Saved** tab.

> **Note:**
> If you save the search to the **Saved** tab, be aware that if you specify exact start and end dates, the search results *will always contains data for the specific dates you use, regardless of the current date*. To avoid this, create your saved searches with one of the predefined relative date ranges available in the **Date Range** list (for example, **Previous Day** or **Previous Week**).

Customizing the Process Instance Search Options

You can customize your search preferences with the **Options** tab on the Search panel

To customize process instance search options

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.
2. Click the **Options** tab.
3. Do any or all of the following:

- Specify which search tab appears by default (Keyword, Advanced, Saved, or Saved - Details).
- Specify a selected saved search. At least one saved search must exist to make this selection, otherwise the list is empty.
- Specify if the selected saved search is to be run automatically when the Process Instances page is opened.

4. Do either of the following:
   - Enter a value in the **Max Results** field to limit the number of tasks returned as search results, or:
   - Select the **No Maximum** check box to return all search results. Use care when selecting this option. For systems running a large number of process instances, response time may slow appreciably.
5. Click **Save**.

Viewing Detailed Information for a Process Instance

Use the following procedure to view the details of process model instance.

To view detailed information for a process instance

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances.**
2. Search for the process model you want to view. For instructions, see "Finding Process Instances Using Keywords" on page 52 and "Finding Process Instances Using an Advanced Search" on page 53.
3. In the search results, locate the process instance you want to work with and click the Process Instance Detail icon in the **Detail** column.

About Process Instance Statuses

The Process Instances page displays a general status icon and a text status for each process instance. Status icons are as follows:

Note:
The status definitions described in this section apply to webMethods-executed and externally executed processes. Although integration processes display the same status text values (for example, Completed or Waiting), user-created services set the integration process status. As a result, the user-created services might have a different definition of the status than the one defined by the webMethods product suite.

The following table lists and describes the general status icons.

Icon	Meaning	Description
	Normal	The process instance is executing or executed without interruption or error.
	Statistically abnormal	The process instance has been suspended, has been stopped, or is running but one or more steps might be executing with errors.
	Out of compliance	The process instance completed, but one or more steps executed with errors.

The following table lists the text statuses along with their meaning.

Process Status	Meaning
Completed	Either the process completed successfully ( ) or the process completed but one or more steps executed with errors ( ).
Failed	Process instance stopped because one or more steps executed with errors, or the process instance is no longer being tracked and is missing information.
Failed (Escalated)	The parent process takes control of the failed process instance. The parent process receives notification of the failed process instance and continues executing. You cannot resubmit a process instance with this status because the parent process is no longer waiting for a response.
Resumed	Execution of the process instance was suspended, but has now been resumed.
Revised	Indicates that a running process instance has been updated to a new process model version during execution. When a process version is updated, the status of the running process instance is automatically changed to Revised regardless of its previous status.
Started	The process has started, but not all steps have completed. A status icon of indicates that one or more steps may be executing with errors.
Stopped	The process instance was stopped/canceled.
Suspended	Process execution was paused ( ). This status is not applicable to external (BAM-only) processes.
Resumed	The process was suspended and has been resumed. This status is not applicable to external (BAM-only) processes.

Process Instance Detailed Information

The Process Instance Detail page provides you with a central location that presents the most important information about a process instance. It offers the following buttons:

- **Previous** and **Next**. These buttons enable you to move through the process instances listed in the search results without having to return to the Process Instances page.
- **Refresh**. Updates the page with the most recent process data available.
- **Close**. Closes the Process Instance Detail page and returns to the Process Instances page.

The Process Instances Detail page provides the following information and controls:

Process Instance Information

The Process Instance Information window displays information that identifies the process instance, as follows in the table below:

Field	Description
Process	Name of the process model associated with the process instance.
Model Version	Name of the model version used for the process instance. This field is only applicable for webMethods-executed processes.
Start Date / Time	Date and time the process instance started.
Last Updated	Date and time of the last change in process instance status.
Instance ID	Unique identifier for the process instance. If a process instance ID consists of multiple parts (for example, order numbers from two or more different order systems), Monitor creates one row for each part.
Parent Instance ID	Unique identifier for the parent process instance, if applicable. Provided only for instances that have been started by other instances.
Instance Iteration	Number of times the process instance has been submitted.
Status	Status of the process instance. For more information, see "About Process Instance Statuses" on page 55.
Duration	Length of time a process instance was active. Duration is calculated based on the status of the process instance: - **Active**. Length of time the process instance has been executing. **Duration** is calculated by subtracting the start time of the process instance from the current system time. - **Inactive**. Length of time the process instance was active. **Duration** is calculated by subtracting the start time of the process instance from the time the process instance became inactive.

Field	Description
	■ **Timed out or stopped**. Length of time the process instance was active. **Duration** is calculated by subtracting the start time of the process instance from the time the cancel or timeout action occurred. **Note:** For steps that execute in parallel, **Duration** does not include the overlapping execution time.
Buttons:	These buttons are available only when a process instance is running (that is, does not have a status of Completed). ■ **Update**. Available when a process instance is running and a new version of the source process model has been enabled. This updates the running instance so that it uses the newly enabled model version for the rest of the process. For more information, see "Updating a Process Instance to a New Model Version" on page 76. ■ **Suspend/Resume**. Click to suspend or resume a process instance. ■ **Stop**. Click to stop the process instance.

Process Stage Timeline

The Process Stage Timeline window displays a graphical presentation of stage activity and status over time. webMethods Optimize for Process must be installed to be able to view stage timeline information.

Note:
When a process contains no stages, or if webMethods Optimize for Process is not installed, this window is collapsed to show only the title bar. It cannot be expanded.

Note:
Process stage logging does not occur instantaneously. This can result in a delay of up to five minutes between stage instance changes and the stage status displayed in the Stage Timeline window. For example, when viewing a recently executed process instance, stages may initially appear as not yet started. Click the **Refresh** button on the Process Instance Detail page to update the Stage Timeline window. For more information, see "About Process Generation and Stage Status Display" in the Building and Uploading chapter of *webMethods BPM Process Development Help*.

By default, stages of all statuses are displayed, as indicated by the stage status display check boxes:

You can filter the list of displayed stages by clearing a stage status check box, which removes stages of that status from the display. Select a check box to display stages of that status.

Note:
Click on a column name to activate the sort controls for that column.

The table below describes the information that the stage list provides.

Column	Description
Status	The current status of the stage: Completed, Running, Incomplete, or Not Started.
Name	The name assigned to the stage.
Cycle Time	Cycle time is based on the status of the stage: ■ **Completed**. The length of time it took for the stage to run from start to completion. ■ **Running**. The length of time since the stage began executing. ■ **Incomplete**. The length of time it took for the stage to run from start to the time it stopped running for any reason other than completion (for example, step failure or process suspended) A stage will have incomplete status if it is still running when the process terminates for any reason. This can happen with normal process execution if the process instance took a split branch that does not terminate. ■ **Not Started**. No cycle time is displayed.
Deviation	Deviation is calculated only when the stage is breached. A value of +/- 0 indicates no deviation. A stage that has breached its defined condition is indicated by an alarm icon . The deviation time is calculated by comparing the Cycle Time to the Condition expression defined for the stage. If the condition specifies <, then the stage is breached when the cycle time exceeds the specified time period. If the condition specifies >, then the stage is breached when the cycle time is less than the specified time period. For example: ■ If the Condition expression is “< 2 seconds” and the stage completes in 3 seconds, the deviation is 1 second, and the stage is breached. ■ If the Condition expression is “> 2 hours” and the stage completes in 1 hour and 30 minutes, the deviation is -30 minutes, and the stage is breached.

You can obtain additional stage information as follows:

- Hover the cursor over any row in the stage list, or over the colored bar for a stage in chart area to view the stage name and description as well as the following:
  - **Status**: The current status of the stage.
  - **Start Milestone**: The point in the process where the stage starts.

- **End Milestone**: The point in the process where the stage ends.
- **SLA Condition**: The expected execution time specified in the stage's Condition expression.
- **Stop Tracking On Breach**: Indicates if this option is selected (Yes or No).

- Clicking any row in the stage list does all of the following:
  - Displays the stage's start milestone and stage end milestone in the Process Diagram window. Only one stage can be displayed in the Process Diagram at any time. Click another row to display the milestones for another stage.

    Note:
    If you have selected the **Stop Tracking On Breach** option for a stage and that stage experiences a breach, all stage tracking ceases and any remaining stages in the process are marked as Incomplete. In this case, it is possible for the steps within the stage to be shown as Complete in the process diagram, while the stage that contains them is shown as Incomplete. This is expected behavior.

  - Filters the Step Summary list to display only those steps that executed within the stage's cycle time. Steps that are in a different path from the stage's start and end milestone steps will also be included if the step's execution time falls within the stage cycle time.
  - Filters the Logged Fields list to display only those logged fields that belong to a step in the filtered Step Summary list.

    Click **Clear Selection** in the Process Stage Timeline window to remove the milestone icons and display all process steps in the Step Summary window and all logged fields in the Logged Fields window.

- To set the time resolution for the chart, select from Year, Month, Week, Day, Hour, or Minute in the **Time Unit** list.

Process Diagram

The Process Diagram window displays an image of the process model as it was designed in Designer, if an image is available for a webMethods-executed or externally executed process model. Because no model exists for an integration process, Monitor cannot display a process diagram for integration processes.

Note:
If you are using Internet Explorer to view process models rendered with Google Web Toolkit (GWT), you must configure the compatibility settings in My webMethods Server for your version of the browser. For more information about this, and about process model rendering in general, see "About Process Model Rendering" on page 107.

The image contains status icons next to steps that have executed or are currently executing.

You can right-click the diagram and use the following menu commands to resize the diagram and change the label display.

- **Fit to screen**. Fits the entire process model diagram on the screen. Click **Restore** to restore the diagram size to 100%. You can also use the slider on the left side of the screen to adjust the zoom level.
- **Hide transition labels** or **Show transition labels**. This menu command toggles to either hide or show labels on the transition lines between events or steps in the process.
- **Show transition descriptions** or **Hide transition descriptions**. This menu command toggles to hide or show transition descriptions. If no description is defined, the transition expression label appears when **Show transition descriptions** is selected. This option is mutually exclusive with **Show/Hide transition expressions**.
- **Show transition expressions** or **Hide transition expressions**. This menu command toggles to show or hide transition expressions. This option is mutually exclusive with **Show/Hide transition descriptions**.
- **Truncate transition expressions/descriptions** or **Expand transition expression/descriptions**. This menu comment toggles to display full or shortened label descriptions.

Step Summary

The Step Summary window displays information about the execution of the steps within the process instance.

Note:
The list of steps can be filtered by stage by selecting the stage in the list portion of the Process Stage Timeline window.

The table below describes the fields that the Step Summary window contains.

Field	Description
Step Name	Name of the step.
Start Date/Time	Date and time the step began executing.
Last Updated	Date and time the step was last updated.
Instance Iteration	Number of times the process instance executed, including resubmissions. For externally executed processes, this value is always 1.
Step Iteration	Number of times the step was executed as a result of: ■ A transition drawn to this step downstream step that causes the step to execute multiple times during a single process execution. ■ A process resubmittal that caused the step to execute again.

Field	Description
Loop Iteration	Applies only to steps that can be configured for standard looping, such as a subprocess step or a call activity step. A step executes a loop iteration only when a loop condition is configured for the step and that condition has been met. **Note:** A step can loop two or more times until the loop condition is no longer met. A separate row exists for each loop iteration. Any time the step iteration increments, the loop iteration count is reset and begins again with 1.
Status	Status of the step. For more information, see "About Process Instance Step Statuses" on page 68.
Duration	Length of time the step took to execute.
Referenced Processes	webMethods-referenced process or BPMN callable process that was executed at run time for this step. Click **child instance(s)..** to view the **Child instances for step** window. This window shows a list with details for all child process instances of the viewed step.
Detail	Click the Step Detail icon in the **Detail** column to view step details, and for processes that are being analyzed, you can also click the KPI Summary icon to view KPI information.

Control Actions

The Control Actions window contains information only when a control action (that is, suspend, resume, resubmit, or stop) has been performed on a webMethods-executed process instance, as follows in the table below:

Field	Description
Date / Time	Date and time the control action was performed.
Action	Action taken on the process instance (suspend, resume, resubmit, or stop).
Instance Iteration	Instance ID for the process instance.
Step Name	Name of the step.
Step Iteration	Number of times the step executed.
Server ID	Server on which the process instance executed.
User	User associated with the process instance.

Activity Messages

The Activity Messages window contains information only when the process instance ran a service that logged user-defined messages.

The following table describes the fields of this window.

Field	Description
Date / Time	Date and time the activity message was logged.
Step Name	Name of the step that logged the activity message.
Entry Type	The message type (Information, Debug, Error, or Warning).
Brief Message	Shortened version of the message.
Full Message	Full version of the message.

Logged Fields

The Logged Fields window contains information only when the process instance logged run-time values for custom fields. Custom logged fields are specified when the process is created in Designer.

Note:
The list of logged fields can be filtered by stage by selecting the stage in the list portion of the Process Stage Timeline window.

The following table describes the fields of this window.

Field	Description
Date / Time	Date and time the custom field was logged.
Step Name	Name of the step that logged the custom field.
Instance Iteration	Number of times the process instance executed, including resubmissions. For externally-executed processes, this value is always 1.
Step Iteration	Number of times the step was executed as a result of: ■ A transition drawn to this step downstream step that causes the step to execute multiple times during a single process execution. ■ The process was resubmitted, causing the step to execute again.
Loop Iteration	Applies only to steps that can be configured for standard looping, such as a subprocess step or a call activity step. A step executes a loop iteration only when a loop condition is configured for the step and that condition has been met. **Note:**

Field	Description
	A step can loop two or more times until the loop condition is no longer met. A separate row exists for each loop iteration. Any time the step iteration increments, the loop iteration count is reset and begins again with 1.
Input/Output	**Input** indicates that any values passed as input to the step are logged. **Output** indicates that output pipeline values for the step were logged.
Field Name	Name of the logged custom field.
Field Value	Value of the logged custom field.

Process Errors

The Process Errors window displays error information for the process instance. Only errors for the entire process instance are shown here; step errors are shown in the Process Step Detail page. For more information, see "Viewing Detailed Information for a Process Step" on page 67.

The following table describes the fields of this window.

Field	Description
Date/Time	Date and time the error occurred.
Error	Type of error that occurred. The following are some typical errors: ■ **Process Timeout**. The process timed out before it could be completed. ■ **Step Timeout.**A step timed out before it could be completed. ■ **Retries Exceeded.** A step was executed more than the defined maximum number of times. ■ **Out of Sequence**. A step tried to execute out of sequence. ■ **Runtime error**. The process instance encountered a run-time event (for example, failure to validate an account) that prevented it from completing successfully.
Error Message	Text of the error message. If the error is a run-time error (for example, if at run time an account was rejected), this column also displays the run-time error code.
Message Detail	Full exception message thrown by one of the services executing within the process instance.
Service Name	Name of the service that threw the exception.
Step Iteration	Number of times the step was executed as a result of: ■ A transition drawn to this step downstream step that causes the step to execute multiple times during a single process execution.

Field	Description
	■ The process was resubmitted, causing the step to execute again.
Step Name	Name of the step associated with the error.
Server ID	For webMethods-executed processes or integration processes, identification for the server on which the error occurred. Server ID is not available for externally executed process errors.

Setting the Priority for User Tasks in a Process Instance

You can set the priority of the user tasks in a selected process instance through the Monitor user interface in My webMethods. When setting the user task priority for a process instance, the priority is set for all user tasks that are not yet completed.

To set the priority of the user tasks in a process instance:

1. In My webMethods Server: **Applications** > **Monitoring** > **Business** > **Process Instances**
2. Click for the process instance in the **Detail** column.
3. In the **Process Instance Information** panel, click **Update User Task Priority**.

   Monitor displays the **Update User Task Priority** window.
4. Select the User Task priority for the current process instance from one of the following options:
   - Critical
   - High
   - Medium
   - Low
   - None
5. Click **Apply**.

   A tooltip is displayed providing the IDs of the user tasks with updated priorities.

Path Forecasting for a Process Instance

You can check the path forecasting for a process instance through the Monitor user interface in My webMethods. Path forecasting is based on aggregated historical data collected by Optimize and is available for currently running process instances that have been enabled for analysis.

When viewing the details for a process instance, you can select a forecast path and view the following estimated data for that path:

- **Estimated Completion Time** - The estimated time of completion if the forecast path is taken.
- **Percentage Complete** - The estimated percentage of completion for the process instance based on the selected forecast path's Average Path Cycle Time.
- **Average Path Cycle Time** - The average duration of the forecast path, calculated based on aggregated average step duration of the forecast path.
- **Average Process Cycle Time** - The average execution duration of previously completed process instances. Process instances that were not fully completed do not contribute to the average cycle time.
- **Path Frequency** - The frequency of the forecast path taken based on samples of historical data.

The estimated time of completion data is displayed for the entire process instance, not just for a single step. As the number of previously completed process instances increases, the accuracy of estimation also improves, because the estimation is based on a larger historical sample.

The path forecasting feature uses Optimize to provide estimations based on previously completed process instances. Optimize has a mode for calculating process and step statistical metrics, which is governed by the Optimize Analytic Engine's Monitor Behavior Setting. This setting has three modes of operation:

- All days are the same - all days of the week contribute to the same average.
- Work days and weekend days -weekdays contribute to one average, while weekend days contribute to a separate average.
- All days are different - each day of the week has its own average.

For more information on specifying statistical intervals, see *Administering webMethods Optimize*.

You should take into account the Optimize statistical mode of operation when checking estimated data for a forecast path. For example, if today is Tuesday and the Optimize statistical mode is "all days are different", the estimation is based on past process instances completed on a Tuesday.

Note:
The Optimize Analytic Engine only calculates the process and step metric after the end of the day and averages do not include process instances for the current day.

Configuring Your System to Path Forecasting for a Process Instance

To configure your system to use path forecasting for process instances:

1. In: **My webMethods** > **System Settings** > **Servers**, select one of the following server environments:

- BAM
- BPM and BAM

2. In **Business** > **Business Processes** open your process for editing and on the Process Details tab select the **Analysis Enabled** check box to enable the process for analysis.

3. In Designer Process Development, open your process.

4. On the Properties tab, click **Run Time**.

5. In the **Quality Of Service** section, set the **Minimum Logging Level** to **2 - Error only** or higher.

Viewing Estimated Data for a Forecast Path

To view the estimated data forecast path of a process instance:

1. In My webMethods Server: **Applications** > **Monitoring** > **Business** > **Process Instances**

2. Click for your process instance in the **Detail** column.

3. In the **Process Diagram** panel, enable **Path Forecasting** by clicking the **On** radio button.

   Monitor displays the **Path Forecasting** bar to the left of the process diagram.

4. Click one of the dots on **Path Forecasting** bar to see a forecast path for the process instance.

   The different forecast paths are sorted from most common to least common by default. You can change the type of sorting from the drop-down list.

5. On the **Path Information** pop-up window, click **Show Stats** to view the estimated data for the forecast path.

When you select a forecast path, for which to view the estimated data, the path is highlighted. This shows whether the different parts of the path are completed (blue highlight) or not completed (black highlight). Forecast paths are always sequential and parallel paths are not taken into account.

Viewing Detailed Information for a Process Step

You can view detailed information for each step listed in the Step Summary area on the Process Instance Detail page.

To view detail information for a step within a process instance

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.

2. Search for the business process instance that contains the step that you want to view.

3. In the search results, locate the process instance you want to work with and click in the **Detail** column.

4. In the Step Summary area of the Process Instance Detail page, locate the step you want to view and click in the **Detail** column.

   Note:
   The icon is not available for rows representing a subprocess or call activity loop iteration, as there is no additional information available.

If the status of the process instance is Resubmitted or Suspended and the input pipeline for the step you are viewing has been logged, you can view the pipeline by clicking **View Pipeline**.

About Process Instance Step Statuses

The Process Instance Detail page displays status information for the execution of steps within a process instance.

The Process Diagram page displays a status icon next for each currently executing or completed step.

The following table lists these icons.

Icon	Meaning
	Step is running.
	Step has completed successfully.
	Step failed.
	Step is waiting (for example, to receive an external document).

The Step Summary area displays a general status icon and a text status for each process step.

General status icons are as follows in the table below:

Icon	Meaning	Description
	Normal	The step is executing or executed without interruption or error.
	Statistically abnormal	The process instance has been suspended or stopped, or the step is running but possibly with errors.
	Out of compliance	The step failed, or completed but executed with an error.

Text statuses are as follows in the table below:

Status	Meaning
Completed	Step has completed processing.
Expired	The time specified for the step in the process model for an event to occur expired before the event was satisfied. For example, the time set for a join condition expired before a document required by the join condition arrived.
Failed	Step stopped because one or more errors occurred, or the step is no longer being tracked and is missing information.
Retries Exceeded	An attempt was made to execute the step more times than is specified.
Unsatisfied Join	A join definition for the step has not been satisfied. That is, some incoming transitions have not arrived.
Started	Step has started but not completed.
Stopped	Step was stopped.
Waiting	Step is waiting for an event to occur (for example, the step might be waiting for a document to arrive).
Interrupted	Step has been interrupted by an interrupting boundary event.

Viewing Steps Within a Subprocess

Each subprocess typically contains one or more steps. When a subprocess is configured for standard looping, looping information is also available.

To view the steps within a subprocess

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.

2. Display detailed information for the process instance that contains the subprocess you want to work with, as described in "Viewing Detailed Information for a Process Instance" on page 55.

3. In the Step Summary area, locate the subprocess you want to view and click the expansion icon ⊞ next to the subprocess name. Monitor displays the steps within the subprocess in the table.

   - If the subprocess is not configured for looping, a table entry is available for each step in the process.
   - If the subprocess is a BPMN subprocess configured for standard looping, each loop iteration of the subprocess is shown as a separate table entry. Click the expansion icon ⊞ next to

the subprocess iteration to see the step summary information for each step in the process during that loop iteration.

> **Note:**
> Any time the parent subprocess step iteration increments, the loop iteration count is reset and begins again with 1.

4. Subprocess that contain any child subprocesses also display an expansion icon. Continue expanding the subprocesses until you locate the step you want to view.

> **Tip:**
> To expand and contract all subprocess instances in the table, click the expansion icon next to the **Step Name** column title.

> **Note:**
> This expandable/collapsible display is not available for deprecated webMethods subprocesses. In this case, all step entries are available at the top level of the table only.

Viewing a Call Activity Step

A call activity step starts an instance of another process model that exists outside of the current process and, if desired, incorporates the results of that referenced process into the pipeline of the current process. When a call activity is configured for standard looping, looping information is also available.

To view the results of a call activity

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.

2. Display detailed information for the process instance that contains the subprocess you want to work with, as described in "Viewing Detailed Information for a Process Instance" on page 55.

3. In the Step Summary area, locate the call activity for which to view details.

   - If the call activity is not configured for looping, Monitor displays a table entry with the step summary information for the call activity.
   - If the call activity is configured for standard looping, Monitor displays each loop iteration of the call activity as a separate table entry. Click the expansion icon ⊞ next to the call activity iteration to see the step summary information for each referenced process during that loop iteration.

     > **Note:**
     > Any time the call activity step iteration increments, the loop iteration count is reset and begins again with 1.

4. You can view the details of the referenced process (and from there, the summary information for the steps within it) in either of two ways:

- In the Step Summary window, click the link in the **Referenced Subprocess** column.
- In the Process Diagram window, click the + icon located within the referenced process step.

Click the **Close** button in the Process Details page for the referenced process to return to the parent process.

Process Step Detailed Information

The Process Step Detail page provides you with a central location that presents the most important information about a process step. It offers the following buttons:

- **Close**. Closes the Process Step Detail page and returns to the Process Instance Detail page.

For specific information about various aspects of step monitoring, see these topics:

- "About Subprocess Detailed Information" on page 73
- "About Subprocess and Call Activity Duration Time" on page 74

The Process Step Detail page provides the following information and controls:

Process Instance Information

The Process Instance Information window displays the same information provided on the Process Instance Detail page.

Step Information

The Step Information window displays information about the selected step. The data displayed here is mostly the same as the information displayed in the Step Summary window.

The following table provides additional information.

Column	Description
Server Type	Identifies the type of server where the step was executed. For webMethods-executed processes, this value is Process Engine.
Server ID	Displays the host name and port number of the executing server. For example, my.integration.server:5555.
Invoked Service	For webMethods-executed processes, identifies the service that the step executed, if the step is a Flow step. This row does not appear for non-service steps or externally executed processes.
Task Detail	For task steps in a webMethods-executed processes, identifies the task that the step executed. This row does not appear for non-task steps or for externally executed processes.

If the step is enabled for resubmission, the following buttons are available in the Step Information window:

- **Save Pipeline to File.** If the pipeline was logged, this button enables you to view the pipeline data or save it to a file.
- **Edit Pipeline.** Enables you to modify the pipeline date prior to resubmission.
- **Resubmit.** Resubmits the process instance to begin executing with this step.

For more information about working with these buttons, and about resubmission in general, see "About Resubmitting Process Instances and Process Steps" on page 76.

Step History

The Step History window provides a list of step status transitions that occurred as the step executed, enabling you to trace the step activity.

The following table describes the fields of this window.

Field	Description
Date/Time	Displays the date and time recorded for the particular status.
Status	Displays each status logged for the step.
User	For webMethods-executed processes that involve webMethods Task Engine for user interaction, identifies the user associated with this step status. For externally executed processes, this column is empty.
Role	For webMethods-executed processes that involve webMethods Task Engine for user interaction, displays the role associated with this step status. For externally executed processes, this column is empty.

Control Actions

The Control Actions window contains information only when a control action has been performed on a webMethods-executed process step, as follows in the table below:

Field	Description
Date / Time	Date and time the control action was performed.
Action	Action taken on the process step.
User	User associated with the process step action (if applied by a user).
Server ID	Server on which the process step executed.

Activity Messages

The Activity Messages window contains information only when the process instance ran a service that logged user-defined messages.

The following table describes the fields of this window.

Field	Description
Date / Time	Date and time the activity message was logged.
Entry Type	The message type (Information, Debug, Error, or Warning).
Brief Message	Shortened version of the message.
Full Message	Full version of the message.

Logged Fields

The Logged Fields window contains information only when the process instance logged run-time values for custom fields. Custom logged fields are specified when the process is created in Designer.

The following table describes the fields of this window.

Field	Description
Date / Time	Date and time the custom field was logged.
Field Name	Name of the logged custom field.
Field Value	Value of the logged custom field.

Step Errors

The Step Errors window displays error information that occurred as the step was executing.

The following table describes the fields of this window.

Field	Description
Date/Time	Date and time the error occurred.
Error Type	Type of error that occurred. For example, PRT.STEP.FAILED. Some errors may not contain an Error Type.
Error Message	Text of the error message.

About Subprocess Detailed Information

When you create a process model on the Software AG Designer Process Development perspective, you can add one or more subprocesses to the model. In addition, you can add a subprocess to a subprocess, nesting as many levels as you need. You can view detailed information for subprocess steps in the Step Summary area, just like any other step. Each subprocess typically contains one or more steps. When a subprocess is configured for standard looping, looping information is also available.

Note:
Do not confuse a subprocess step with a webMethods-referenced process step or a BPMN call activity step. Referenced processes and BPMN callable processes execute another process model, with the process model name shown in the **Referenced Process** column of the Step Summary area.

To view the steps within a subprocess, see "Viewing Steps Within a Subprocess" on page 69.

About Subprocess and Call Activity Duration Time

The Step Summary area contains a **Duration** column, which displays the length of time that the step took to execute. This value represents the time to execute for the activity:

- For a subprocess, this is the time to execute all of the steps in the subprocess.
- For a call activity, this is the time to execute the callable process.

Therefore, it would seem logical that the duration time of the subprocess or call activity would equal the sum of the durations of all the steps in the subprocess, or the duration of the callable process.

However, the duration may actually be greater than the sum of its internal objects. This is because the subprocess or call activity duration also includes communication time, as well as time taken by minor transport delays caused by network latency.

In addition, when the *WaitForSubprocess* value in the service pub.prt.SubprocessModel is set to `false` for a dynamic reference process, the call activity step in the parent is shown as ended as soon as the child instance is started. Therefore, the duration for that step iteration in the parent is initially very short. After the child instance completes, the duration value changes to the actual duration of the child instance. Monitor computes the duration from the start record to whatever the last record is for a step.

In cases where you use the dynamic reference process method to initiate multiple instances of a child model, the duration value represents the *duration of the longest of the child instances that were run*, and not the sum of all the durations of all the children.

Viewing KPI Data for Process Instances

When you work with webMethods Optimize for Process, you can create *key performance indicators* (KPIs) to measure critical success factors of a process instance. For example, in an order management process, you might define KPIs for how many orders were received, their dollar amounts, and whether they were processed successfully.

Optimize for Process creates and stores data for each KPI. You can view the KPI values to see how they have varied over time. For more information about KPIs, see *webMethods Optimize User's Guide*.

To view KPI data that is associated with a process instance

1. Search for the process instance for which you want to view KPI data, as described in "Finding Process Instances Using Keywords" on page 52 and "Finding Process Instances Using an Advanced Search" on page 53.

2. In the search results, locate the process instance you want to work with and click the KPI Summary icon in the **Detail** column.

Viewing KPI Data Associated with a Process Step

In the Step Summary area on the Process Instance Detail page, you can view KPI data that is associated with specific steps in a process instance. For more information about KPIs, see *webMethods Optimize User's Guide*.

To view KPI data that is associated with a step within a process instance

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.

2. Display detailed information for the process instance that you want to work with, as described in "Viewing Detailed Information for a Process Instance" on page 55.

3. In the Step Summary area of the Process Instance Detail page, locate the step you want to view and click **KPI Summary**.

Stopping, Suspending, or Resuming a Process Instance

You can apply the following actions to a process instance:

- Permanently stop a process instance from executing. You can stop any process instance *except* those with the status, Completed, Failed, Failed (Escalated), or Stopped.
- Suspend a process instance to temporarily pause its execution. You can pause any process instance *except* those with the status, Completed, Failed, Failed (Escalated), or Suspended.
- Resume a currently suspended process instance (that is, any process with the status, Suspended).

Note:
In My webMethods, you can stop, suspend, or resume only webMethods-executed process instances.

To stop, suspend, or resume a process instance

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.

2. Display detailed information for the process instance that you want to work with, as described in "Viewing Detailed Information for a Process Instance" on page 55.

3. On the Process Instance Detail page, click **Stop**, **Suspend**, or **Resume**, as needed.

Updating a Process Instance to a New Model Version

When you update a process model to create a new version and then enable that new version, you can also update any or all of the currently running instances of that model so that they start using the newer version. You can update:

- **All running instances**. Do this by responding to prompts when you enable the new process model version. For more information, see "Enabling and Disabling Process Model Versions" on page 97.
- **A single instance**. Do this by clicking the **Update** button on the Process Instance Detail page, as described in the procedure below.

When you update a running instance, the instance begins to use the newer version with the next step that is executed.

Important:
If the new version is incompatible with the one the instance was originally using, the process instance will fail. For example, the instance might fail if the new instance no longer contains the next step that the instance was to execute.

For more information about creating a new process model version, see the topic "Working with Process Versions" in the *webMethods BPM Process Development Help*.

To update a process instance to a new model version

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.

2. Display information for the process instance that you want to upgrade as described in "Viewing Detailed Information for a Process Instance" on page 55.

3. Click **Update** in the Process Instance Information area on the Process Instance Detail page.

Monitor changes the process instance status to Revised, regardless of its previous status.

About Resubmitting Process Instances and Process Steps

You can use the resubmit feature in Monitor to repair and re-execute process instances that have failed, or to re-execute completed process instances. Note that only webMethods-executed processes can be resubmitted.

A key component of the resubmit feature is the ability to edit the pipeline associated with a step within the process instance prior to the resubmit operation. This is useful for modifying pipeline data that could be causing the failure.

The Process Instance page contains two buttons for resubmitting process instances, **Resubmit Closest** and **Resubmit Earliest**. These buttons are active only when applicable conditions exist. Note that the specific behavior of the system in response to a process resubmit request depends on which button you click and on the state of the process, step, and model, as explained in the following paragraphs.

> **Note:**
> The resubmit buttons and their associated table check boxes are displayed only when an Integration Server (IS) is available.

You can resubmit one or more process instances from the **Navigate > Applications > Monitoring > Business > Process Instances** page. In most cases, you will need to use the Search feature to find processes with the appropriate status within the desired time frame. The resubmitted process begins executing at the appropriate step, depending on the factors described in the following paragraphs and the resubmit button that you choose.

A process with a "Completed" status may be resubmitted based on whether the process has steps that are resubmit-enabled, as follows:

- **Process contains a resubmit-enabled step** may be resubmitted using either button. In each case, the process instance execution is re-initiated.
- **Process does not contain a resubmit-enabled step** cannot be resubmitted. Clicking either button has no effect.
- **Process contains a resubmit-enabled step** but **the step was not executed during the original run** cannot be resubmitted. Clicking either button has no effect.

A process with a "Suspended" status resumes from the point at which it was suspended if you click either the **Resubmit Closest**, **Resubmit Earliest**, or **Restart** buttons. When you resubmit a step of a suspended process, webMethods Monitor first resumes the process and then resubmits the suspended process from the requested step.

Process instances with a "Failed" status may be resubmitted, with the following order of precedence:

- The **Resubmit Closest** button does the following based upon whether the failed step is resubmit-enabled:
  - **Failed step is resubmit-enabled**, the process instance resumes executing at the failed step.
  - **Failed step is not resubmit-enabled**, the process instance begins executing the closest step to the failed step that is resubmit-enabled that has been executed previously.
  - **Failed steps are not resubmit-enabled** and there are no steps that have been executed previously, no action is taken.
- The **Resubmit Earliest** button executes the remainder of the process instance from the first step that is resubmit-enabled. If no steps are resubmit-enabled, no action is taken.

- The **Restart** button ignores the resubmit-enabled status of the steps in the process instance and resubmits all of the failed steps.

You resubmit a process instance by resubmitting a step within that process instance. The resubmittal procedure varies depending on the status of the steps within the process:

- A process instance with a "Completed" status may be resubmitted from any step in the instance, provided you have first enabled the step for resubmission before the process instance began executing.
- A process instance with a "Failed" status may be resubmitted from any process instance step that has failed. A failed step is automatically enabled for resubmission, and the pipeline at the point of failure is automatically logged by the Process Engine and made available to Monitor. No manual configuration is needed. Process instances with Failed (Escalated) status cannot be resubmitted as the parent process is no longer waiting for a response.

Enabling Your System to Resubmit Processes

When the Integration Server that hosts Monitor is not connected to any messaging product and has no process model fragments, and you want to be able to resubmit processes, follow the steps below:

1. In Integration Server Administrator, go to the **Package > Management** page and click for the WmMonitor package.
2. Clear the **Resubmit to local IS** check box, click Submit, and then reload the package.

Requirements for Submitting Process Instances

The following requirements apply to resubmitting completed or failed processes:

- Only webMethods-executed processes can be resubmitted. You must have a BPM or a BPM and BAM server environment selected in the **Server** list at the top of the Process Instances page.
- For a completed process, the step from which you want to resubmit the process instance must be enabled for resubmission before the process instance began executing. For more information about resubmit enabling, see "Enabling and Disabling a Step for Resubmission" on page 101.
- The logging level for the process model from which the instance was started must be set to a level that will log the input pipelines for those steps, as described in "About Process Model Logging Levels" on page 94.
- You must have privileges to resubmit process instances and any referenced processes that it may call. For more information, see "Granting Users Access to Monitor " on page 22 and "Identifying the Audit Data on Which Users Can Perform Actions" on page 27.
- An Integration Server must be available. The resubmit buttons and their associated table check boxes are displayed only when this is the case.

- You have appropriate Integration Server ACL settings to resubmit the process instances to a remote Integration Server in a non-clustered environment. If not:

  1. Create an ACL using Integration Server Administrator. Go to the **Security > Access Control Lists** page and create a new ACL to specify which user groups can access the remote server.

  2. Go to the **Settings > Remote Servers > Create Alias** page, select the ACL created in the previous step from the **Execute ACL** list, and create a remote server alias.

  When the Integration Servers are clustered, Monitor resubmits the process instances to its host Integration Server.

Resubmittal Behavior in the Run Time

The behavior that occurs when you click the **Resubmit** button depends on the status of the process instance itself.

For a process instance in "Failed" status, resubmitting a failed step in the instance:

- Marks the process instance as "Resubmitted"
- Executes that step with the pipeline that is passed in from Monitor.
- Leaves the instance iteration of the process instance unchanged.
- Increments the step iteration of the resubmitted step from its previous iteration.
- Executes the remainder of the process instance from the point of resubmission.

For a process instance in "Completed" status, resubmitting a step in the instance:

- Creates a completely new iteration of the process instance.
- Increments the process instance iteration.
- Executes the remainder of the process instance from the point of resubmission.

Resubmitting Processes from a Step and Optionally Editing Pipeline Information

To resubmit a step in a process instance and optionally edit pipeline information

1. In My webMethods: **Navigate > Applications > Monitoring > Business > Process Instances**.

2. Search for the process instance you want to resubmit, as described in "Finding Process Instances Using Keywords" on page 52 and "Finding Process Instances Using an Advanced Search" on page 53.

3. In the search results, locate the process instance that you want to resubmit and click **Detail**.

> **Note:**
> If your process instance has failed in multiple places, it may be necessary to resubmit multiple steps to get the instance running correctly. If this is the case, repeat the resubmit procedures for each step that requires resubmission.

4. On the Step Summary window of the Process Instance Detail page, locate the step that you want to resubmit and click **Detail**.

> **Note:**
> At this point, you can view or save the step pipeline data in XML format for additional analysis. Click **Save Pipeline To File**, and then click **Open** to view the file or click **Save** to save the file, then click **OK**.

5. To edit the input pipeline for a step on the Process Step Detail page:

   a. In the Step Information window, click **Edit Pipeline**.

   > **Important:**
   > After you update the pipeline, the changes you make are available *only* while you are still on the Process Step Detail page. If you leave this page, the changes you make are not saved. If you want to edit the pipeline before resubmitting, you must make your edits and then *immediately* resubmit the process step.

   > **Note:**
   > The pipeline for the original iteration of the step is left unchanged. Monitor logs the updated pipeline with the new iteration of the step.

   b. On the Edit Step page, for each field value you want to change:

      a. Click the hyperlinked field value that you want to change. If the value of a field is not displayed with an active hyperlink, you cannot change that field value.

      b. Specify a new value for the field.

      c. Click **OK**.

      d. After you modify all the fields you want to change, click **Save**. Monitor returns to the Process Step Detail page.

6. Click **Resubmit**.

Example Resubmittal Use Cases

The following use cases demonstrate the most common scenarios for resubmitting a process instance.

Unhandled Exception

One of the most common use cases for resubmission is when a step in a process fails and there is no exception handler for that step, referred to as an *unhandled exception*. If an unhandled exception occurs, both the step and process instance are marked as "Failed" in Monitor and the pipeline for the failed step is saved at the point of failure.

The process instance continues to execute if there other tracks in the process that were executing prior to the failure, until the instance reaches a logical stopping point.

You can then locate the failed instance and step, edit the pipeline, and resubmit the failed step. The process instance will then continue executing from the point of the unhandled exception to its logical conclusion using the modified pipeline values.

Handled Exceptions

In this case, the process model implementation includes an exception handler. For example, you can model a boundary error event that transitions to an error handler sequence flow in the event of step failure at run-time. This is considered a *handled exception*.

The resubmit feature can still be used to resubmit failed steps in this case. However, this is not a common pattern, as the purpose of the exception handler is to execute whatever logical behavior is needed to prevent the failure of the process instance in the event of a step failure.

In other words, if your exception handler is designed properly and working as expected, you would not normally need to resubmit the process instance. If you have to resubmit failed steps in process instances with an exception handler, you are advised to modify your exception handling logic to correctly process the source of the failure.

Steps Enabled for Resubmission

Resubmission is not limited to steps that have failed. For example, you might want to resubmit a successfully completed step that is situated earlier in the process flow than the actual failed step to repair a failed process instance. Editing the pipeline at this previous, successfully completed step might be necessary to enable the failed step to complete successfully.

Monitor supports this only if the completed step you want to resubmit was previously enabled for resubmission. You must enable the step for resubmission prior to the execution of the process instance. To enable a step for resubmission, see "Enabling and Disabling a Step for Resubmission" on page 101.

Resubmitting Completed Processes

A less common but supported pattern is the resubmission of a successfully completed process instance. Note that a successfully completed process can still contain a failed step, if the process model also contains an exception handler step that was able to complete the process successfully. In this case, you can resubmit either the failed step or any other step that is enabled for resubmission. A totally new instance of the process is started and execution begins from the point of resubmission. Depending on the process model, it may be necessary to resubmit multiple steps to complete the process.

6 Working with Process Models

Finding Process Model Information in My webMethods

Process model information is available in the following location in My webMethods:

On the Business Processes page in My webMethods: **Navigate > Applications > Monitoring > Administration > Business > Business Processes**.

Using the Business Processes Page

To access the Business Processes page: **Navigate > Applications > Administration > Business > Business Processes**.

The Business Processes page provides a paged list all of the available process models you are authorized to view. Initially, the process model list might be empty, depending on your search options.

You can search for process models by keyword or by defining advanced search criteria, and you can configure the following search options and do the following actions on this page:

> **Note:**
> Viewing information about instances, threshold status and enabling or disabling a process model for execution and analysis is not allowed in BAM only mode and for BAM only process models.

- Specify the search tab to display when initially opening the page.
- Specify the search to execute, if any, when initially opening the page.
- Define the search results display. You can sort the results, define the maximum number of rows to display, and define the columns to display.
- View a process model diagram.
- View information about running, failed, and other instances by hovering the cursor over the number of instances in those columns.
- Set thresholds for the archiving policy based on the number of process instances or data in the database.
- Enable or disable a process model for execution or analysis.
- View detailed information about each process model and modify some attributes of the process model, including logging level and whether process steps can be resubmitted, as described in "Viewing and Modifying Process Model Information" on page 89.

You can save the search criteria for any search you create, so you can re-execute those saved searches without having to reconstruct them. You can also add any search results to a workspace in My webMethods.

For instructions for all these tasks, see the PDF publication *Working with My webMethods*.

The Business Processes page displays the following information about available process models:

Column	Description
Process Name	Name assigned to the process model in the application where the model was created. Click the link to view detailed information about each process model as described in "Viewing and Modifying Process Model Information" on page 89. If two or more versions of the process model are available, the information for the most recent version appears when you click this link.
Model Version	Click the expansion icon ⊞ next to the process name to view versions of the process model, if available. Model versions are applicable only for webMethods-executed process models. Click the link to view detailed information about each version as described in "Viewing and Modifying Process Model Information" on page 89.
Analysis	Indicates whether the process is made available to Optimize for monitoring and analysis. This column is available only when Optimize for Process is installed. For more information, see "Enabling or Disabling a Process Model for Analysis" on page 99. **Note:** This field is applicable only to webMethods-executed and externally-executed process models. You cannot enable an integration process for analysis. Click the enabled icon ✔ or the disabled icon ○ in the row for the process model to enable or disable for analysis.
Execution Enabled	Indicates whether Process Engine uses this version of a webMethods-executed process model to start new process instances. You can only enable one version of a webMethods-executed process model at a time. This field is not applicable to externally executed and integration processes. For more information, see "Enabling and Disabling Process Model Versions" on page 97. Click the enabled icon ✔ or the disabled icon ○ in the row for the process model to enable or disable for execution.
Used	Indicates whether the process model version has been used for at least one process instance. ■ **Yes**. The process model version has been used at least once. ■ **No**. The process model version has never been used; that is, there have never been any instances of this model version. When a version of a model has never been used, you can delete the model version.

Column	Description
	■ **Unavailable**. Monitor is currently not able to provide use information because the server used to log instances of the model version is not available.
Total Instances	Displays the total number of instances of the process model. Hover your mouse over the total number of instances to see statistics about the number of Running, Completed, Failed and Other instances. The color of the indicator is different based on the threshold number: ■ Green - when the number of instances run is less than 1/4 of the threshold. ■ Yellow - when the number of instances run is between 1/4 and 3/4 of the threshold. ■ Red - when the number of instances run is more than 3/4 of the threshold. The indicator shows you when to archive the instances based on the threshold set for that process.
Date Deployed	Date and time the process model was last updated in the Process Audit Log database component.

Finding Process Models Using Keywords

You can search for process models by specifying one or more keywords found in the following:

- For all process model types: Search results include all process models where the keyword is contained in the process model name.
- For webMethods-executed process models only: Search results include all process models where the keyword is contained in the process model name or the description of the process model.

Important:
Whether the search is case-sensitive or case-insensitive depends on how the underlying database (for example, Oracle, DB2, or SQL Server) handles the queries that Monitor issues to obtain data.

To find process models using a keyword search

1. In My webMethods: **Navigate > Applications > Administration > Business > Business Processes**.

2. Click the **Keyword** search tab.

3. If Monitor is configured to use multiple Integration Servers, use the **Server** selection box to specify the server to search.

   Note:
   The **Server** selection box defaults to the last server specified by the current user. If the current user has never selected a server, the default server set by the My webMethods Server administrator on the System Settings page is used.

4. Type a keyword in the **Keywords** field. The keyword search compares the specified search term with the fields described above (depending on the process type) and returns all matching process models you are authorized to view (up to the maximum rows setting).

   - If you want to return all process models your are authorized to view, leave the **Keywords** field empty.
   - Use * as a wildcard character
   - Use " " for an exact phrase
   - Outside of an exact phrase, a space character is treated as a logical AND.

5. Click **Search**. Optional: Click **Save** to retain the search criteria on the **Saved** tab.

Finding Process Models Using an Advanced Search

Use an advanced search to specify multiple criteria to find specific process models.

Important:
Whether the search is case-sensitive or case-insensitive depends on how the underlying database (such as Oracle, DB2, or SQL Server) handles the queries that Monitor issues to obtain data.

To find process models using an advanced search

1. In My webMethods: **Navigate > Applications > Administration > Business > Business Processes**.

2. Click the **Advanced** search tab.

3. If Monitor is configured to use multiple Integration Servers, use the **Server** selection box to specify the server to search.

   Note:
   The **Server** selection box defaults to the last server specified by the current user. If the current user has never selected a server, the default server set by the My webMethods Server administrator on the System Settings page is used.

4. You can search for process models using these fields:

- By keyword, as described in "Finding Process Models Using Keywords" on page 86.
- By process name. Click the **Process Name** field to select an entry from a scrollable list of all available process models.
- By **Process Description**. Type a word or phrase from a process description.
- By execution enabled status. Click the **Execution Enable** list to select a value of **All**, **Yes** (models that are enabled), or **No** (models that are not enabled).
- By analysis enabled status. Click the **Analysis Enable** list to select a value of **All**, **Yes** (models that are enabled), or **No** (models that are not enabled).

If two or more fields are specified, the values are combined with a logical AND.

5. Click **Search**. Optional: Click **Save** to retain the search criteria on the **Saved** tab.

Customizing the Process Model Search Options

You can customize your search preferences with the **Options** tab on the Search panel

To customize process instance search options

1. In My webMethods: **Navigate > Applications > Administration > Business > Business Processes**.
2. Click the **Options** tab.
3. Do any or all of the following:
   - Specify which search tab appears by default (Keyword, Advanced, Saved, or Saved - Details).
   - Specify a selected saved search. At least one saved search must exist to make this selection. Otherwise, the list is empty.
   - Specify if the selected saved search is to be run automatically when the Business Processes page is opened.
4. Do either of the following:
   - Enter a value in the **Max Results** field to limit the number of tasks returned as search results, or:
   - Select the **No Maximum** check box to return all search results. Use care when selecting this option. For systems running a large number of process instances, response time may slow appreciably.
5. Click **Save**.

Viewing and Modifying Process Model Information

You can view detailed information for a process model as well as modify various run-time settings for the model.

To view detailed information for a process model and modify run-time settings

1. In My webMethods: **Navigate > Applications > Administration > Business > Business Processes.**

2. Search for the process model you want to view. For instructions, see "Finding Process Models Using Keywords" on page 86 and "Finding Process Models Using an Advanced Search" on page 87.

   Do any of the following:

   - To view the process model detailed information, click the process model name or click **Edit** to open the Edit Process page.
   - To sort the table by the contents of a column, click on a column name to activate the sort controls for that column.

Working with the Edit Process Page

Monitor displays detailed process model information on the Edit Process page. The available information depends on the type of process (for example, webMethods-executed, externally executed, or integration) and the webMethods products you have installed.

In addition to viewing read-only information about the process, you can also add, modify, or delete various run-time settings as well as process stages. These changes are written to the Process Audit database when you save your work. To update your process model in Software AG Designer with these changes, you must use the appropriate **Synchronize** button available in the process model Properties view to apply the new database values to the process model in Designer.

Note that if you apply changes to your process model in Software AG Designer, Designer will write those changes to the Process Audit database during the build and upload procedure, and they will be applied to the next instance run from the process model. For more information, see the *webMethods BPM Process Development Help*.

Process Information Administration Window

This window contains the following tabs:

Process Details Tab

The **Process Details** tab displays fields that identify the process model, as follows in the table below:

Field	Description
Process Name	Name assigned to the process model in the tool where the model was created.
Model Version	Version of the process model, as set in Designer. The model version is applicable only for webMethods-executed process models. You can only have a single version of externally executed and integration processes, and the model version for these types of processes is always set to 1.
Description	Description of the process model defined in the tool in which the model was created. **Note:** For integration processes, this is the value of the input variable, *processLabel* in the pub.monitor.integrationProcessLogging:createProcessMetadata service.
Created By	User name of the user who created the process model. **Note:** For integration processes, this is the user that invoked the pub.monitor.integrationProcessLogging:createProcessMetadata service.
Date Deployed	Date and time the process model was last updated in the Process Audit Log database component: ■ For a webMethods-executed process model version, the last time the version was built and uploaded for execution from Designer. ■ For an externally executed process model, the last time the process model was uploaded for analysis from Designer. ■ For an integration process, when the service, pub.monitor.integrationProcessLogging:createProcessMetadata, logged information about the process.
Execution Enabled	Whether Process Engine uses this version of a webMethods-executed process model to start new process instances. You can only enable one version of a webMethods-executed process model at a time. This field is not applicable to externally executed and integration processes because the Process Engine does not manage the execution of these types of processes. For more information, see "Enabling and Disabling Process Model Versions" on page 97.
Analysis Enabled	Whether the process is made available to Optimize for monitoring and analysis. This column is available only when Optimize for Process is installed. For more information, see "Enabling or Disabling a Process Model for Analysis" on page 99. **Note:**

Field	Description
	This field is applicable only to webMethods-executed and externally executed process models. You cannot enable an integration process for analysis.
Used	Indicates whether the process model version has been used for at least one process instance. ■ **Yes**. The process model version has been used at least once. ■ **No**. The process model version has never been used; that is, there have never been any instances of this model version. When a version of a model has never been used, you can delete the model version. ■ **Unavailable.** Monitor is currently not able to provide use information because the server used to log instances of the model version is not available.

Process Settings Tab

The **Process Settings** tab displays logging settings that apply to webMethods-executed process models only. You can specify the amount and type of data you want the Process Engine to log for process instances that use this process model version, whether you want the Process Engine to log process transitions, and whether diagnostic logging is enabled.

For more information, see "About Process Model Data Logging" on page 93 and "Configuring Logging Settings for a Process Model Version" on page 95.

The Process Settings Tab contains the information from the following table:

Field	Description
Logging Level	For webMethods-executed processes only. Specifies how much data Process Engine logs for process instances that use this process model version. For more information, see "About Process Model Data Logging" on page 93.
Instances Threshold	The number of process instances that can run for this process model before you need to archive. The indicator for each process model on the Business Processes page is based on this number.
Log Transitions	For webMethods-executed processes only. Specifies if Process Engine logs process transitions. For more information, see "About Process Model Data Logging" on page 93.
Diagnostic Logging	Specifies that log messages from instances started from a process model are logged to a separate file for diagnostic purposes.

Step Settings Tab

The **Step Settings** tab display the steps contained in the process model version of a webMethods-executed process model, and enables you to specify whether a step can be used to resubmit a process instance executed from the model version. For more information, see "Enabling and Disabling a Step for Resubmission" on page 101.

Instance Analytics Tab

The **Instance Analytics** tab is available only when Optimize for Process is installed. This panel displays a list of error types associated with the current process. Additionally, you can create standard and custom error types. For more information, see the PDF publication *Administering webMethods Optimize*.

Process Stages and EDA Events Window

This window contains the following tabs:

Stages Tab

Although process stages are typically defined at design time in Designer, you can modify and delete those stages, or create new stages in the process model. After you save your changes in Monitor, you can synchronize these changes to your process model in Designer, as described in "Working with the Edit Process Page" on page 89.

To sort the stages table by the contents of a column, click on a column name to activate the sort controls for that column.

For information about working with stages and milestones in Monitor, see "Working with Stages and Milestones" on page 102. For detailed information about stages and milestones, see the *webMethods BPM Process Development Help*.

Events Tab

The Events tab enables you to enable or disable event emission for predefined EDA events from Process Engine. For more information, see "Enabling and Disabling EDA Event Emission" on page 100.

Process Diagram Window

The Process Diagram window displays an image of the process model as it was designed in Designer, if an image is available for a webMethods-executed or externally executed process model. For information about how to ensure that process models render correctly, see the section, "About Process Model Rendering" on page 107.

Note:
Because integration processes are not modeled, Monitor cannot display a model image for an integration process.

Tip:
You can right-click the diagram and use the following menu commands to resize the diagram and change the label display.

- **Fit to screen**. Fits the entire process model diagram on the screen. Click **Restore** to restore the diagram size to 100%. You can also use the slider on the left side of the screen to adjust the zoom level.
- **Hide transition labels** or **Show transition labels**. This menu command toggles to either hide or show labels on the transition lines between events or steps in the process.
- **Show transition descriptions** or **Hide transition descriptions**. This menu command toggles to hide or show transition descriptions. If no description is defined, the transition expression label appears when **Show transition descriptions** is selected. This option is mutually exclusive with **Show/Hide transition expressions**.
- **Show transition expressions** or **Hide transition expressions**. This menu command toggles to show or hide transition expressions. This option is mutually exclusive with **Show/Hide transition descriptions**.
- **Truncate transition expressions/descriptions** or **Expand transition expression/descriptions**. This menu command toggles to display full or shortened label descriptions.

About Process Model Data Logging

A Process Engine can log data for webMethods-executed process instances. You can view this data and perform actions on it in Monitor.

- For each process model version, you specify the amount and type of data to log in the **Logging Level** setting.

  Note:
  If you want to be able to resubmit process instances from Monitor at certain steps, you must set the logging level to a level that will log the input pipelines for those steps, and you must enable resubmission for each of those steps using the **Resubmit Enabled** setting, as described in "Enabling and Disabling a Step for Resubmission" on page 101.

- When Monitor renders a process diagram, it shows all the possible paths that can be taken within a process instances. If you want to see the path the process instances actually took at run time, use the **Log Transitions** setting to enable process transition logging for the process model version. The lines for the path that was actually executed are displayed as heavier lines.

- You can select the **Diagnostic Logging** option to specify that log messages from instances started from a process model are logged to a separate file for diagnostic purposes. For more information, see "Enabling and Disabling Process Instance Diagnostic Logging" on page 96.

Important:
When you regenerate a process model version, the logging settings return to the default values, and you must reset them if you want different settings.

About Process Model Logging Levels

The following table describes the available process model logging levels, with additional information about choosing a particular level.

You want to log	If in Monitor, you want to be able to			Set to
	View process status?	View step status?	Resubmit a process?	
■ Nothing (that is, disable process logging)	No	No	No	**1 - None**
■ Process status when steps fail ■ Input pipelines for failed steps ■ Run-time values for document fields	At failed step	No	At failed step	**2 - Errors only**
■ Process status ■ Input pipelines for failed steps ■ Run-time values for document fields ■ Optionally, transitions	Yes	No	At failed step	**3 - Process only**
■ Process status and start step status ■ Input pipelines for start steps and failed steps ■ Run-time values for document fields ■ Optionally, transitions	Yes	For start step	At start or failed step	**4 - Process and start events**
■ Process status and all step statuses ■ Input pipelines for every step ■ Run-time values for document fields ■ Optionally, transitions	Yes	For all steps	At any step that has logged input pipeline	**5 - Process and all events, activities, and looped activities**

You want to log	If in Monitor, you want to be able to			Set to
	View process status?	View step status?	Resubmit a process?	
■ Loop count and loop iteration status for all processed steps.				

Further information about choosing a logging level is found in "Improving Process Logging Performance" on page 95.

Improving Process Logging Performance

To improve process logging performance, consider the following:

- Choose **2 - Errors only**, **3 - Process only**, or **4 - Process and start steps** as your logging level. Choose **5 - Process and all steps** only when you need ultimate quality of service.
- Store input pipelines only when absolutely necessary. It is usually sufficient to store pipelines for failed steps only. Remove all unnecessary data from pipelines to minimize the volume of data to store.
- For process steps that run services, there are two areas in which you could inadvertently log the same information twice:
  - Process Engines can write start and successful completion or failure log entries for process steps that run services. Services can write log entries that convey the same information.
  - Process Engines can store input pipelines for services that are run by process steps. Services can also log input pipelines.

  Coordinate your logging for these services to avoid logging the same information twice.

  Note:
  When coordinating logging, consider that when a service is run by a process step, that service is actually called by a wrapper service, making it a nested service (as opposed to a top-level service).

For instructions on setting up service logging, and for complete information on logging in general, see *webMethods Audit Logging Guide*.

Configuring Logging Settings for a Process Model Version

To configure the logging settings for a process model version:

1. In My webMethods click **Navigate > Applications > Administration > Business > Business Processes**.

2. Find the process model version to work with and then click **Edit** to open the Edit Process page.

3. In the Process Information Administration window, click the **Process Settings** tab, and do any or all of the following:

   - Select one of the available **Logging Level** values, as described in "About Process Model Logging Levels" on page 94.

     Important:
     A **Minimum Logging Level** setting is specified in Designer for each process model version. This setting in Designer controls the lowest logging level that you can set in Monitor.

   - Select the **Log Transitions** check box if you want to log process transitions for display in the process diagram. This requires a logging level that enables you to log transitions

   - Select the **Diagnostic Logging** check box if you want to log messages from a process instance to a separate log file for diagnostic purposes. For more information, see "Enabling and Disabling Process Instance Diagnostic Logging" on page 96, below.

4. Click **Save**.

Enabling and Disabling Process Instance Diagnostic Logging

Log messages from individual process instances are always sent to the Integration Server server.log file. However, these messages are mixed together with Process Engine messages as well as messages from other process instances, so it can be hard to find the specific messages you are looking for.

The Process Engine also supports process instance diagnostic logging for individual process models. When you enable process instance diagnostic logging for a process model, the process instance log messages are sent to both the server.log file and to a separate process instance log file. You can then access the process instance log file to see the messages from an individual process instance.

To enable or disable process model instance diagnostic logging

1. In My webMethods: **Navigate > Applications > Administration > Business > Business Processes**.

2. Search for the process model you want to enable or disable. For instructions, see "Finding Process Models Using Keywords" on page 86 and "Finding Process Models Using an Advanced Search" on page 87.

3. In the search results, click **Edit** for the webMethods-executed model version or externally executed process model that you want to enable for analysis.

> **Note:**
> Enabling this option increases the processing overhead for all instances of this process model, which may have an impact on performance. You are advised to disable this option as soon as you have completed your diagnostic activities.

4. In the **Diagnostic Settings** window:
   - Select the **Diagnostic Logging Enabled** check box to enable process instance logging.
   - Clear the **Diagnostic Logging Enabled** check box to disable process instance logging.
5. Click **Save**.

See the following topic for information about viewing the diagnostic file.

Viewing a Process Instance Diagnostic Logging File

When a process model is enabled for process instance diagnostic logging, a log file is created in the directory *Software AG_directory* /IntegrationServer/serverName/instances/*instance_name*/packages/WmPRT/log, with a file name of *processInstanceID*.log. If the process is running on multiple servers (for example, in a distributed environment), a diagnostic log file is created *on each server that runs the process*. In this case, you must view all of the instance logs to get a complete picture of the instance activity. The log file name is the same on each server.

You can view the file in any text editor, or you can dynamically monitor the file in a command session using the `tail` command. The `tail` command is available on all Linux and UNIX systems, and on some Windows systems. If your Windows system does not offer the `tail` command, you can download it from the following locations:

- As part of the Windows Server 2003 Resource Kit Tools available from Microsoft.
- As an executable from Sourceforge.

Message entries are formatted as follows:

[*timestamp messageID threadID*] *processInstanceID:iteration stepID message*

All messages that are not pertinent to the process instance (for example, correlation of an incoming document) are sent to the server.log file.

Enabling and Disabling Process Model Versions

You can enable or disable webMethods-executed process model versions for execution. Only one version of a process model can be enabled at any time.

When you initially build and upload a process model from Designer, the version is disabled by default. However, you can specify that the process model is automatically enabled during the build and upload process in Designer. Otherwise, you must enable the process in Monitor.

When a model version is disabled, Process Engine will not use the model version for new process instances. If no versions of a process model are enabled, the Process Engine will not start any process instances of the process model.

Important:
If all versions of a model that is called by another process are disabled, the parent process will fail at the step that calls the disabled process model.

When you enable a model version, if another version was previously enabled, Monitor disables it, and Process Engine uses the newly enabled version when starting new process instances. When you enable a new version of a process model, you can optionally upgrade all running process instances that use the model, so that they immediately start using the newer version with the next step to be executed. For additional information, see "Updating a Process Instance to a New Model Version" on page 76.

To enable or disable a webMethods-executed model version

1. In My webMethods: **Navigate > Applications > Administration > Business > Business Processes**.

2. Search for the model version to enable or disable. In the search results, click **Edit** for the model version to open the Edit Process page.

3. Click the **Process Details** tab and do either of the following:

   - Select the **Execution Enabled** check box to enable the model version.

     Monitor asks if you want to upgrade running process instances to this version of the process model so that they immediately start using the newly enabled version with the next step to be executed. If you want to upgrade *all* running instances of the process model to use the new model version, click **Yes**.

     If you want to upgrade only individual process instances, you can do so after you enable the model version, when you are viewing the details for a process instance. For more information, see "Updating a Process Instance to a New Model Version" on page 76.

   - Clear the **Execution Enabled** check box to disable the model version.

4. Click **Save**.

Determining the Enabled Version of a Process Model

You can have multiple versions of a webMethods-executed process model; however, only one version can be enabled at a time. To determine which version is enabled (if any), perform the following procedure.

To determine the enabled version of a webMethods-executed process model

1. In My webMethods: **Navigate > Applications > Administration > Business > Business Processes**.

2. Click the **Advanced** search tab.

3. In the **Process Name** list, select the name of the webMethods-executed process model you want to examine.

4. Click **Search**. In the search results, the enabled version (if any) has a ✓ check mark in the **Execution Enabled** column.

Enabling or Disabling a Process Model for Analysis

Before you can enable a process model for analysis, Optimize for Process must be installed. You can only enable the version of a process model that is execution-enabled.

Note:
Externally-executed process models are always enabled for analysis. Integration processes cannot be enabled for execution.

If a process model is enabled for analysis, when a running process instance of the model version is executing, Optimize for Process collects metrics about the running process instance, such as how long it took the process instance to run, how long it took each step in the process to run, the number of errors that occurred. You can then view the collected metrics on the **Navigate > Applications > Monitoring > Business > Process Overview** page. For more information, see the PDF publication *webMethods Optimize User's Guide*.

To enable or disable a process model version for analysis

1. In My webMethods: **Navigate > Applications > Administration > Business > Business Processes**.

2. Search for the process model you want to enable or disable.

3. In the search results, click **Edit** to open the Edit Process page for the webMethods-executed model version or externally executed process model that you want to work with.

4. Click the **Process Details** tab and do either of the following:

   - Select the **Analysis Enabled** check box to enable a process model version for analysis.
   - Clear the **Analysis Enabled** check box to disable a process model version for analysis.

5. Click **Save**.

Enabling and Disabling EDA Event Emission

You can enable and disable the emission of predefined Process Engine EDA events for individual process models on an event-by-event basis. When event type emission is disabled, no predefined events are emitted by the model for the disabled event type.

EDA event emission for all predefined event types is disabled by default. To enable event emission, you must manually enable the predefined EDA events you want to emit for each individual process model. The event emission enablement settings are maintained with the process through the build and upload procedure. They are part of the process asset description, and are sent with the process when deployed with webMethods Deployer.

Note:
Disabling Process Engine EDA events has no effect on custom event types applied to the process model or steps within it.

In addition to enabling EDA event emission in Monitor as described below, you can also enable events in Software AG Designer. For more information about enabling and disabling EDA event emission in Designer, and about EDA events in general, see *webMethods BPM Process Development Help*.

Note:
Before you modify process model EDA event settings, be aware of the interaction of these settings between Designer and the Edit Process page. For more information, see "About Synchronizing Stages and Events with Software AG Designer " on page 105.

To enable or disable EDA event emission in Monitor

1. On the Business Processes page, locate the process model you want to work with and click **Edit**.

2. In the Process Stages and EDA Events window, click the **Events** tab. The following predefined EDA event types appear in the **Available EDA Events** list:

   - **Process Instance** controls whether to emit an event when the status of a process instance changes, for example from started to completed.
   - **Process Instance Log Message** controls whether to emit an event when a message is logged from a process model step, for example, warnings, messages, and errors.
   - **Process Instance Error** controls whether to emit an event when an error occurs in a process instance.
   - **Process Instance Log Custom ID** controls whether to emit an event when the service pub.prt.log:logCustomId is invoked in a process model step to associate a custom ID with the process model instance.
   - **Step Instance** controls whether to emit an event when the status of a step changes, for example, from started to waiting.

- **Step Loop Instance** controls whether to emit an event when a step loop starts or completes.
- **Step Instance Transition** controls whether to emit an event when one step transitions to another step.
- **Step Instance Error** controls whether to emit an event when an error occurs during a process step.

Note:
For detailed information about each of these EDA event types, see the chapter "EDA Event Types" in *webMethods BPM Process Development Help*.

3. Select the EDA event you want to enable or disable for emission.

4. Use the available buttons to move EDA events between the **Available EDA Events** list and the **Selected EDA Events** list:

   - Click or to move a single event selection or a multiple, non-contiguous selection made by selecting events with the CTRL key pressed.
   - Click or to move all EDA events in a list to the other list.

5. Click **Save**.

Enabling and Disabling a Step for Resubmission

A process instance with a "Completed" status may be resubmitted from any step in the instance, provided you have first enabled the step for resubmission before the process instance began executing. This setting has no effect on process instances that are currently running or that have already stopped running.

In addition to enabling steps for resubmission in Monitor as described below, you can also enable steps for resubmission in Software AG Designer. For more information about enabling and disabling steps for resubmission in Designer, see *webMethods BPM Process Development Help*.

Note:
If you want to be able to resubmit process instances from Monitor at certain steps, you must set the process model logging level to a level that will log the input pipelines for those steps. For more information, see "About Process Model Logging Levels" on page 94.

To enable or disable a step for resubmission in Monitor

1. On the Business Processes page, search for the process model that contains the steps you want to enable.

2. In the search results, locate the process model you want to work with and click **Edit**.

3. On the Edit Process page, click the **Step Settings** tab, and then do either of the following:

- Select the corresponding check box for a step you want to enable for resubmission.
- Clear the corresponding check box for a step to disable resubmission capability.

Note:
When a step is enabled for resubmission, the pipeline data for that step is saved. Extensive use of resubmittal enablement (for example, enabling all steps for all process models for resubmittal) may result in a reduction in performance.

4. Click **Save**.

Working with Stages and Milestones

You can create, delete, and modify process stages in Monitor. For additional information, see "Stages Tab" on page 92.

Note:
You must have a BPM or a BPM and BAM server environment selected in the Server list at the top of the Process Instances page before you can add stages in Monitor.

Adding a Stage

Note:
If you leave the **Stages** tab while adding a stage and before you have clicked **Save**, your changes will be discarded.

Note:
Before you make modifications to process model stage settings, be aware of the interaction of these settings between Designer and webMethods Monitor. For more information, see "About Synchronizing Stages and Events with Software AG Designer " on page 105.

To add a stage

1. On the Business Processes page, locate the process model that you want to work with and click **Edit**.

2. In the Process Stages and EDA Events window, click the **Stages** tab.

3. Click **Add Stage**. A new row appears in the stage list, populated with default information.

4. Configure the following fields to define the stage.

   Note:
   Any data entry validation errors are displayed within the stage row.

The table below explains how to configure the fields.

Column	Description
Name	Type a name for the stage. **Note:** The **Name** is not editable after you click **Save**. If you want to rename a stage, you must delete it and then recreate it with the new name. **Note:** There is an 80-character limit for the stage name when double-byte characters are used in an IBM DB2 database. If you are not using DB2, or if your characters are single byte, then the stage name is limited to 255 characters.
Description	Optional. Type a description of the stage.
Start Milestone	Click the list and select a milestone. Optionally, you can type characters in the text box to filter the list. The **Start Milestone** and **End Milestone** selections must be different. Click the list to the right of the milestone selection, and click **Start** or **Complete** to specify the start or the completion of the selected milestone.
End Milestone	Click the list and select a milestone. Optionally, you can type characters in the text box to filter the list. The **Start Milestone** and **End Milestone** selections must be different. Click the list to the right of the milestone selection, and click **Start** or **Complete** to specify the start or the completion of the selected milestone.

Condition	■ Select **<** (less than) or **>** (greater than). Default is **<**.
	■ Enter a positive whole number. The maximum supported values are as follows:
	  ■ 2,777,777 hours
	  ■ 166,666,666 minutes
	  ■ 9,999,999,999 seconds
	  ■ 9,999,999,999,999 milliseconds
	■ Default is **1**.
	■ Select **hours**, **minutes**, **seconds**, or **milliseconds**. Default is **hours**.
	
	The result is a condition. If the condition specifies <, then the stage is breached when the cycle time exceeds the specified time period. If the condition specifies >, then the stage is breached when the cycle time is less than the specified time period. For example:

Column	Description
	`< 1 hours` means that the stage must complete in less than 1 hour or a `ProcessStageBreached` event will be emitted.
Stop Tracking On Breach	Stops stage processing for all remaining stages in the process instance when a stage breach occurs in this stage, and only one stage breached EDA event is emitted. Remaining stages are not tracked and will be shown as Incomplete in Monitor. The check box is cleared by default.

5. Click **Save**.

For information about synchronizing the new stage with the process model in Designer, see "About Synchronizing Stages and Events with Software AG Designer " on page 105.

Modifying a Stage

Note:
If you leave the **Stages** tab while modifying a stage and before you have clicked **Save**, your changes will be discarded.

Note:
Before you make modifications to process model stage settings, be aware of the interaction of these settings between Designer and webMethods Monitor. For more information, see "About Synchronizing Stages and Events with Software AG Designer " on page 105.

You cannot modify a stage name. If you want to rename a stage, you must delete it and then recreate it with the new name.

Otherwise, all other stage and milestone information can be modified as described in "Adding a Stage" on page 102.

Deleting a Stage

Note:
If you leave the **Stages** tab after deleting a stage and before you have clicked **Save**, the deletion will be discarded.

Note:
Before you make modifications to process model stage settings, be aware of the interaction of these settings between Designer and webMethods Monitor. For more information, see "About Synchronizing Stages and Events with Software AG Designer " on page 105.

To delete a stage

1. On the Business Processes page, locate the process model that you want to work with and click **Edit**.

2. In the Process Stages and EDA Events window, click the **Stages** tab.

3. Click the option button next to stage name for the stage you want to delete. To clear your selection, click the option button again

4. Click **Delete**.

5. Click **Save**.

Viewing Stages in the Process Diagram

You can display a stage's start milestone and stage end milestone in the Process Diagram window. Only one stage can be displayed in the Process Diagram at any time.

To view a stage in the process diagram, select any row in the stage list as follows:

- Click the option button next to stage name to select a stage row. The icon becomes active and the milestone pins are rendered in diagram.
- To disable the pin display in the diagram for a stage row and keep the row selected, click the icon.
- To disable the pin display in the diagram for a stage row and clear the row selection, click the option button .
- For a selected stage row with pin display disabled, click to display the pins again.
- Click the option button of another stage row to switch stage selection. The milestone pins on diagram are updated for the newly selected stage.

About Synchronizing Stages and Events with Software AG Designer

You can create, modify, and delete stages, and enable/disable EDA events, in two locations:

- On the Edit Process page in My webMethods.
- In Software AG Designer, as part of the process model development functionality. For more information, see the *webMethods BPM Process Development Help*.

In both cases, any changes to the stage or event definitions in a process model can be saved to the Process Audit database. The saved changes overwrite whatever previous stage information was present in the database. As a best practice, you should ensure that your process model stage and event settings are always synchronized between the two locations. For the Edit Process page, the following conditions apply:

- When you open the Edit Process page to view the process model settings in the Process Stages and EDA Events window, the stage and event settings saved in the Process Audit database are retrieved and displayed.
- When you save the process model stage and event settings on Edit Process page, the settings are written to the Process Audit database, overwriting whatever stage and event settings are stored there.
- Designer writes stage and event settings to the database when the process model is built and uploaded, overwriting whatever stage and event settings are stored there.

Therefore, be aware that it is possible for someone working in Designer to modify the database stage or event settings after you have opened Edit Process page. In this case, when you click **Save** on the Edit Process page, the stage and event settings in the Edit Process page will overwrite the settings written to the database by the person working in Designer.

Similarly, if you save stage and event settings to the database from the Edit Process page, it is possible for a Designer user to overwrite those settings at any time.

A **Synchronize** button is available in Designer for both stages and events, enabling the Designer user to apply the current database stage settings to the process model in Designer. You are advised to establish procedures to ensure that stage settings event settings are managed without conflict between Designer users and Edit Process page users.

To help ensure that you are working with the latest settings, you are advised to either refresh the Edit Process page or re-open the Edit Process page immediately before modifying and saving stage and event settings.

Deleting Unused Process Models

If a process model has not been used, you can delete information about the process model from the Process Audit Log database and the Monitor display. Before you can delete a webMethods-executed process model version, you must first disable that process model (see "Enabling and Disabling Process Model Versions" on page 97). You can delete any type of process (webMethods-executed, externally executed, or integration process) as long as that process has never been used for a process instance.

Note:
To delete a process model you must have My webMethods Server administrator privileges.

To delete unused process models

1. On the Business Processes page, search for the process model you want to delete.
2. In the search results, select a check box for each process model you want to delete. To select all models, click in the table header.
3. Click **Delete**.

Upgrading Process Models

If you have processes that were created in an earlier version of Designer, you can upgrade a single process, or multiple processes at one time. Upgrading a process may be necessary to take advantage of newer features in Monitor.

- **To upgrade an individual process**. Import the process model into Designer version 9.6 or later, and then save, build, and upload it. For more information, see the topic, "About Importing and Exporting Processes" in *webMethods BPM Process Development Help*.
- **To migrate multiple process models**. To migrate from an existing Integration Server installation, see the PDF publication, *Upgrading Software AG Products*

About Process Model Rendering

Beginning with version 9.6, Monitor uses Google Web Toolkit (GWT) to render process model diagrams, both for models created in Designer 9.6 and later, and for process model diagrams that were created in Designer from version 8.2.2 to version 9.5.1. There may be minor cosmetic differences between how process model diagrams are rendered in Flash or GWT, but there are no significant functional differences.

To ensure that the browser loads the GWT rendering component properly, before viewing models with GWT, clear your browser's cache and history.

Minor Functional Differences Between GWT and Flash

When you view a process model that has been upgraded from an earlier version in GWT, note the following differences in functionality:

- Models displayed with GWT do not support the ability to grab and drag the image as is possible in a Flash-based display.
- In GWT, use the slider to change the zoom and focus of the diagram.

Using Internet Explorer

To use Internet Explorer version 9 and later to view process models rendered with GWT, you must configure the compatibility settings in My webMethods Server for your version of the browser, as well as modify certain browser settings. For more information, see "Modifying the Default Header Rule for Internet Explorer" on page 107.

For general information on configuring response headers, see the section, "Working with Response Header Rules" in *Administering My webMethods Server*.

Modifying the Default Header Rule for Internet Explorer

To be able to use Internet Explorer to view a process diagram on the Process Details page, you must modify the default header rule in My webMethods Server as well as certain browser settings.

To apply modifications for using Internet Explorer version 9 and later

1. Log in to My webMethods Server as sysadmin.
2. In My webMethods Server: **Administrative Folders > Administration Dashboard > User Interface > Manage Response Header Rules**.
3. Click the link for the default rule **IE - parameter for compatibility mode**.
4. Make sure the **Is Enabled** check box is selected.
5. In the **Result** area, click the **X-UA-Compatible** parameter and then click **Update**.
6. In the Response Header dialog box, change the **Value** field to `IE=edge`.
7. Click **Apply**, then click **Update Rule**.
8. In Internet Explorer: **Tools > Internet Options > Security > Trusted Sites.**
9. Click **Sites** and add the URL for the My webMethods Server host to the list of trusted sites.
10. Click the **Advanced** tab in Internet Options.
11. In the **Accelerated graphics** list, select the check box for **Use Software rendering instead of GPU rendering**.
12. Click **OK**.

7 Defining, Executing, and Logging Integration Processes

Logging an Integration Process

You determine the data that is logged for an integration process by invoking built-in services that are provided with Monitor in the pub.monitor.integrationProcessLogging folder of the WmMonitor package. When data is logged for an integration process, you can monitor and track its progress using Monitor pages in My webMethods. For complete information about the services in the pub.monitor.integrationProcessLogging folder, see *webMethods Monitor Built-In Services Reference*.

Defining an Integration Process

This example shows how you can use the pub.monitor.integrationProcessLogging services to define an integration process, the steps in the process, and the fields to log for the process.

The following table describes how to define an integration process.

Flow	Description
1.	To add a definition for the integration process, invoke the pub.monitor.integrationProcessLogging:createProcessMetadata service. The process definition is comparable to an entry for a process model that was created with Designer.
2.	After you add a process definition, you can add steps to the process. To add a definition for a step, invoke the pub.monitor.integrationProcessLogging:createStepMetadata service.
3.	After you add a step definition, you can add definitions for one or more logged fields that are associated with the step. To add a logged field and associate it with a step, invoke the pub.monitor.integrationProcessLogging:createLoggedFieldMetadata service. Logged field definitions are comparable to custom data in a process model that was created with Designer. In the example above, two logged fields are defined for and associated with the first step of the integration process.
4.	Continue adding steps and logged fields for the steps. In the example above, a second and third step are added. One logged field is defined for and associated with the third step.

Executing an Integration Process

After you define an integration process, including process steps and logged fields, you can create running instances of the integration process.

The following example shows how you might use the pub.monitor.integrationProcessLogging services to create and execute an instance of an integration process that you have already defined.

> **Note:**
> Although you can view information about integration processes using Monitor, integration processes are not eligible to be stopped, suspended, resumed, or resubmitted.

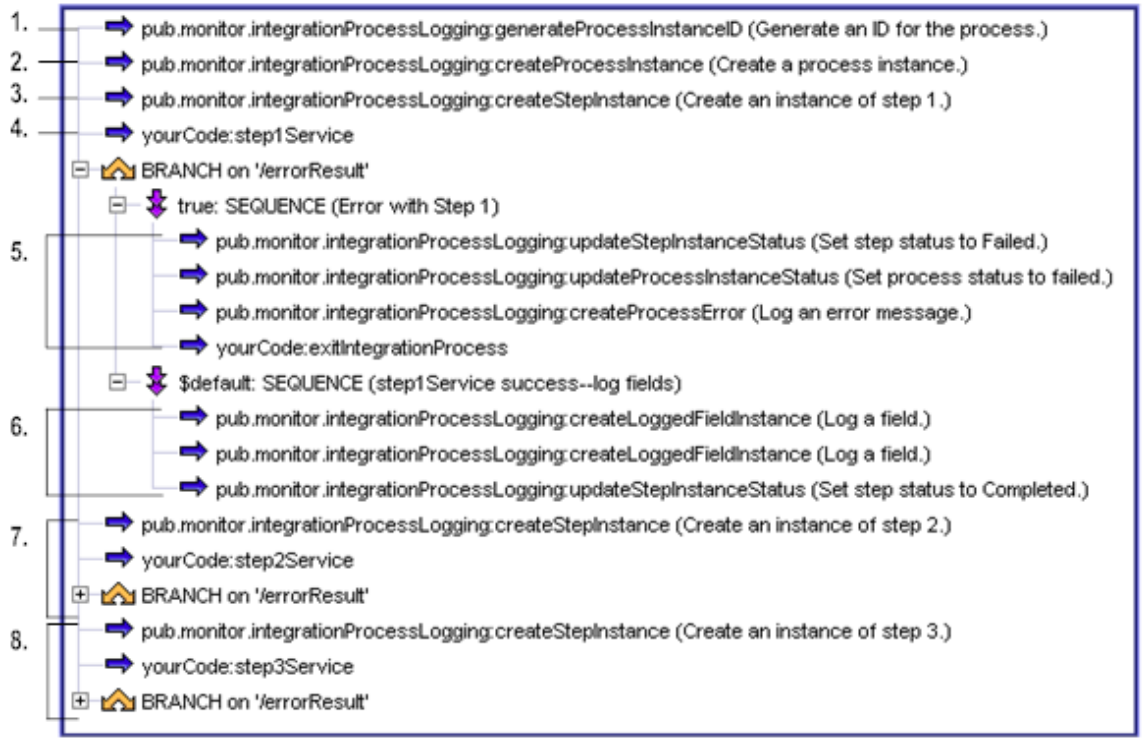

The following table describes how to create and execute an instance of an integration process.

Flow	Description
1.	Each process requires a unique process instance ID. You can create one on your own or use the pub.monitor.integrationProcessLogging:generateProcessInstanceID service to generate a unique process instance ID. **Note:** This service does not add data to the Process Audit Log database. That is, it does not start a new instance of the integration process.

Flow	Description
2.	To log a running instance of an integration process, create a process instance. To create a new process instance, invoke the pub.monitor.integrationProcessLogging:createProcessInstance service. As input to this service, you must identify the process definition for which you are creating the instance. The process definition must already exist in the Process Audit Log database, which you add using the pub.monitor.integrationProcessLogging:createProcessMetadata service, as shown in the sample in "Defining an Integration Process" on page 110. After you create the process instance, its process status is Started. Use the pub.monitor.integrationProcessLogging:updateProcessInstanceStatus service if you want to change the process status.
3.	To start logging information about the first step in the integration process, create an instance of the first step. To create a step instance, invoke the pub.monitor.integrationProcessLogging:createStepInstance service. As input to this service, you identify: ■ The process instance in which the step is running. Specify the same process instance ID that you specified as input to the pub.monitor.integrationProcessLogging:createProcessInstance service in the previous INVOKE flow operation. ■ The step definition for which you are creating the instance. The step definition must already exist in the Process Audit Log database, which you add using the pub.monitor.integrationProcessLogging:createStepMetadata service. After you create the step instance, the step status is Started. Use the pub.monitor.integrationProcessLogging:updateStepInstanceStatus service if you want to change the step status.
4.	After the step instance is logged, execute the service(s) for the first step of the integration process. The next flow operations are based on the outcome of this step.
5.	This series of flow operations shows the logic you might perform when a step fails: ■ The step status is set to Failed using the pub.monitor.integrationProcessLogging:updateStepInstanceStatus service. ■ The process status is set to Failed or Failed (Escalated) using the pub.monitor.integrationProcessLogging:updateProcessInstanceStatus service. ■ An error message is logged for the process using the pub.monitor.integrationProcessLogging:createProcessError service. ■ A service you create is invoked to exit the integration process.
6.	This series of flow operations shows the logic you might perform when a step is successful:

Flow	Description
	■ Values are set for logged fields that are associated with the step using the pub.monitor.integrationProcessLogging:createLoggedFieldInstance service. ■ The step status is set to Completed using the pub.monitor.integrationProcessLogging:updateStepInstanceStatus service.
7.	This series of flow operations handles the second step in the integration process. The flow operations are similar to those described above in steps 4, 5, and 6.
8.	This series of flow operations handles the third step in the integration process. The flow operations are similar to those described above in steps 4, 5, and 6.

8 Archiving or Deleting Data in an Archive Database

Overview of Archiving or Deleting Data in an Archive Database

Before you can archive or delete data, you must configure the archiving and deletion feature as described in "Audit Data Archiving and Deletion in Monitor " on page 29.

Software AG recommends that you remove audit data from the IS Core Audit Log and the Process Audit Log databases regularly to keep your logging at peak performance. To remove audit data, you can either archive or delete it, as follows:

- **Archive**. Moves audit data to the Archive database and then removes it from the source database.
- **Delete**. Removes audit data from the source database, but does not move it to any other location.

As an alternative to using an Archive database, you can store process audit log data in a separate partition. For information about configuring partitions, see "Overview of Using Partitions for Process Audit Log Data" on page 126.

After you archive or delete audit data, you can no longer view it in My webMethods. However, if you archive data, you can still query the Archive database using SQL. The following table lists the methods to archive or delete document, process, service, and Integration Server data from the IS Core Audit Log and Process Audit Log databases.

Method for archiving or deleting data	For more information, see
Using the Monitor user interface	"Archiving or Deleting Audit Data Using the Monitor User Interface" on page 116.
Using Monitor built-in services in the pub.monitor.archive folder	"Archiving or Deleting Audit Data Using Built-in Services" on page 119.
Scheduling Monitor archive built-in services in Integration Server	"Archiving or Deleting Audit Data Using Built-in Services" on page 119.
Running stored procedures from your Oracle, SQL Server, or DB2 database client	"Archiving or Deleting Audit Data Using Stored Procedures" on page 119.

Archiving or Deleting Audit Data Using the Monitor User Interface

You can archive or delete data from the Monitor user interface using this procedure.

To archive or delete audit data using the Monitor user interface

1. In My webMethods: **Navigate > Applications > Administration > Business > Data Management > Archive Audit Data**.

2. Specify how long to keep data in IS Core Audit Log and the Process Audit Log tables. Monitor archives or deletes data that is older than the retention period you specify.

The following table lists and describes the options.

Option	Description
Number of days to retain (ending with today)	Monitor keeps data for the number of days (including the current date) that you specify. For example, if you specify 15, Monitor keeps data that is 15 or fewer days old and archives/deletes data that is 16 or more days old.
Retention period start date (ending with today)	Date of the oldest data to keep. The period ends with and includes the current date. For example, if you specify 6/3/2012, Monitor keeps data from 6/3/2012 through the current date and archives/deletes any data logged before 6/3/2012. Use the calendar picker and hour and minute fields to set the date.

Note:
Processes and services have a start timestamp and an end timestamp; Monitor archives or deletes process based on the end timestamp. Documents and server data have a single timestamp, and Monitor archives or deletes them based on that timestamp.

3. In Data Types, specify which data to archive or delete, as follows.

   a. Select the check boxes corresponding to the types of data to archive or delete, as follows in the table below.

Select	To archive
Processes	For the process model or models selected, Monitor archives or deletes: ■ Process log entries, input pipelines, error data, and run-time values for user-specified input and output document fields. ■ Referenced processes, process-related service data (services, service error data, and user-defined messages). ■ Process control data (resubmit, suspend, and resume actions).
Services	Service log entries, input pipelines, error data, user-defined messages, and service control data (resubmit actions). **Note:**Monitor can archive user-defined messages for a service only if customized logging is set up for the service. That is, if service logging is globally enabled in Integration

Select	To archive
	Server, but customized logging is not set up for the service in Designer, then Monitor cannot archive user-defined messages written by the service.
Documents	Logged documents for all webMethods Broker (deprecated) clients and document control data (resubmit actions). If selected, Monitor archives or deletes all logged documents.
Server Data	Integration Server session and guaranteed delivery log entries, and error log entries that are not associated with logged processes, services, or documents (for example, errors that occur during startup or during the run of unlogged processes, services, activations, and documents). If selected, Monitor archives or deletes all server data.

b. To archive processes, select a process model name or **All** to archive all process models.

> **Note:**
> You can only archive by process model name when Monitor uses stored procedures to perform archiving. When Monitor is configured to use JDBC pools, all process models are archived.

c. Next, for processes or services, select which model status to archive or delete.

The following table lists the two options.

Option	Archives or deletes
Completed	Audit data for processes or services with a Completed status.
Completed-Failed	■ **For services**. Data with the status, Completed, Failed, Failed (Escalated), and Resubmitted. ■ **For processes.** Data with the status, Completed, Failed, Stopped, and Resubmitted.

> **Note:**
> You can only choose the status when archiving or deleting processes and services. Documents and Integration Server data do not have statuses and therefore cannot be archived selectively.

Monitor archives or deletes only data that matches all of the values that you specify.

4. In the **Archiving Batch Size** field, indicate the number of primary items and accompanying items to archive or delete at a time. For example, to archive or delete 100 processes and accompanying services, activity logs, and errors at a time, choose a number that takes the size of each record and other performance factors into consideration. If the record size is large,

consider reducing the batch size; if the record size is small, increasing the batch size might increase the speed of the archive or delete.

5. To archive or delete data, do one of the following:

   - Click **Archive and Delete** to copy the data from the IS Core Audit Log and Process Audit Log tables to the Archive database and then delete the data from the IS Core Audit Log and Process Audit Log tables.
   - Click **Delete only** to delete the data from the IS Core Audit Log and Process Audit Log tables.

Archiving or Deleting Audit Data Using Built-in Services

The pub.monitor.archive folder in the WmMonitor package contains services for archiving and deleting audit data. You can run the services in Designer, or you can use Integration Server Administrator to schedule the services to run automatically at specific times or intervals. Running the services regularly minimizes the time required to process the data.

To schedule a service for archiving or deleting, build a wrapper service that calls the service and sets the service's input parameters, then run the wrapper service as a scheduled task from Integration Server Administrator. The wrapper service executes for the first time immediately after you schedule the task. For instructions on scheduling services, see *webMethods Integration Server Administrator's Guide*.

The archive and delete services are located in the directory, *Integration Server_directory* \instances\instance_name\packages\WmMonitor\pub\ monitor\archive, are described in *webMethods Monitor Built-In Services Reference*.

Note:
For Oracle, SQL Server, MySQL, PostgreSQL, and DB2, the default values for all archive and delete service parameters are stored in OPERATION_PARAMETER in the Archive database. You can change the defaults in the table by running the pub.monitor.archive:setOperationParameters service. You can override the defaults for specific archive or delete actions by specifying values on the relevant parameters when you run the archive and delete services.

Archiving or Deleting Audit Data Using Stored Procedures

When you installed the Archive database, you also installed stored procedures for archiving or deleting data between databases installed in Oracle, SQL Server, or DB2. You access the stored procedures for archiving through your database application.

The stored procedures for archiving data to the Archive database from the IS Core Audit Log database, the Process Audit Log database, or both are listed in the following table.

Procedure	Description
Oracle: **DOCUMENT_ARCHIVE. START_DOCUMENTARCHIVE** **SQL Server and DB2:** **DOCUMENT_ARCHIVE_ START_DOCUMENTARCHIVE** **MySQL and PostgreSQL:** **DOCUMENT_ARCHIVE_ START_DOCUMENTARCHIVE**	In IS Core Audit Log, archives or deletes documents logged for webMethods Broker (deprecated) clients from the WMDOCUMENT table. In Process Audit Log, archives and deletes document control data (such as resubmit actions) from the PRA_PROCESS_ACTION table.
Oracle: PROCESS_ARCHIVE. START_PROCESSARCHIVE **SQL Server and DB2:** PROCESS_ARCHIVE_ START_PROCESSARCHIVE **MySQL and PostgreSQL:** PROCESS_ARCHIVE_ START_PROCESSARCHIVE	In IS Core Audit Log, archives or deletes process-related service data from the tables: WMSERVICEACTIVITYLOG and WMERROR In Process Audit Log, archives or deletes process and control data (such as resubmit actions) from the tables: ■ PRA_PROCESS_CUSTOM ■ PRA_STEP_LOGGED_FIELD ■ PRA_STEP_LOOP_LOGGED_FIELD ■ PRA_PROCESS ■ PRA_PROCESS_ACTION ■ PRA_PROCESS_RECENT ■ PRA_PROCESS_STEP ■ PRA_STEP_TRANSITION ■ WMCUSTOMFIELDDEFINITION ■ WMPROCESSDEFINITION ■ WMPROCESSIMAGE ■ WMSTEPDEFINITION ■ WMSTEPTRANSITIONDEFINTION
Oracle: SERVER_ARCHIVE. START_SERVERARCHIVE	In IS Core Audit Log, archives or deletes Integration Server data from the tables: WMERROR, WMSESSION, WMTXIN, and WMTXOUT

Procedure	Description
SQL Server and DB2: SERVER_ARCHIVE_START_SERVERARCHIVE **MySQL and PostgreSQL:** SERVER_ARCHIVE_START_SERVERARCHIVE	Integration Server data consists of session and guaranteed delivery log entries, and error log entries that are not associated with logged processes, services, or documents (for example, errors that occur during startup or during the run of unlogged processes, services, activations, and documents). **Note:** Archiving Integration Server data archives only top-level errors. To archive lower-level errors associated with services or processes, you must first archive those services or processes.
Oracle: SERVICE_ARCHIVE.START_SERVICEARCHIVE **SQL Server and DB2:** SERVICE_ARCHIVE_START_SERVICEARCHIVE **MySQL and PostgreSQL:** SERVICE_ARCHIVE_START_SERVICEARCHIVE	In IS Core Audit Log, archives or deletes service log entries, input pipelines, error data, and user-defined messages from the following tables: ■ WMSERVICE ■ WMSERVICE_MIN_MAX ■ WMERRORTBL ■ WMSERVICEACTIVITYLOGTBL ■ WMSERVICEASSOCTBL ■ WMSERVICECUSTOMFLDSTBL In Process Audit Log, archives or deletes service control data (such as resubmit actions) from the WMCONTROL table.

Note:
Indexes for the archive schema are removed when archiving with stored procedures.

Parameters

The following table lists and describes the parameters.

p_retaindays **String** Number of days to keep the indicated type of data in the source database, ending with and including the current date.

Suppose the current date is September 30, to archive or delete data for services that finished running more than 15 days ago (that is, before September 15), you would specify this parameter as 15.

Supply either *p_retaindays* or *p_retaindate*. Do not supply both.

p_retaindate **Long** Start date for the period to keep the indicated type of data in the source database. The period ends with and includes the current date. *p_retaindate* is in epoch time (milliseconds since midnight, 01/01/1970).

Supply either *p_retaindate* or *p_retaindays*. Do not supply both.

p_action **String** Indicates whether to archive or delete the indicated type of data from the source database.

Set to...	To...
`ARCHIVE`	Default. Copy the data from the source to the Archive database and delete the data in the source database.
`DELETE`	Delete the data from the source database.

p_status **String** Used by the PROCESSARCHIVE or SERVICEARCHIVE procedure, indicates which data to archive or delete based on the specified status. Process and service data with a status other than those specified is retained in the source database.

Code	Archive or Delete data for
2	Processes or services with a Completed status.
4	Processes or services with a Failed status.
1024	Processes with a Stopped status.
32768	Processes and services with a Resubmitted status.
32776	Services with an Activity status.

p_batchsize **String** Number of primary items and accompanying items to archive or delete at a time. To archive or delete 100 processes and accompanying services, activity logs, and errors at a time, choose a number that takes the size of each record and other performance factors into consideration. If the record size is large, consider reducing the batch size; if the record size is small, increasing the batch size may increase the speed of the archive or delete.

p_modelid **String array** Used by the PROCESSARCHIVE procedure only. Complete model ID of the model for the process instance(s) you want to archive. Use a comma to separate multiple *modelid* string values. You can retrieve model IDs by invoking the pub.monitor.process.instance:getProcessList service and using the value returned in the processNames/PROCESSKEY output parameter. If null, all process models are archived.

Deleting User Information From Database Tables

For organizations that have to strictly comply with the GDPR regulations, use the SQL scripts recommended below to delete specific user information from the Monitor database tables.

To delete user information from database tables you can use a simple SQL statement, for example, update <TableName> set <ColumnName> = 'Unknown' where <ColumnName> = <Old Value>.

The following database tables contain specific user information in the specified columns:

Table	Column
PRA_PROCESS_ACTION	USERNAME
PRA_PROCESS_STEP	USERNAME
WMCONTROL	USERNAME
WMPROCESSDEFINITION	CREATEDBY
WMPROCESSTASK	CREATEDBY, UPDATEDBY
WMPROCESSTASKUSER	TASKUSER

Viewing the Results of an Archive Operation

If you installed the Archive database in Oracle, SQL Server, or DB2, you can view the progress of an archive operation on the **Navigate > Applications > Administration > Business > Data Management > Archive History** page in My webMethods.

9 Archiving Data Using Partitions

Overview of Using Partitions for Process Audit Log Data

In the default stored procedure method of archiving, the stored procedures search for the records to archive (or delete) row by row, based on the input criteria. This is generally not a problem for smaller databases, but the process can be time-consuming for a large database with many audit records to be archived.

As an alternative to using stored procedures to archive and delete Process Audit data, you can use database partitioning, an option that greatly decreases the time required to archive and delete data. The database partitions themselves are a standard feature of each database vendor, although you may need to purchase a separate partitioning license from your database vendor if you do not already have one. Monitor provides Oracle, Microsoft SQl, and IBM DB2 database scripts to configure and manage your partitions.

Note:
Partition archiving support is only provided for Process Audit data. You must continue to use stored procedures for all other audit data.

To archive or delete audit data with partitioning, the first step is to define the needed partitions. Then, when you archive a partition, the script moves it from your active Process Audit database to the archive Process Audit database, and operation that typically takes seconds to complete, compared with archiving by stored procedures, which can take hours. To delete data, you drop the relevant partition.

Each partition stores only those records that fall within the partition's date range based on the column, ATRESTTIMESTAMP. When creating partitions, adhere to the following rules:

- Create as many partitions as you need.
- Configure each partition with a non-overlapping date range.
- Define every Process Audit database table with identical partitions.

Monitor stores process instances that are still running in a partition named WM_FUTURE (Oracle and DB2) or partition 1 (MS SQL). As audit data is written to the Process Audit tables, Monitor automatically writes audit data to this partition. This partition stores all audit data that is not yet considered complete. When a process instance completes, Monitor updates the ATRESTTIMESTAMP with the final completion date and moves all associated audit entries to the appropriate partition. This guarantees that all related audit data for a process instance exists in the same partition.

Configuring Partitions

You can define as few or as many partitions as you require based on your data volume and archiving needs.

Note:
If you are have implemented partitioning in a previous version of Monitor and plan to upgrade to the current version, you must first migrate your ProcessAudit schema to the current version

using the Database Component Configurator tool. You can then use the provided partition scripts to partition the ProcessAudit schema.

To create and manage partitions for Process Audit Log data, refer the readme.txt file for your database in the following directories:

- For Oracle: < *Software AG_directory* >\common\db\scripts\oracle\processaudit\75\partition_support
- For IBM DB2: < *Software AG_directory* >\common\db\scripts\db2\processaudit\75\partition_support
- For Microsoft SQL: < *Software AG_directory* >\common\db\scripts\mssql\processaudit\75\partition_support

10 Using Mobile Monitor

Understanding Mobile Monitor

Mobile Monitor is an application for mobile devices and tablets that provides an alternative for using the Monitor user interface, providing real-time information about the status of executed business processes.

Mobile Monitor is only available to users on Integration Server version 9.5.1 and above and can only be used with Integration Servers that are exposed to the Internet either through a DMZ or a VPN connection.

The sidebar in Mobile Monitor provides quick navigation to view the number of business processes executed, the total number of process instances, the number of process instances by status, the number of failed instances (in **Notifications**), and the active server. The icons at the top and bottom of the screen provide navigation to the app settings, server configuration, the list of business processes, and the search feature.

Using MobileMonitor, you can do the following:

- **View business processes executed**. The Business Processes screen displays a list of the business processes executed. Individual status icons display a count of the number of processes in running, completed, failed, and resubmitted status.

  Tap the **List** icon at the bottom of the screen or tap **Business Processes** in the sidebar to see a list of all the business processes and the count of processes by status.

- **View business process model details**. The Business Process Details screen displays the process model diagram and details about the process model, such as whether the model is enabled, the user who created it, the model description, and the date the model was deployed.

  Click the arrow beside a business process to view the Business Process Details screen.

- **View a list of executed business process instances**. The Process Instances screen displays a list of executed process instances for each business process and a summary about the process instance. Process instances are listed in order by the audittimestamp (start date and time that the process was executed). Individual status icons show a count of the process instances in running, completed, failed, and resubmitted status. You can configure up to three fields to display about the process instance. For more information about configuring the fields on the Process Instances screen, see "Configuring Mobile Monitor " on page 132.

  Tap the process instance on the Business Processes screen or **Process Instances** in the sidebar to see a list of process instances, sorted by audittimestamp, and the count of process instances by status.

- **View the details of a business process instance**. The Process Instance Details screen displays the details of a business process instance, including the process model diagram, the version number, the date and time that the model was executed, the date the model was last updated, the custom ID, the iteration of the instance, the status of the instance, and how long the instance has been running or ran before it completed (duration).

  Tap the process instance on the Process Instances screen to view the instance details and process model diagram.

- **Pin a process instance**. Use the pin feature to keep a business process at the top of the list of business processes. This is useful when you want to monitor a specific process.

  On the Business Processes screen, tap the **Pin** icon to pin a business process.

- **View failed processes**. The **Notifications** icon on the Business Processes screen displays a count of the business processes that failed. You can configure how often Mobile Monitor polls for notifications, as often as every minute or hourly. Or, you can disable notification polling.

  Tap the **Notifications** icon on the Business Processes screen or **Notifications** in the sidebar to display the list of failed business processes on the Notifications screen. For more information about notifications, see "About Notifications" on page 131.

- **Manage server connections**. The Server Settings screen displays the configured servers and the server status. Mobile Monitor pulls business process information from one server at a time, known as the *active server*. You can add and delete servers and specify which server is the active server.

  Tap the **Servers** icon at the bottom of the screen or tap **Servers** in the sidebar to view the Server Settings screen. For more information about configuring servers, see "Configuring Mobile Monitor " on page 132.

- **Search for business processes and filter the list of process instances**. You can search for a specific business process by name. Or, you can limit the list of process instances using a filter. In the process instance filter, you can specify an execution date range, a custom process ID, and a process instance status. The process instance status can be specified as started, completed, failed, stopped, revised, failed/escalated, suspended, resumed, resubmitted, or all, to include instances with any status.

  To search for a specific business process, tap the **Search** icon on the Business Processes screen. To filter the list of process instances, tap the **Search** icon on the Process Instances screen or tap **Search Instances** in the sidebar.

- **Configure application settings**. The Application Settings screen displays all the configuration options available in the Mobile Monitor app. You can configure the fields that display on the Process Instances screen and the Notifications screen. You can also set how often the app polls for notifications and whether logging is enabled.

  Tap the **Gear** icon at the bottom of the screen to access the Application Settings screen. For more information about configuration settings, see "Configuring Mobile Monitor " on page 132.

About Notifications

Mobile Monitor notifies you when a process instance fails by adding a count (alert) to the **Notifications** icon on the Business Processes screen. If there are any failed processes, the **Notification** icon displays a count of the number of failed processes, up to a maximum count of 100. By default, the app polls for new notifications every minute.

To view the list of failed process instances, tap **Notifications** in the sidebar or the **Notifications** icon at the top of the Business Processes screen. Tap the process instance to view the process instance details.

The Notifications screen shows a maximum of ten instances per screen. To navigate through the list, tap **Show More** or **Show Previous**.

Unread process instances display with a red background. Mobile Monitor retains the list of failed notifications until you clear it. When you exit and relaunch the app, Mobile Monitor retrieves new notifications and adds them to the existing list of failed notifications. To clear the **Notification** counter, tap **Mark all as read** or the **X** icon at the top of the Notifications screen.

You can configure which information about a process instance displays in the Notifications screen, the notification polling interval, and whether notifications are enabled. For information about configuring notifications, see "Configuring Mobile Monitor " on page 132.

Configuring Mobile Monitor

Log on Mobile Monitor with the same user name and password that you use to access the Monitor user interface. Your user name must already be defined in Integration Server as a member of the ACL group MonitorAdministrators, MonitorUsers, or Administrators.

The first time you log on, Mobile Monitor authenticates you. You can save your password so that you do not need to retype it every time you open the application.

You must add a server in the app to retrieve business process and process instance statuses. Obtain the Integration Server IP address (or host name) and port information from your administrator.

You can configure notification options in Mobile Monitor, including the following:

- **Configure notification polling**. Configure how often the app polls for notifications or disable notification polling.
- **Configure notification options**. Select the information that displays about a process instance in the Notifications screen, including instance ID, custom ID, version, and date the model was last updated.
- **Configure process instance fields**. Select the information that displays about a process instance on the Process Instances screen. You can select up to three fields from the following: instance ID, custom ID, version, date the model was last updated, instance start time, instance iteration, and duration. The Process Instance Details screen displays all of this information about the process instance.
- **Configure logging**. Enable or disable logging for business processes. Use the on/off slider to set **Enable Logging**.

To configure Mobile Monitor

1. Enter your user name and password in Mobile Monitor.
2. Tap the **Servers** icon at the bottom of the Business Processes screen to configure the server connection.
   a. Tap the **Add Server** icon to add a server.

> **Note:**
> Mobile Monitor only shows processes and processes instances from the active server. To monitor processes on a different server, you can change which server is the active server.

b. Type the name of the server in the **Server Name** field.

c. In the **IS Host** and **IS Port** fields, type the IP address and port (provided by your administrator).

d. Complete the **Username** and **Password** fields.

e. In **Active**, set the slider to **On** to pull data from this server.

> **Tip:**
> Tap **Ping** to verify that the server details are correct and that the server is running.

f. Tap **Save** to save the server details.

3. Tap the **Gear** icon to configure global application settings.

a. Tap **Application Settings**.

b. In **Notification Polling**, set how often the app polls for notifications.

> **Note:**
> This setting only controls notifications. It does not control how often Mobile Monitor fetches business process statuses. To refresh the data, pull down on the Business Processes screen.

c. In **Notification Options**, select the information to display about a business process on the Notifications screen.

d. In **Process Instance Field Options**, select up to three fields to display about a process instance in the Process Instances screen.

e. In **Enable Logging**, set the field on or off to enable or disable logging.

f. Click **Save**.

A Archive Tables

Overview of Archive Tables

This section lists the database tables for which you have to set permissions before you run the data archive process. Make sure you have the permission to archive the tables specified for the Process Engine version you are using.

Process Archive Tables

Some Process Audit Archive tables were renamed in release 9.6. The following table compares the table names in release 9.5 SP1 with 9.6.

9.5 SP1 Table Name	9.6 Table Name	Name Changed?
WMCONTROL	PRA_PROCESS_ACTION	Yes
WMCUSTOMFIELDDEFINITION	WMCUSTOMFIELDDEFINITION	No
WMCUSTOMLOOPDATA	PRA_STEP_LOOP_LOGGED_FIELD	Yes
WMCUSTOMPROCESSDATA	PRA_STEP_LOGGED_FIELD	Yes
WMERROR	PRA_ERROR	Yes
WMPROCESS	PRA_PROCESS	Yes
WMPROCESSASSOC	PRA_PROCESS_CUSTOM	Yes
WMPROCESSATREST	PRA_PROCESS_AT_REST	Yes
WMPROCESSBLOCKAGE	Table is removed.	Yes
WMPROCESSDEFINITION	WMPROCESSDEFINITION	No
WMPROCESSIMAGE	WMPROCESSIMAGE	No
WMPROCESSRECENT	PRA_PROCESS_RECENT	Yes
WMPROCESSSTEP	PRA_PROCESS_STEP	Yes
WMPROCESSSTEPLOOP	PRA_PROCESS_STEP_LOOP	Yes
WMPROCESSTASK	WMPROCESSTASK	No
WMPROCESSTASKSTEP	WMPROCESSTASKSTEP	No
WMPROCESSTASKUSER	WMPROCESSTASKUSER	No
WMPROCESSTRANSITION	PRA_STEP_TRANSITION	Yes
WMSERVICEACTIVITYLOG	PRA_STEP_MESSAGE	Yes
WMSERVICE_MIN_MAX	WMSERVICE_MIN_MAX	No

9.5 SP1 Table Name	9.6 Table Name	Name Changed?
WMSTEPDEFINITION	WMSTEPDEFINITION	No
WMSTEPTRANSITION DEFINITION	WMSTEPTRANSITION DEFINITION	No

Note:
The following tables still exist as a part of the ISCoreAuditSchema: WMCONTROL, WMSERVICEACTIVITYLOG and WMSERVICE_MIN_MAX.

The following table lists the Process Audit tables for releases prior to 9.6.

Table Name	7.1.2	8.0 SP3, 8.2 SP1, 9.0 SP1	9.5 SP1
WMCONTROL	X	X	X
WMCUSTOMFIELDDEFINITION	X	X	X
WMCUSTOMPROCESSDATA	X	X	X
WMERROR	X	X	X
WMPROCESS	X	X	X
WMPROCESSASSOC	X	X	X
WMPROCESSATREST			X
WMPROCESSDEFINITION	X	X	X
WMPROCESSIMAGE	X	X	X
WMPROCESSRECENT	X	X	X
WMPROCESSSTEP	X	X	X
WMPROCESSTASK		X	X
WMPROCESSTASKSTEP		X	X
WMPROCESSTASKUSER		X	X
WMPROCESSTRANSITION	X	X	X
WMSERVICE	X	X	
WMSERVICEACTIVITYLOG	X	X	X
WMSERVICE_MIN_MAX	X	X	X
WMSTEPDEFINITION	X	X	X

Table Name	7.1.2	8.0 SP3, 8.2 SP1, 9.0 SP1	9.5 SP1
WMSTEPTRANSITIONDEFINITION	X	X	X

Server Archive Tables

The following table lists the Server Archive Tables for releases prior to 9.6.

Table Name	7.1.2	8.0 SP3, 8.2 SP1, 9.0 SP1	9.5 SP1
WMERROR	X	X	X
WMPROCESS	X	X	X
WMSECURITY	X	X	X
WMSERVICE	X	X	X
WMSERVICEACTIVITYLOG	X	X	X
WMSESSION	X	X	X
WMTXIN	X	X	X
WMTXOUT	X	X	X

Service Archive Tables

The following table lists the Service Archive Tables for releases prior to 9.6.

Table Name	7.1.28.0 SP3	8.2 SP1, 9.0 SP1	9.5 SP1
WMCONTROL	X	X	X
WMERROR	X	X	X
WMSERVICE	X	X	X
WMSERVICEACTIVITYLOG	X	X	X
WMSERVICEASSOC		X	X
WMSERVICE_MIN_MAX	X	X	X

Document Archive Tables

The following table lists the Document Archive Tables for releases prior to 9.6.

Table Name	7.1.2	8.0 SP3, 8.2 SP1, 9.0 SP1	9.5 SP1
WMCONTROL	X	X	X
WMDOCUMENT	X	X	X