software AG

Installing Software AG Products

Version 10.1

October 2017

Table of Contents

About this Guide

This guide provides pre-installation, installation, and uninstallation instructions for Software AG products, specifically Apama, Terracotta, and webMethods.

> **Note:** The information provided in this guide for Terracotta and Universal Messaging applies to when you are using those products with webMethods products. For instructions on installing Terracotta and Universal Messaging in other contexts, and for instructions on installing products not covered in this guide, see the documentation for those products.

This guide is intended for use with *Using the Software AG Installer*. That guide explains how to prepare your machine to use the Software AG Installer, and how to use the Software AG Installer and Software AG Uninstaller to install and uninstall your products.

> **Important:** If you want to upgrade products, you must use the instructions in *Upgrading Software AG Products*. The instructions in this installation guide do not cover upgrades, so following the instructions in this installation guide for upgrades would have unpredictable results.

Document Conventions

Convention	Description
Bold	Identifies elements on a screen.
Narrowfont	Identifies storage locations for services on webMethods Integration Server, using the convention *folder.subfolder:service* .
UPPERCASE	Identifies keyboard keys. Keys you must press simultaneously are joined with a plus sign (+).
Italic	Identifies variables for which you must supply values specific to your own situation or environment. Identifies new terms the first time they occur in the text.
`Monospace font`	Identifies text you must type or messages displayed by the system.

Convention	Description
{ }	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

Software AG Documentation Website

You can find documentation on the Software AG Documentation website at http://documentation.softwareag.com. The site requires Empower credentials. If you do not have Empower credentials, you must use the TECHcommunity website.

Software AG Empower Product Support Website

You can find product information on the Software AG Empower Product Support website at https://empower.softwareag.com.

To submit feature/enhancement requests, get information about product availability, and download products, go to Products.

To get information about fixes and to read early warnings, technical papers, and knowledge base articles, go to the Knowledge Center.

Software AG TECHcommunity

You can find documentation and other technical information on the Software AG TECHcommunity website at http://techcommunity.softwareag.com. You can:

- Access product documentation, if you have TECHcommunity credentials. If you do not, you will need to register and specify "Documentation" as an area of interest.

- Access articles, code samples, demos, and tutorials.

- Use the online discussion forums, moderated by Software AG professionals, to ask questions, discuss best practices, and learn how other customers are using Software AG technology.

- Link to external websites that discuss open standards and web technology.

1   Software AG Installer, Recommendations, and Product Licenses

Software AG Installer

This guide provides product-specific preparation, installation, and uninstallation instructions. These instructions are intended for use with *Using the Software AG Installer*. That guide explains how to prepare your machine to use the Software AG Installer, and how to use the Software AG Installer to install and uninstall your products. It describes the various methods you can use to install and uninstall, such as using the installer wizard, the installer console mode, an installation script, or an installation image. It describes JDK support for the installer, the uninstaller, and the products, and explains installation and uninstallation logging

Recommendations

- Software AG strongly recommends that you create installation images of your products and store them on your internal network. Create an image for each operating system on which you plan to install that is covered by your license agreement. Storing images enables you to ensure consistency across installations over time; the installer provides only the latest product releases, while installation images contain the product releases you download into them. In addition, installing from images helps you reduce WAN traffic and improve installation speeds. For instructions on creating an installation image, see *Using the Software AG Installer*.

- When you need multiples of the same release of a product on one machine, Software AG recommends that you create multiple instances of the product in the same installation rather than creating multiple installations of the same product on the machine. Most products now offer this functionality. The valid use cases for creating multiple installations on the same machine are when those installations are from different releases (for example, side-by-side upgrade), different environments (for example, production environment on fix level 1 and staging environment on fix level 2), or different projects (for example, installations used by different departments within an organization).

  > **Note:** If you are using a virtualization layer such as VMWare, the above use cases are typically handled using different VM instances on a single powerful VM server, with a single installation on each VM instance.

Product Licenses

You must provide license files to install and, later, to start some products. Software AG sends these license files with your installation message; save the license files in a directory whose name and path do not contain any spaces.

When you provide the license files during installation, the installer validates that:

■   The license file is the correct one for the product.

■   The license file has not been changed.

■   The license file has not expired.

■   The product is licensed on the operating system on which you are installing.

You will only be able to install a product if its license file passes these checks. The products will repeat these checks at startup.

2   Installing Products Using the Software AG Installer or Command Central

Installing Products Using the Software AG Installer or Command Central

You can install products using different methods. The methods can involve Software AG Installer, Command Central, and Platform Manager. Platform Manager is the agent that enables Command Central to install, administer, and manage your products.

- You can use the Software AG Installer to install all products except Command Central on the local machine. The installer automatically installs Platform Manager with the products. For instructions, see *Using the Software AG Installer* and this guide.

- You must use the Command Central bootstrapper to install Command Central, and then use the bootstrapper or Command Central to install Platform Manager in the directories into which you want to install your products. After Platform Manager is installed, you can use Command Central to install your products on remote machines.

  For instructions on installing Command Central and Platform Manager and then using Command Central to install products, see the *Software AG Command Central Help*. For a list of products that Command Central can install, see the *Software AG Command Central Feature Support Matrix*.

If you uninstall products using the Software AG Installer or Command Central, regardless of which tool you used to install them.

3 Hardware Requirements

Minimum and Recommended Hardware Requirements

The table below lists the minimum and recommended hardware requirements for your products. Recommended amounts are shown in parentheses. Products not listed have minimal or no hardware requirements (for example, products that are hosted on Integration Server).

> **Important:** The hardware requirements below are for the listed products only. They do not include 2GB of hard drive space for Software AG infrastructure shared libraries that are required by most products and that are installed once per installation directory. They do not include requirements for supporting software such as RDBMSs. They do not include the additional requirements listed in "Additional Hardware Requirements" on page 24. You must determine the overall hardware requirements for each of your machines based on the overall requirements of your operating system and software.

Product	Hard Drive Space	RAM	CPUs
ActiveTransfer Server	100MB (200MB)	512MB	1
AgileApps Cloud	20GB	8GB	1 (4)
Apama			
Server	750MB	4GB	2
Capital Markets Adapters (each)	5MB	4GB	2
Capital Markets Foundation	55MB	4GB	2
Dashboard Viewer	400MB	4GB	2
Platform Manager Plug-in	5MB	4GB	2
Predictive Analytics Plug-in	25MB	4GB	2
API Gateway*	1.5GB	1.75GB	1 (2)

Product	Hard Drive Space	RAM	CPUs
API Portal	20GB	8GB	4
ApplinX			
Server	100MB	80MB +300KB per user	2 (4)
Administration	75MB	40MB	
JSP Framework	50MB	2GB	
C#.NET Framework	55MB	2GB	
VB.NET Framework	30MB	2GB	
Visual Studio Add-In	1MB	2GB	
webMethods Broker	2GB (10GB)	512MB (2GB)	1
CentraSite			
Registry Repository	3GB	4GB (8GB)	1 (2)
Application Server Tier	3GB	4GB (8GB)	1 (2)
Content Service Platform Server	1GB	2GB	1
Software AG Designer Eclipse platform	500MB (1.5GB)	1.5GB (2GB)	1 (2)
EntireX			
Core Files	100MB	2GB	1
Administration and Monitoring	10MB	1GB	1

Product	Hard Drive Space	RAM	CPUs
Broker	20MB Windows, 30MB UNIX		
Development Tools	5MB Windows, 10MB UNIX	1GB	1
Mini Runtime	10MB	1GB	1
Web Services deployed in Software AG Runtime**	140MB Windows, 150MB UNIX	1GB	1
EntireX plug-in to Software AG Designer	50MB	50MB	
Integration Agent	275MB (500MB)	512MB (1GB)	1
Integration Server	300MB (500MB)	1GB (2GB)	1 (2)
Locate	20GB (40GB)	512MB (4GB)	1 (3)
MashZone NextGen	2000MB	2GB	2
Microservices Container	250MB (400MB)	1 GB	1
Mobile			
Mobile Administrator	300MB	4GB	2
Mobile Designer	1GB	1GB	1
My webMethods Server	300MB (500MB)	1GB (2GB)	1

Product	Hard Drive Space	RAM	CPUs
OneData			
Server	512MB	2GB (4GB)	1 (2)
MDR Modules	50MB		
Customer MDM Template	20MB		
Optimize			
Analytic Engine	1GB	4GB (8GB)	2 (4)
Infrastructure Data Collector	300MB	2GB*	1
Web Service Data Collector	100MB	256MB	1
Platform Manager	500MB	128MB	1
ARIS PPM	10GB	7GB (16GB)	2 (4)
ARIS PPM Analysis GUI	5GB	4GB	1
TerracottaBigMemory Max when used with webMethods products	200MB (500MB)	3GB	1
Trading Networks Server	50MB		1
Universal Messaging realm server when used with webMethods products	500MB	1GB	1
Zementis Predictive Analytics	5GB	4GB	1

*Includes 200MB hard drive space and 256MB RAM for Kibana and 300MB hard drive space and 512MB RAM for Event Data Store.

Product	Hard Drive Space	RAM	CPUs

**Includes 140MB hard drive space, 1GB RAM, and 1 CPU for the required and automatically installed Web Services Stack.

Additional Hardware Requirements

Product	Additional Hardware Requirements
ActiveTransfer Server	ActiveTransfer Server needs additional hard drive space if you use the hard drive as your virtual file system. The space required depends on the total size of documents ActiveTransfer Server needs to store, upload, and download. ActiveTransfer Server might need additional hard drive space for log files.
AgileApps Cloud	See the AgileApps wiki.
Apama	Hard drive space, RAM, and CPU requirements for Apama vary significantly depending on the Apama application you run. You will need to determine processing and memory requirements by measuring the application, since these requirements vary. Contact Software AG Global Support for help.
webMethods Broker	webMethods Broker might need additional hard drive space if your documents are large, or if your clients use many guaranteed documents.
Software AG Designer	*Using the Software AG Installer* lists free space the Software AG Installer requires in its system temp (Windows) or temporary (UNIX) directory. For Software AG Designer Eclipse platform, the installer needs 700MB additional free hard drive space in that directory. Each product plug-ins you install in Software AG Designer can require from 50K to 150MB of additional free hard drive space.
Deployer	Hard drive space for the deployment projects you create. When allocating space, allow for the number of projects to increase over time.
Locate	The hard disk space for Locate must be in a RAID configuration. The range of hardware requirements in the table is intended to

Product	Additional Hardware Requirements
	accommodate your needs up to the largest dataset you could install.
MashZone NextGen	The RAM requirement correlates with the amount of data processed and might need to be increased based on use case. You might need to increase CPU cores based on the number of simultaneous users.
Mobile Administrator	You will need additional hard drive space depending on the number of applications you plan to add, the size of their source code, builds, screen shots, and logs.
Mobile Designer	You will need additional hard drive space to accommodate the SDKs for the mobile devices for which you want to develop applications.
OneData	OneData needs additional hard drive space for temp, import, and export files in the *Software AG_directory*\OneData\config directory.
Optimize	
Analytic Engine	The Analytic Engine needs 1GB virtual swap space and a disk subsystem with 10GB available free space. In a production environment, the disk subsystem must have a redundant array of independent disks (RAID 5), an UltraWide SCSI, and 10K to 15K RPM drives.
	Also in a production environment, the Analytic Engine needs additional hard drive space for log files. The recommended amount is 100MB; the engine needs 50MB of space for each log file.
Web Service Data Collector	Each Web Service Data Collector needs 128MB of virtual swap space. In a production environment, each Web Service Data Collector needs additional hard drive space for log files. The recommended amount is 75MB; each Web Service Data Collector needs 5MB of hard drive space for each log file. More hard drive space might be needed if you use debug level or higher logging.
ARIS PPM	You might need additional RAM and hard drive space, depending on the number of process instances you import and store in ARIS PPM. The more process instances, the more RAM and hard drive space you will need.

Product	Additional Hardware Requirements
Trading Networks	To use large document handling, Trading Networks Server needs additional hard drive space on which to temporarily save documents (instead of storing them in memory). For detailed information, see *webMethods Trading Networks User's Guide*.
Universal Messaging	If you use persistent topics or queues, or persistent messages, Universal Messaging needs additional hard drive space to persist the published data. The amount of space required would be loosely based on this equation: messages per second x message size x message time to live. If you store messages in memory only, Universal Messaging needs additional heap within the JVM to hold references to these messages. The heap size required is based on the same equation.
	Universal Messaging realm servers support high-performance spin locks. If you enable spin locking for a realm server, and then add the realm server to a cluster, the realm server needs two additional CPUs to handle high-performance cluster event processing.
	If you enable a Universal Messaging realm server to use shared memory, the realm server needs two additional CPUs for each client that connects to it with shared memory.
Zementis Predictive Analytics	Additional storage may be required for working and temporary space. The amount of the total required disk space depends on the size of files (models, data, etc.) expected to be uploaded to and processed by Zementis Predictive Analytics.

4 Operating System Support, Browser Support, and RDBMS Support and Connection Information

Operating System Support

For information on operating systemsthat are supported by your products, see *System Requirements for Software AG Products*.

For Apama, also see *Apama Supported Platforms*.

If the vendor for your operating system recommends installing the latest service packs, patch-bundles, and fixes, Software AG encourages you to follow that recommendation. If you experience problems with installed service packs, patch-bundles, or fixes, report those problems directly to the vendor.

Browser Support

For information on browsers that are supported by your products, see *System Requirements for Software AG Products*.

RDBMS Support and Connection Information

For information on RDBMSs that are supported by your products, see *System Requirements for Software AG Products*. For information on supported database drivers, see "Database Components" on page 94.

Some products require you to supply database connection information during installation. Sample URL formats for supported database drivers are shown in the product panels. Keep in the mind the following:

- Most products use the DataDirect Connect JDBC 5.1 driver. For information about options supported by this driver, see DataDirect Connect documentation, available on the Software AG Documentation website.

- If you are using the installer GUI mode, for ease of use, the database connection values you enter on one panel are reflected on the next as you go forward through the installer the first time.

- Use the DataDirect Connect connection option MaxPooledStatements=35 on all database URLs except those for Trading Networks. This connection option improves performance by caching prepared statements. (Trading Networks caches its prepared statements using its own pooling mechanism).

- If the database user and password do not yet exist, specify the database user and password you will create after installation. You can use one database user for multiple database components, or you can use a different database user for each database component.

- For DB2 database URLs:

■ You must specify schema names using all uppercase letters. The tool you use
   to create database components later in this guide creates schema names in all
   uppercase letters, and DB2 is case sensitive for schema names.

■ if your product will connect to a schema other than the default schema for
   the specified database user, you must specify these connection options in the
   database URL, where `AlternateID` is the name of the default schema that is
   used to qualify unqualified database objects in dynamically prepared SQL
   statements:

```
;AlternateId=schema;"InitializationString=(SET CURRENT
PATH=current_path,schema)";MaxPooledStatements=35
```

5 Shut Down Products

Shut down all non-Software AG applications that are running on the machine on which you are going to install. If these applications are not shut down, product files might become corrupted.

If you are going to install products into an existing Software AG product installation directory (that contains products from the same release), shut down running products in that directory so the installer can update key files that are locked by the operating system. For instructions, see the product documentation.

6   Installing Deployer and Asset Build Environment

Deployer and Asset Build Environment Installation

You can deploy assets from one environment to another (for example, from an integration environment to a staging environment or from a staging environment to a production environment) using Deployer and Asset Build Environment. In runtime-based deployment, you use Deployer to deploy assets from source servers to target servers. In repository-based deployment, you use Asset Build Environment to build assets from source servers or a VCS to a file-based repository, and then use Deployer to deploy the assets from the repository to target servers. Software AG recommends using repository-based deployment.

If no firewall separates your integration, staging, and production environments, install one Deployer to handle all environments. If a firewall separates your integration and staging environments from your production environments, install one Deployer to handle your integration and staging environments and a separate Deployer to handle your production environments. The Deployer host machine must have access to all source servers from which you want to deploy, and all target servers to which you want to deploy. Restrict access to the Deployer host machine to only those people who need to use Deployer, such as release managers, infrastructure engineers, system administrators, and operators.

Install Asset Build Environment on integration machines on which assets are created so you can verify successful builds before checking code into your VCS. For centralized builds, install Asset Build Environment on a separate machine, known as a continuous integration server, where you will use it to build assets into components that can be deployed. Restrict access to the Asset Build Environment host machine to only those people who need to use Asset Build Environment, such as release managers, infrastructure engineers, system administrators, and operators.

Deployer Host Server

Deployer runs on Integration Server or Microservices Container. For information on software requirements, machine preparation, installation, and configuration for these products, see the relevant sections in "Installing and Uninstalling Application Integration, B2B Integration, Dynamic Apps, or Microservices Platform" on page 41. When the installer offers a choice of database for Integration Server or Microservices Container, accept the default choice of **Embedded database**.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Register Daemons and Set File Permissions

On a UNIX system, if you created or installed from a script, or if you were asked whether to use sudo and you chose not to, the installer was not able to register daemons and set file permissions. Perform these tasks now by running the script *Software AG_directory* /bin/ afterInstallAsRoot.sh as a user who is in the sudoers configuration.

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

> **Important:** If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Uninstall Deployer Host Server (Integration Server or Microservices Container)

See this section in "Installing and Uninstalling Application Integration, B2B Integration, Dynamic Apps, or Microservices Platform" on page 41.

7 Installing and Uninstalling Apama

Typical Development Installation

The Software AG Installer offers typical development installations of products that enable you to perform a certain task. When you select a typical development installation, the installer automatically selects all products that make up that installation. The typical development installation for this product is called Apama.

> **Note:** The Apama typical development installation does not include Apama Platform Manager Plug-in. If you want to manage your Apama installation from Command Central, manually select the plug-in when selecting the Apama components to install.

To create production environments, work with your administrators, Software AG Global Consulting Services, and best practices documentation.

Software Requirements and Considerations

- To build samples and develop Apama applications using Apama APIs, you might need to install additional compilers. For details, see the Apama documentation.

- To build samples and develop applications on supported Linux platforms using the C/C++ API, install the GCC-C++ package and all its dependencies. These are typically provided on the installation media as part of your operating system distribution.

- Apama does not support Security Enhanced Linux (SELinux). If you are going to install on a Linux system, turn off this option.

- If you are going to use the compiled Apama runtime, install the binutils package. For information about the compiled runtime, see the Apama documentation about starting the event correlator.

- The file system to which the correlator writes its log files and persistence database can significantly affect the performance of latency-critical applications. File systems such as ext3, ext4, NFS, ZFS and GFS2 can have poor worst-case latency, especially when the system is loaded. For latency-critical applications, Software AG recommends using XFS where possible.

Prepare Your Machine

The Apama Dashboard Data Server port must be accessible to the Apama Dashboard Viewer. If you are going to install on a Windows system, and the firewall is enabled, unblock network access for this port. The default value for the port is 3278. For security reasons, you should never change firewall settings such that this port is exposed to untrusted clients.

If you are going to install on a Linux system:

■ Apama does not support Security Enhanced Linux (SELinux). If you are going to install on a Linux system, turn off this option.

■ Linux systems are often run with their core file size limit set to zero. If problems occur, no core file can be written, which makes it difficult to determine the cause of the problem. Set the core file size to unlimited in the login configuration for each Apama user.

Install Products

Start the Installer, Provide General Information, and Choose the Products to Install

Read the requirements in *Using the Software AG Installer* (for example, the requirement to create a user account that has the proper privileges for Windows and that is non-root for UNIX). Follow the instructions in that guide to start the installer and provide general information such as proxy server, release to install, installation directory, and how to use the product selection tree.

> **Important:** Unless otherwise stated for a specific product, do not install products from this release into a Software AG directory that contains products from any other release. Unless otherwise stated for a specific product, you cannot mix products from different releases in the same installation directory; if you do so, you will experience problems or be unable to access functionality.

On the product selection tree, select the Apama components to install.

Supply Product Configuration Information

Apama

If you are installing Apama Server, the installer asks for a license file. If you provide one, the installer will copy it into the Apama work directory. If you do not provide one, the Apama correlator will run with reduced capabilities. See the Apama documentation for information on running without a license.

If you are installing the Apama Server or Dashboard Viewer, the installer asks for the path to the work directory in which Apama will store its user-created content, such as log files. Specify an empty directory or a directory that does not yet exist. If you specify a directory that does not yet exist, the installer will create it and populate it with standard files and directories. You can include non-English characters in the work directory location, but they must be within the set of characters from the operating system's default character encoding.

Predictive Analytics Engine

If you are installing the Apama Predictive Analytics Plug-in, the installer asks for the Predictive Analytics Engine licence file.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Specify the Predictive Analytics Engine License File

If you installed the Predictive Analytics Engine but did not specify the license file during installation, store the license file in the directory you specified as the Apama work directory. The Predictive Analytics Engine cannot start without the license file in that location.

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

> **Important:** If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Uninstall Products

Follow the instructions in *Using the Software AG Installer*, with the additional guidelines below.

If you are going to uninstall Software AG Designer Eclipse, and you want to review or preserve Eclipse installation information, do so before starting the uninstaller, because the uninstaller will delete the entire Software AG Designer installation directory. User-created data that is stored in other directories, such as your workspace, will remain untouched.

8 Installing and Uninstalling Application Integration, B2B Integration, Dynamic Apps, or Microservices Platform

Typical Development Installations

The Software AG Installer offers typical development installations of products that enable you to perform a certain task. When you select a typical development installation, the installer automatically selects all products you have licensed that make up that installation. You can select more products or deselect products if desired. For detailed information on how the products in each typical development installation work together, see *Understanding Software AG Products*.

For some typical development installations, you must create certain *database components*. A database component is a grouping of database objects that is used by one or more products. The table below lists these database components. For complete information about each database component, see "Creating and Dropping Database Components" on page 93.

To create production environments, work with your administrators, Software AG Global Consulting Services, and best practices documentation.

Typical Installation	Database Components
ActiveTransfer	ActiveTransfer, all for Integration Server, all for My webMethods Server
Application Integration Development	All for Integration Server, all for My webMethods Server, all for Trading Networks
Application Platform Development	All for Integration Server, all for My webMethods Server
CloudStreams Service Development	CloudStreamsEvents, all for Integration Server
Composite Application Development	All for My webMethods Server
Dynamic Apps Platform	All for Integration Server, BPM, My webMethods Server, and Optimize; optionally, Staging and Reporting

Software Requirements and Considerations

Product	Software Requirements and Considerations
AgileApps Cloud	See the AgileApps wiki.
Integration Server and Microservices Container	On Mac OS X systems, Integration Servers cannot connect to webMethods Brokers using SSL.

Prepare Your Machine

Prepare to Install Integration Agent, Integration Server, or Microservices Container on a Mac OS X or Other UNIX System

The ability of Integration Agent, Integration Server, and Microservices Container to handle traffic is constrained by the number of file descriptors available to the product's process. On most systems, 64 file descriptors are available to each process by default. If you are going to install the product on a Mac OS X or other UNIX system, Software AG recommends that you ask your system administrator to increase the number of file descriptors available to the product process to at least 1024.

> **Important:** You might have to increase this number depending on the number of files the product needs to have open at one time. It is dangerous to set the rlim_fd_max value higher than 1024 because of limitations with the select function, so if the product requires more file descriptors, ask the system administrator to set the setrlimit value directly.

Prepare to Install My webMethods Server

If you are going to install the My webMethods Server on a UNIX system, make sure you have set sufficient user limits for the shell you use to start the installation and the product daemons. For example, if you have multiple language packs installed for My webMethods Server, or if there are multiple products on the machine, the daemons for My webMethods Server will run out of memory if they are started from the shell with a low user limit for data. If your system policy allows it, Software AG recommends setting the value for coredump, data, file, memory, and threads to at least 32768, and the value of nofiles to 8192. For more information about setting and displaying the ulimits, read the man page or ask your system administrator.

Install Products

Start the Installer, Provide General Information, and Choose the Products to Install

As you go through the Software AG Installer, you will be asked for various types of information, such as product license files and database connection parameters. Gather this information before starting the installer.

Read the requirements in *Using the Software AG Installer* (for example, the requirement to create a user account that has the proper privileges for Windows and that is non-root for UNIX). Follow the instructions in that guide to start the installer and provide general information such as proxy server, release to install, installation directory, and how to use the product selection tree.

> **Note:** When you install on a UNIX system, the installer might ask whether to use sudo and related scripts. If you want to have the installer run the script, the user under which you run the installer must be in the sudoers configuration. If you are creating or installing from an installation script, the option to use sudo is not available because the installer does not store the sudo password in scripts, for security reasons. Instead you will perform certain tasks after installation as described later in this chapter.

> **Important:** Unless otherwise stated for a specific product, do not install products from this release into a Software AG directory that contains products from any other release. Unless otherwise stated for a specific product, you cannot mix products from different releases in the same installation directory; if you do so, you will experience problems or be unable to access functionality.

On the product selection tree, select the products to install.

- If you select packages for Integration Server or Microservices Container, the installer will copy the packages into a package repository in the product installation. The installer will also create a server instance and will ask whether to install all packages in the package repository on the server instance. On the language pack selection tree, the installer will copy all language packs you select into the package repository and will always install them on the server instance. After installation, you can create additional server instances and install packages and language packs from the package repository on those server instances and the server instance created during installation. For instructions, see the *webMethods Integration Server Administrator's Guide*.

- The licenses for CloudStreams or Enterprise Gateway allow only partial functionality of their host server (that is, Integration Server or Microservices Container). Install each product on its own host server. Do not install other products on that host server because those other products might not work properly.

■   If you install My webMethods Server, and later install Business Console, Task Engine, or My webMethods user interfaces in a separate run of the installer, those components are installed on all My webMethods Server instances in the target installation directory.

■   If you install **Designer > Composite Applications**, also select **Infrastructure > Shared Platform > Platform**.

After the product selection tree, the installer displays the language pack selection tree. For information on language packs, see the international operating environments appendix in this guide. The installer then displays panels (GUI mode) or prompts (console mode) that ask for product configuration information. Most are self-explanatory, so this section shows only the panels that require explanation. The information and fields on the prompts is identical to the information and fields on the panels.

> **Important:** Make sure all ports you specify on panels or in response to prompts are not already being used by other products on your machine. The installer cannot detect whether ports are in use when products are shut down, and the shutting down of products is a requirement for running the installer. See the list of default ports used by Software AG products.

Supply Product Configuration Information

Most product panels or prompts are self-explanatory. This section provides information for those products that require additional explanation.

Integration Server or Microservices Container

If you are installing CloudStreams or Enterprise Gateway, specify the license for that product, not a license for Integration Server.

If you are installing Integration Server or Microservices Container with the Application Platform package and Software AG Designer with the Application Platform plug-ins in the same installation directory for development purposes, do not register for automatic startup. The Application Platform plug-ins invoke scripts that start and stop the server from Software AG Designer.

To determine where to store data written by the server, see " Integration Server or Microservices Container Data Storage" on page 95.

If you want to use the MySQL Community Edition database, select **Embedded**. You will configure the MySQL database connection after installation using instructions in the product documentation.

My webMethods Server

To determine where to store data written by My webMethods Server, see " My webMethods Server Data Storage" on page 98.

If you want to use the MySQL Community Edition database, select **Embedded**. You will configure the MySQL database connection after installation using instructions in the product documentation.

Trading Networks

If you want to use the MySQL Community Edition database, accept the default (Oracle) You will configure the MySQL database connection after installation using instructions in the product documentation.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Register Daemons and Set File Permissions

On a UNIX system, if you created or installed from a script, or if you were asked whether to use sudo and you chose not to, the installer was not able to register daemons and set file permissions. Perform these tasks now by running the script *Software AG_directory*/bin/ afterInstallAsRoot.sh as a user who is in the sudoers configuration.

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

> **Important:** If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Complete the Content Service Platform Installation

1. Create the database user and password you want to use to create the Content Service Platform database tables. The database user must have privileges to create tables and indices. For Oracle, grant the database user CONNECT and RESOURCE privileges only.

2. Set up the Content Service Platform Server. Open a command window, go to the *Software AG_directory*/CSP/server directory, and run this command:

   ```
   cspconfigurator.{bat|sh} -Dsetup.env.db.user=database_user
   ```

```
-Dsetup.env.db.password=password
```

> **Note:** Do not run this command more than once.

3. If you installed the Content Service Platform Windows Client, go to the *Software AG_directory*\CSP\winclient directory and locate the Software AG webMethods CSP Client Setup 9.0.exe file. If the file is on a Windows system, double-click the .exe file. If the file is on a non-Windows system, copy the file to a Windows system and then double-click the .exe. The client installer wizard opens; follow the instructions to install the client.

> **Note:** The .exe file name might contain additional numbers.

4. On a Windows system, you were given the choice to install Content Service Platform as an application or a service. If you installed Content Service Platform as an application, the Windows Start menu includes entries for starting and stopping Content Service Platform. If you installed Content Service Platform as a service, you must open a command window, go to the *Software AG_directory*/CSP/server directory and run the command `csp install` to complete the installation of the service and create Windows Start menu entries.

Uninstall Products

Follow the instructions in *Using the Software AG Installer*, with the additional guidelines below.

If you are going to uninstall Software AG Designer Eclipse, and you want to review or preserve Eclipse installation information, do so before starting the uninstaller, because the uninstaller will delete the entire Software AG Designer installation directory. User-created data that is stored in other directories, such as your workspace, will remain untouched.

If you are going to uninstall Integration Server or Microservices Container, retract all assets you have published to CentraSite. For instructions, see the *webMethods BPM and CAF CentraSite Metadata Help*.

If you are going to uninstall from a Mac OS or other UNIX system, and you used sudo to register daemons and set file permissions during or after installation, you must also use sudo to unregister daemons and remove file permissions, as follows:

- If you are not going to use an uninstallation script to uninstall, you can choose the uninstaller option to perform this task, or you can perform this task before starting the uninstaller by executing the script *Software AG_directory*/bin/beforeUninstallAsRoot.sh as a user who is in the sudoers configuration.

- If you are going to use an uninstallation script, the uninstaller cannot execute the script because it does not store the sudo password, for security reasons. You must execute the script yourself before starting the uninstaller.

9 Installing and Uninstalling Analytics & Decisions

Overview

The information provided in this chapter for Terracotta and Universal Messaging applies to when you are using those products with webMethods products. For instructions on installing those products in other contexts, see the product documentation.

Typical Installation

The Software AG Installer offers typical development installations of products that enable you to perform a certain task. When you select a typical development installation, the installer automatically selects all products you have licensed that make up that installation. You can select more products or deselect products if desired. The typical installation for these products is called Analytics & Decisions.

For this typical installation, you must create certain *database components*. A database component is a grouping of database objects that is used by one or more products. Depending on the products you have licensed, you must create all the database components for Integration Server, My webMethods Server, and Optimize. For complete information about each database component, see "Creating and Dropping Database Components" on page 93.

To create production environments, work with your administrators, Software AG Global Consulting Services, and best practices documentation.

Software Requirements and Considerations

Product	Software Requirements and Considerations
ARIS Process Performance Manager	You cannot install ARIS Process Performance Manager, also known as ARIS PPM, on integrated network drives or on substituted drives.
Universal Messaging	■ If you anticipate large-scale numbers of client connections or throughput, Software AG recommends using a 64-bit JVM for Universal Messaging realm servers, to enable larger heap sizes. ■ On HP-UX systems, shared memory drivers are currently not supported due to an implementation problem with the HP JVM.

Prepare Your Machine

Prepare to Install MashZone NextGen on a Linux System

- Check the setting for shared memory (kernel parameter shmmax) by executing the command `sysctl -a | fgrep kernel.shmmax`. If the value is less than 629145600, log on as root user and increase the value by executing `sysctl -w kernel.shmmax=629145600` or `echo "kernel.shmmax=629145600" >> /etc/sysctl.conf`, then activate the new value by executing `sysctl -p`.

- Check the settings for the system-wide maximum number of file descriptors (kernel parameter fs.file-max) by executing the command `sysctl -a | fgrep fs.file-max`. If the value is less than 200000, log on as the root user and increase the value by executing `sysctl -w fs.file-max=200000` or `echo "fs.file-max=200000" >> /etc/sysctl.conf`, then activate the new value by executing `sysctl -p`.

- Check the user, group, and process settings for the maximum number of open file descriptors by executing the command `ulimit -Hn and ulimit -Sn`, where -Hn is the hard limit and -Sn is the soft limit. If the value is less than 200000, log on as a non-root user and increase the value by executing `ulimit -n 200000`. To permanently save this setting for the user, execute:

```
echo "<user name> soft nofile 200000" >> /etc/security/limits.conf
echo "<user name> hard nofile 200000" >> /etc/security/limits.conf
```

- Enter the host name of the machine on which you are installing products in the DNS of the network or in the file /etc/hosts.

Prepare to Install Optimize

If you are going to install the Optimize Analytic Engine on a UNIX system, make sure you have set sufficient user limits for the shell you use to start the installation and the product daemons. For example, the daemons for the Optimize infrastructure will run out of memory if they are started from the shell with a low user limit for data. If your system policy allows it, Software AG recommends setting the value for coredump, data, file, memory, and threads to at least 32768, and the value of nofiles to 8192. For more information about setting and displaying the ulimits, read the man page or ask your system administrator.

Prepare to Install ARIS PPM

Prepare to Install ARIS PPM on a Windows System

Windows systems offer only 5000 ports for TCP/RMI connections. This number might not be sufficient for error-free communication if the system is too busy (for example, distributed systems, use of Performance Dashboard). To change the parameter, add

an entry like the one below to the Windows registry. This example entry increases the number of available ports to 8192.

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters,
"MaxUserPort"=dword:00002000
```

Prepare to Install ARIS PPM on a Linux System

- Check the setting for shared memory (kernel parameter shmmax) by executing the command `sysctl -a | fgrep kernel.shmmax`. If the value is less than 629145600, log on as the root user and increase the value by executing `sysctl -w kernel.shmmax=629145600` or `echo "kernel.shmmax=629145600" >> /etc/sysctl.conf`, then activate the new value by executing `sysctl -p`.

- Check the settings for the system-wide maximum number of file descriptors (kernel parameter fs.file-max) by executing the command `sysctl -a | fgrep fs.file-max`. If the value is less than 200000, log on as the root user and increase the value by executing `sysctl -w fs.file-max=200000` or `echo "fs.file-max=200000" >> /etc/sysctl.conf`, then activate the new value by executing `sysctl -p`.

- Check the user, group, and process settings for the maximum number of open file descriptors by executing the command `ulimit -Hn` and `ulimit -Sn`, where -Hn is the hard limit and -Sn is the soft limit. If the value is less than 200000, log on as a non-root user and increase the value by executing `ulimit -n 200000`. To permanently save this setting for the user, execute:

  ```
  echo "<user name> soft nofile 200000" >> /etc/security/limits.conf
  echo "<user name> hard nofile 200000" >> /etc/security/limits.conf
  ```

- Enter the host name of the machine on which you are installing products in the DNS of the network or in the file /etc/hosts.

Prepare to Install Zementis Predictive Analytics

Set an environment variable with the name ADAPA_HOME to the absolute path to the directory where Zementis Predictive Analytics will store all uploaded artifacts, or update the Zementis Predictive Analytics custom_wrapper.conf file in the sag-installation-folder/profiles/CTP/configuration directory by adding set.ADAPA_HOME=%OSGI_CONFIGURATION_AREA%/adapa-folder. These artifacts include PMML files, custom resources, and the product license key file (zementis.license).

Install Products

Start the Installer, Provide General Information, and Choose the Products to Install

As you go through the Software AG Installer, you will be asked for various types of information, such as product license files and database connection parameters. Gather this information before starting the installer.

Read the requirements in *Using the Software AG Installer* (for example, the requirement to create a user account that has the proper privileges for Windows and that is non-root for UNIX). Follow the instructions in that guide to start the installer and provide general information such as proxy server, release to install, installation directory, and how to use the product selection tree.

> **Note:** When you install on a UNIX system, the installer might ask whether to use sudo and related scripts. If you want to have the installer run the script, the user under which you run the installer must be in the sudoers configuration. If you are creating or installing from an installation script, the option to use sudo is not available because the installer does not store the sudo password in scripts, for security reasons. Instead you will perform certain tasks after installation as described later in this chapter.

> **Important:** Unless otherwise stated for a specific product, do not install products from this release into a Software AG directory that contains products from any other release. Unless otherwise stated for a specific product, you cannot mix products from different releases in the same installation directory; if you do so, you will experience problems or be unable to access functionality.

If you are installing MashZone NextGen or ARIS PPM on a Windows system, you must install using the Windows Administrator user account. To do so, instead of double-clicking the installer .exe file to start the installer, right-click the .exe file and click **Run as administrator**. Other user accounts do not work, even if they were assigned administrator privileges.

If you are installing on a Windows system and you want to install CentraSite Application Server Tier and Zementis Predictive Analytics Integrated Server in the same installation directory, you must install the two products in two separate runs of the installer.

On the product selection tree, choose the products to install. If you are installing Optimize, you also need a package called Optimize Support that provides data about business processes and KPIs to Optimize. You can install this package on Integration Server or Microservices Container. If you install the package on Microservices Container, data is provided for unorchestrated business processes only. Data is not provided for orchestrated business processes or events.

After the product selection tree, the installer displays the language pack selection tree. For information on language packs, see the international operating environments appendix in this guide. The installer then displays panels (GUI mode) or prompts (console mode) that ask for product configuration information. Most are self-explanatory, so this section shows only the panels that require explanation. The information and fields on the prompts is identical to the information and fields on the panels.

> **Important:** Make sure all ports you specify on panels or in response to prompts are not already being used by other products on your machine. The installer cannot detect whether ports are in use when products are shut down, and

> the shutting down of products is a requirement for running the installer. See the list of default ports used by Software AG products.

Supply Product Configuration Information

Most product panels or prompts are self-explanatory. This section provides information for those products that require additional explanation.

Optimize

If you have multiple network interfaces, identify the one to use; the default is the default network interface for local machine. Do not use localhost.

ARIS PPM

Select a system configuration for ARIS PPM that can handle your expected usage needs. In each use case below, the values are valid only for the specified number of ARIS PPM clients. The database is not included in the estimate.

System Configuration	ARIS PPM Clients	Users	System Hardware
Small	1	Up to 5	20 GB free hard drive space, 4 GB RAM, 2 CPUs
Medium	1	Up to 10	Depends on ARIS PPM customization
Large	Several	Up to 100	Depends on ARIS PPM customization

Universal Messaging

If you choose to install a Universal Messaging realm server (alone or with Enterprise Manager or Template Applications), provide the information below.

Field	Entry
License file	The contents of this field varies as explained below.
	■ If you have the standard Universal Messaging license, specify the full path to that license file.

Field	Entry
	■ If you do not have a Universal Messaging license file, leave the field blank. You can use Universal Messaging on a trial basis for up to 90 days.
	If you obtain a Universal Messaging license file after installation, you can add it at that time. For instructions, see the Universal Messaging documentation.
NHP interface binding	If your machine has multiple IP addresses and you want Universal Messaging to listen to a specific one, select that IP address from the list.
NHP interface port	The installer sets up a bootstrap Universal Messaging interface to which all sample Universal Messaging applications and tools connect by default. Specify the port to which to bind the bootstrap interface.
Realm server name	Name for the realm server instance. If you are also installing an Enterprise Manager instance, Template Applications instance, or both, those instances will also use this name. **Note:** Different types of instances can have the same name. However, the name you supply here must be unique among instances of the same type (that is, among instances of realm servers).
Data directory	Full path to the directory in which to store the data for the realm server instance. **Note:** If you use a non-default location for the data directory, you will have to manually migrate the data directory when you upgrade Universal Messaging. If you install multiple realm server instances on the same machine, use a different data directory for each instance.
Configuration	Initial configuration for the Universal Messaging realm server. Select the configuration for webMethods suite use cases.

If you choose to install Universal Messaging Enterprise Manager, Template Applications, or both, without also choosing to install a Universal Messaging realm server, provide the information below.

Field	Entry
Realm name/ client profile	Name for the Enterprise Manager instance, Template Applications instance, or both. You might want to use the name of the realm server instance to which the tools will connect. The list shows the names of any realm server instances that already exist in the installation directory, but you can also type a name. **Note:** Different types of instances can have the same name. However, the name you supply here must be unique among instances of the same type (that is, among instances of Enterprise Managers, and among instances of Template Applications).
Realm server URL (RNAME)	URL for the realm server instance to which the Enterprise Manager instance, Template Applications instance, or both should connect. You can specify a realm server that is already installed or that will be installed later.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Register Daemons and Set File Permissions

On a UNIX system, if you created or installed from a script, or if you were asked whether to use sudo and you chose not to, the installer was not able to register daemons and set file permissions. Perform these tasks now by running the script *Software AG_directory* /bin/ afterInstallAsRoot.sh as a user who is in the sudoers configuration.

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

Important: If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Complete the ARIS PPM Installation

The ARIS PPM Cloud Agent is installed automatically on Windows systems. On Linux systems, you must install it manually, as a daemon. Go to the *Software AG_directory*/ppm/server/bin directory and run the call `./CloudAgentApp.sh install` with sudo or root privileges. Then start Cloud Agent by running the call `./CloudAgentApp.sh start`.

Complete the Terracotta Installation

For instructions on configuring Terracotta when you are using it with webMethods products, see *Using Terracotta with webMethods Products*. For instructions on configuring Terracotta in other contexts, see the Terracotta product documentation.

Uninstall Products

Follow the instructions in *Using the Software AG Installer*, with the additional guidelines below.

If you are going to uninstall Software AG Designer Eclipse, and you want to review or preserve Eclipse installation information, do so before starting the uninstaller, because the uninstaller will delete the entire Software AG Designer installation directory. User-created data that is stored in other directories, such as your workspace, will remain untouched.

If you are going to uninstall ARIS PPM, the Cloud Agent will be uninstalled automatically from Windows systems. From Linux systems, you must uninstall the Cloud Agent daemon manually. After shutting down all products and before starting the uninstaller, go to the *Software AG_directory*/ppm/server/bin directory and run this scripts with sudo or root privileges:

```
./CloudAgentApp.sh remove
```

If you are going to uninstall from a Mac OS or other UNIX system, and you used sudo to register daemons and set file permissions during or after installation, you must also use sudo to unregister daemons and remove file permissions, as follows:

- If you are not going to use an uninstallation script to uninstall, you can choose the uninstaller option to perform this task, or you can perform this task before starting the uninstaller by executing the script *Software AG_directory*/bin/beforeUninstallAsRoot.sh as a user who is in the sudoers configuration.

- If you are going to use an uninstallation script, the uninstaller cannot execute the script because it does not store the sudo password, for security reasons. You must execute the script yourself before starting the uninstaller.

10   Installing and Uninstalling Master Data Management

Software Requirements and Considerations

On Windows systems, install Microsoft Visual C++ Runtime 2005 or higher in any directory on the machine that will host Locate.

Install Products

Start the Installer, Provide General Information, and Choose the Products to Install

As you go through the Software AG Installer, you will be asked for various types of information, such as product license files and database connection parameters. Gather this information before starting the installer.

Read the requirements in *Using the Software AG Installer* (for example, the requirement to create a user account that has the proper privileges for Windows and that is non-root for UNIX). Follow the instructions in that guide to start the installer and provide general information such as proxy server, release to install, installation directory, and how to use the product selection tree.

> **Important:** Unless otherwise stated for a specific product, do not install products from this release into a Software AG directory that contains products from any other release. Unless otherwise stated for a specific product, you cannot mix products from different releases in the same installation directory; if you do so, you will experience problems or be unable to access functionality.

On the product selection tree, select the products to install. You can also select Event Data Store, which is Software AG's packaging of Elasticsearch.

After the product selection tree, the installer displays the language pack selection tree. For information on language packs, see the international operating environments appendix in this guide. The installer then displays panels (GUI mode) or prompts (console mode) that ask for product configuration information. Most are self-explanatory, so this section shows only the panels that require explanation. The information and fields on the prompts is identical to the information and fields on the panels.

> **Important:** Make sure all ports you specify on panels or in response to prompts are not already being used by other products on your machine. The installer cannot detect whether ports are in use when products are shut down, and the shutting down of products is a requirement for running the installer. See the list of default ports used by Software AG products.

Supply Product Configuration Information

Field	Entry
Repository ID	OneData provides a repository in which to store your master data. Provide an ID for the repository (for example, DevRepo). The ID will be used internally in OneData.
Repository name	Provide a name for the repository (for example, OneData Development Repository). The name will appear in the OneData user interface.
Connection Prefix	OneData has three storage areas - Metadata, Work Area, and Release Area. The prefix you specify will be added to the schema connection names (that is, *prefix* _md, *prefix* _wa, and *prefix* _ra). If you override the default value for this field, you will have to edit the database connections for the Work Area and Release Area after you start OneData (see the OneData documentation).
URL	The URL you specify must include the option catalogOptions=1 so that OneData can retrieve database metadata information for various operations (for example, jdbc:wm:oracle://localhost:1521;serviceName=XE;catalogOptions=1)
Database user and **Password**	Provide a different database user and password for each of the three schemas.
Schema name	You must specify schema names if you use a SQL Server RDBMS.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

Important: If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Uninstall Products

Follow the instructions in *Using the Software AG Installer*.

11   **Installing and Uninstalling API Management**

Typical Installation

The Software AG Installer offers typical development installations of products that enable you to perform a certain task. When you select a typical development installation, the installer automatically selects all products you have licensed that make up that installation. You can select more products or deselect products if desired. For CentraSite, the typical development installation is called SOA Governance. For API Portal and API Gateway, the typical development installation is called API Management.

If you want to publish events emitted by Mediator, you must create the MediatorEvents database component. If you want to publish events emitted by API Gateway, you must create the APIGatewayEvents database component. You must also create all of the database components for the Integration Server that hosts Mediator or API Gateway. A database component is a grouping of database objects that is used by one or more products. For complete information about each database component, see "Creating and Dropping Database Components" on page 93.

To create production environments, work with your administrators, Software AG Global Consulting Services, and best practices documentation.

API Gateway and Mediator Host Server

API Gateway runs on Integration Server or Microservices Container. Mediator runs on Integration Server. For information on software requirements, machine preparation, installation, and configuration for these products, see the relevant sections in "Installing and Uninstalling Application Integration, B2B Integration, Dynamic Apps, or Microservices Platform" on page 41. For Microservices Container, when the installer offers a choice of database for Integration Server or Microservices Container, accept the default choice of **Embedded database**.

Software Requirements and Considerations

For CentraSite:

■ If you are going to install on a Linux system, the library libcrypt.so must be present to support the security infrastructure. If you do not see the libcrypt.so file in the /usr/lib64 directory for Linux x86_64, install the rpm package glibc-devel.

■ If you are going to install on a Linux x86_64 system, install the rpm package compat-libstdc++ from your operating system distribution using the vendor instructions.

Prepare Your Machine

Prepare to Install API Portal on a UNIX System

■ Check the setting for shared memory (kernel parameter shmmax) by executing the command `sysctl -a | fgrep kernel.shmmax`. If the value is less than 629145600, log on as root user and increase the value by executing `sysctl -w kernel.shmmax=629145600` or `echo "kernel.shmmax=629145600" >> /etc/sysctl.conf`, then activate the new value by executing `sysctl -p`.

■ Check the settings for the system-wide maximum number of file descriptors (kernel parameter fs.file-max) by executing the command `sysctl -a | fgrep fs.file-max`. If the value is less than 200000, log on as the root user and increase the value by executing `sysctl -w fs.file-max=200000` or `echo "fs.file-max=200000" >> /etc/sysctl.conf`, then activate the new value by executing `sysctl -p`.

■ Check the user, group, and process settings for the maximum number of open file descriptors by executing the command `ulimit -Hn and ulimit -Sn`, where -Hn is the hard limit and -Sn is the soft limit. If the value is less than 200000, log on as a non-root user and increase the value by executing `ulimit -n 200000`. To permanently save this setting for the user, execute:

```
echo "<user name> soft nofile 200000" >> /etc/security/limits.conf
echo "<user name> hard nofile 200000" >> /etc/security/limits.conf
```

■ Enter the host name of the machine on which you are installing products in the DNS of the network or in the file /etc/hosts.

Prepare to Install CentraSite on a UNIX System

■ Make sure you have set sufficient user limits for the shell you use to start the installation and the product daemons. For example, the daemons for the CentraSite database and infrastructure will run out of memory if they are started from the shell with a low user limit for data. If your system policy allows it, Software AG recommends setting the value for coredump, data, file, memory, and threads to unlimited, and the value of nofiles to 8192. For more information about setting and displaying the ulimits, read the man page or ask your system administrator.

■ Set the shmmax parameter to the memory (RAM) that is physically available on your machine. Otherwise you might have problems during CentraSite startup.

Install Products

Start the Installer, Provide General Information, and Choose the Products to Install

As you go through the Software AG Installer, you will be asked for various types of information, such as product license files and database connection parameters. Gather this information before starting the installer.

Read the requirements in *Using the Software AG Installer* (for example, the requirement to create a user account that has the proper privileges for Windows and that is non-root for UNIX). Follow the instructions in that guide to start the installer and provide general information such as proxy server, release to install, installation directory, and how to use the product selection tree.

> **Note:** When you install on a UNIX system, the installer might ask whether to use sudo and related scripts. If you want to have the installer run the script, the user under which you run the installer must be in the sudoers configuration. If you are creating or installing from an installation script, the option to use sudo is not available because the installer does not store the sudo password in scripts, for security reasons. Instead you will perform certain tasks after installation as described later in this chapter.

> **Important:** Unless otherwise stated for a specific product, do not install products from this release into a Software AG directory that contains products from any other release. Unless otherwise stated for a specific product, you cannot mix products from different releases in the same installation directory; if you do so, you will experience problems or be unable to access functionality.

You can install the CentraSite Registry Repository and Application Server Tier on different machines, or in the same directory on the same machine. If you install the components on different machines, or in the same directory on the same machine but at different times, you must install the Registry Repository first.

If you are installing on a Windows system and you want to install CentraSite Application Server Tier and Zementis Predictive Analytics Integrated Server in the same installation directory, you must install the two products in two separate runs of the installer.

On the installation directory panel, the installer auto-detects a host name or IP address for the machine on which you are running the installer. If you are installing CentraSite, you might need to change that value, as follows:

■ If you install the Application Server Tier and Registry Repository on different machines, the value when you install the Registry Repository must specify a host name or IP address that the Application Server Tier can use to access the Registry Repository.

■ If you are installing CentraSite in the cloud, the installer auto-detects the internal IP address of the cloud instance. If you want CentraSite to be accessible from outside the cloud instance (for example, from another cloud instance running other products), replace the auto-detected value with an externally accessible host name or IP address.

On the product selection tree, select the products to install. The licenses for API Gateway, Mediator, and Enterprise Gateway allow only partial functionality of their host server (that is, Integration Server or Microservices Container). Install each product on its own host server. Do not install other products on that host server or those other products might not work properly. See this section in "Installing and Uninstalling Application Integration, B2B Integration, Dynamic Apps, or Microservices Platform" on page 41 for more information about Integration Server and Microservices Container.

On the language pack selection tree, if you select the CentraSite Application Server Tier language pack, you must also select the CentraSite Shared Files language pack.

After the product selection tree, the installer displays the language pack selection tree. For information on language packs, see the international operating environments appendix in this guide. The installer then displays panels (GUI mode) or prompts (console mode) that ask for product configuration information. Most are self-explanatory, so this section shows only the panels that require explanation. The information and fields on the prompts is identical to the information and fields on the panels.

> **Important:** Make sure all ports you specify on panels or in response to prompts are not already being used by other products on your machine. The installer cannot detect whether ports are in use when products are shut down, and the shutting down of products is a requirement for running the installer. See the list of default ports used by Software AG products.

Supply Product Configuration Information

Most product panels or prompts are self-explanatory. This section provides information for those products that require additional explanation.

Integration Server or Microservices Container (Host Server)

Specify the license for API Gateway, Mediator, or Enterprise Gateway, not a license for Integration Server or Microservices Container. To determine where to store data written by API Gateway or Mediator, and the host server, see "Data Storage" on page 95.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Register Daemons and Set File Permissions

On a UNIX system, if you created or installed from a script, or if you were asked whether to use sudo and you chose not to, the installer was not able to register daemons and set file permissions. Perform these tasks now by running the script *Software AG_directory* /bin/ afterInstallAsRoot.sh as a user who is in the sudoers configuration.

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

> **Important:** If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Uninstall Products

Follow the instructions in *Using the Software AG Installer*, with the additional guidelines below.

If you are going to uninstall Software AG Designer Eclipse, and you want to review or preserve Eclipse installation information, do so before starting the uninstaller, because the uninstaller will delete the entire Software AG Designer installation directory. User-created data that is stored in other directories, such as your workspace, will remain untouched.

If you are going to uninstall from a Mac OS or other UNIX system, and you used sudo to register daemons and set file permissions during or after installation, you must also use sudo to unregister daemons and remove file permissions, as follows:

■   If you are not going to use an uninstallation script to uninstall, you can choose the uninstaller option to perform this task, or you can perform this task before starting the uninstaller by executing the script *Software AG_directory*/bin/ beforeUninstallAsRoot.sh as a user who is in the sudoers configuration.

■ If you are going to use an uninstallation script, the uninstaller cannot execute the script because it does not store the sudo password, for security reasons. You must execute the script yourself before starting the uninstaller.

If you are going to uninstall CentraSite:

■ If the CentraSite Registry Repository is on a different machine than the Application Server Tier, uninstall the Application Server Tier first, and then uninstall the Registry Repository.

■ The CentraSite Registry Repository database will not be deleted after uninstallation is complete. There are many circumstances in which you should retain the database after uninstalling CentraSite; for example, when the database contains information that is important for your enterprise. If you do want to delete the CentraSite database after uninstallation is complete, Software AG strongly recommends making a backup copy of its contents first. Then delete the *Software AG_directory*/CentraSite/ data directory.

12   Installing and Uninstalling Mobile Products

Typical Development Installation

The Software AG Installer offers typical development installations of products that enable you to perform a certain task. When you select a typical development installation, the installer automatically selects all products you have licensed that make up that installation. You can select more products or deselect products if desired. For Mobile products, the typical development installation is called Mobile Application Development.

For this typical development installation, you must create certain *database components*. A database component is a grouping of database objects that is used by one or more products. Depending on the products you have licensed, you must create the MobileSupport database component, and all of the database components for the Integration Server that hosts Mobile Support. For complete information about each database component, see "Creating and Dropping Database Components" on page 93.

To create production environments, work with your administrators, Software AG Global Consulting Services, and best practices documentation.

Mobile Support Host Server

Mobile Support runs on Integration Server. For information on software requirements, machine preparation, installation, and configuration for Integration Server, see the relevant sections in "Installing and Uninstalling Application Integration, B2B Integration, Dynamic Apps, or Microservices Platform" on page 41.

Software Requirements and Considerations

Product	Software Requirements and Considerations

Mobile Administrator	■ When installing Mobile Administrator on a physical machine, you need some common development tools. Install those tools by running the appropriate commands below.
	
	■ Red Hat Enterprise Linux / CentOS:
	
	```yum install gettext make gcc gcc-c++ openssl-devel \```
	``` libicu-devel zlib-devel readline-devel autoconf \```
	``` perl-ExtUtils-MakeMaker```
	
	■ Debian:
	
	```apt-get install gettext make gcc g++ libssl-dev curl \```
	``` libicu-dev zlib1g-dev libreadline6-dev autoconf \```
	``` libmodule-build-perl libmodule-install-perl libc6-dev \```
	``` libmysql++-dev libsqlite3-dev build-essential \```
	``` libssl-dev```

Product	Software Requirements and Considerations
	■ SuSE Enterprise Linux Server: ```zypper install gettext-tools make gcc libopenssl-devel \ curl libicu-devel zlib-devel readline-devel autoconf \ perl-ExtUtils-PkgConfig libxslt-devel libMagickWand1 \ libcurl-devel p7zip ImageMagick-Devel libxml2-devel \ libmysqlclient-devel apache2-worker \ apache2-mod_xsendfile mysql-client gcc-c++ mysql``` ■ When installing Mobile Administrator on a virtual machine (used when hosting Mobile Administrator in the cloud or running Mobile Administrator for demos), install Oracle VM VirtualBox or Vagrant.
Mobile Designer	■ Install an integrated development environment (IDE). Software AG recommends using the Eclipse provided with Software AG Designer on the Software AG Installer. ■ If you are going to use Mobile Designer separately from Software AG Designer, install Apache ANT 1.8 or higher using the instructions in the vendor documentation.

Install Products

Start the Installer, Provide General Information, and Choose the Products to Install

As you go through the Software AG Installer, you will be asked for various types of information, such as product license files and database connection parameters. Gather this information before starting the installer.

Read the requirements in *Using the Software AG Installer* (for example, the requirement to create a user account that has the proper privileges for Windows and that is non-root for UNIX). Follow the instructions in that guide to start the installer and provide general information such as proxy server, release to install, installation directory, and how to use the product selection tree.

Important: Unless otherwise stated for a specific product, do not install products from this release into a Software AG directory that contains products from any other release. Unless otherwise stated for a specific product, you cannot mix products from different releases in the same installation directory; if you do so, you will experience problems or be unable to access functionality.

On the product selection tree, select the products to install. See this section in "Installing and Uninstalling Application Integration, B2B Integration, Dynamic Apps, or Microservices Platform" on page 41 for more information about Integration Server.

After the product selection tree, the installer displays the language pack selection tree. For information on language packs, see the international operating environments appendix in this guide. The installer then displays panels (GUI mode) or prompts (console mode) that ask for product configuration information. Most are self-explanatory, so this section shows only the panels that require explanation. The information and fields on the prompts is identical to the information and fields on the panels.

> **Important:** Make sure all ports you specify on panels or in response to prompts are not already being used by other products on your machine. The installer cannot detect whether ports are in use when products are shut down, and the shutting down of products is a requirement for running the installer. See the list of default ports used by Software AG products.

Supply Product Configuration Information

Most product panels or prompts are self-explanatory. This section provides information for those products that require additional explanation.

Integration Server (Mobile Support Host)

To determine where to store data written by Integration Server, see " Integration Server or Microservices Container Data Storage" on page 95.

My webMethods Server

To determine where to store data written by My webMethods Server, see " My webMethods Server Data Storage" on page 98.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Complete the Mobile Administrator Installation

This step varies depending on whether you installed Mobile Administrator on a physical or virtual machine.

■ If you installed on a physical machine, go to the *Software AG_directory*/ MobileAdministrator directory and run the command `./install.sh` as root. If the installation is successful, you will see the message `Notice: Finished catalog run`

in 2594.71 seconds, and the installed Mobile Administrator will be available via HTTP or HTTPS at the IP address of the host machine. If an error occurs, you will see the complete error message.

If Mobile Administrator cannot connect to the Apache web server on ports 80 (HTTP) and 443 (HTTPS), make sure your firewall is configured correctly. If the host machine has SELinux enabled, go to the /var/www/appfresh/log directory and open the apache_error.log file. If you see error messages similar to (13) Permission denied: proxy: HTTP: attempt to connect to 127.0.0.1:5000 (localhost) failed, run the command below as root:

```
# setsebool -P httpd_can_network_connect 1
```

◾ If you installed on a virtual machine, the Mobile Administrator virtual installer is a Vagrantfile. The installer sets up a virtual machine with a Linux distribution and installs Mobile Administrator. Go to the *Software AG_directory*/MobileAdministrator directory and execute Vagrantfile by running the appropriate command below.

Distribution	Command
Debian 6.0.7	`vagrant up standalone`
SuSE Enterprise Linux Server 11	`APPFRESH_LINUX=sles vagrant up standalone`
CentOS 6.5	`APPFRESH_LINUX=centos vagrant up standalone`
Distribution	Command

If the installation is successful, you will see the message Notice: Finished catalog run in 2594.71 seconds, and the installed Mobile Administrator will be available at https://192.168.10.10. If an error occurs, you will see the complete error message.

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

> **Important:** If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Uninstall Products

Uninstall Mobile Administrator

To uninstall Mobile Administrator from a virtual machine, follow the Oracle VM VirtualBox or Vagrant instructions.

To uninstall Mobile Administrator from a physical machine, follow the instructions below.

1. Back up your data files (for example, application binaries or screenshots that have been generated by Mobile Administrator). The data files are stored in the /var/www/appfresh/public/datafiles directory.

2. Export your MySQL database. For instructions, see the vendor documentation.

3. Go to the *Software AG_directory* /MobileAdministrator directory.

4. Stop Mobile Administrator by running the command `/etc/init.d/appfresh-mobile stop`

5. Remove the appfresh directory by running the command `rm -rf /var/www/appfresh`

6. Remove the appfresh-mobile directory by running the command `rm -rf /var/run/appfresh-mobile`

7. Remove the database by running the command `mysql -uroot -e 'drop database appfreshmobile'`

8. Remove the Apache virtual host by running the command `rm /etc/apache2/sites-enabled/25-appfresh*`

Uninstall All Other Products

Follow the instructions in *Using the Software AG Installer*, with the additional guidelines below.

If you are going to uninstall Software AG Designer Eclipse, and you want to review or preserve Eclipse installation information, do so before starting the uninstaller, because the uninstaller will delete the entire Software AG Designer installation directory. User-created data that is stored in other directories, such as your workspace, will remain untouched.

13 Installing and Uninstalling webMethods Broker

Software Requirements and Considerations

■ On Solaris, HP-UX, and AIX systems, Software AG offers 64-bit and 32-bit (LP32 Object Code) versions of the webMethods Broker C API. If you select the C API for installation, the installer installs both versions, in the *Software AG_directory*/Broker/lib and /lib32 directories, respectively.

■ If you are going to install on a Solaris 11 system, install the system/xopen/xcu4 package.

■ If you are going to install on a Linux system, the library libcrypt.so must be present to support the security infrastructure. If you do not see the libcrypt.so file in the /usr/lib64 directory for Linux x86_64, install the rpm package glibc-devel.

■ If you are going to install on a Linux x86_64 system, install the rpm package compat-libstdc++ from your operating system distribution using the vendor instructions.

Prepare Your Machine

If you are installing webMethods Broker on a UNIX system:

■ Information about webMethods Broker installations is stored in the webMethods Broker installations.txt file in the /var/opt/webmethods directory. Make sure the user that you are going to use to install has write permission to that directory.

■ Broker Monitor and Broker Server require certain minimum system and user limits on UNIX systems. Ask your system administrator to set the following values for each process:

Limit	Value
Maximum threads	512
Maximum open files	8192
Core dump size	unlimited

■ If you are going to configure Broker Server to create data files on an NFS-mounted partition, mount the partition using the command below. The options on the command help protect data integrity.

```
mount -o hard,nointr,proto=tcp host:/path /mount_point
```

Option	Description
`hard`	By default, if the NFS server does not respond to requests from the Broker Server, the Broker Server tries a few times and then fails. The `hard` option forces the Broker Server to keep retrying until the NFS server responds. The Broker Server hangs during this period.
`nointr`	Prevents users from shutting down the Broker Server while it waits for the NFS server to respond.
`proto=tcp`	By default, the NFS-mounted partition uses the user datagram protocol (UDP) protocol to communicate with the Broker Server. The proto=tcp option forces the partition to use transmission control protocol (TCP) instead. TCP is more reliable than UDP.

Suppose the partition you want to mount is on a machine named netappca, the data directory on the partition is wmbroker_data, and you want to mount the partition on the local file system at /var/opt/wmbroker_data. The command would be as follows:

```
mount -o hard,nointr,proto=tcp netappca:/wmbroker_data /var/opt/wmbroker_data
```

For information on other ways to protect data integrity, see your NFS server documentation.

■ If you are going to install Broker Server on an NFS file system, the Basic Authentication feature will only work if you mount the NFS file system with the suid option, which allows set-user-identifier or set-group-identifier bits to take effect.

Install Products

Start the Installer, Provide General Information, and Choose the Products to Install

Read the requirements in *Using the Software AG Installer* (for example, the requirement to create a user account that has the proper privileges for Windows and that is non-root for UNIX). Follow the instructions in that guide to start the installer and provide general information such as proxy server, release to install, installation directory, and how to use the product selection tree.

Important: Unless otherwise stated for a specific product, do not install products from this release into a Software AG directory that contains products from any other release. Unless otherwise stated for a specific product, you cannot mix products from different releases in the same installation directory; if you do so, you will experience problems or be unable to access functionality.

On the product selection tree, select the products to install.

After the product selection tree, the installer displays the language pack selection tree. For information on language packs, see the international operating environments appendix in this guide. The installer then displays panels (GUI mode) or prompts (console mode) that ask for product configuration information. Most are self-explanatory, so this section shows only the panels that require explanation. The information and fields on the prompts is identical to the information and fields on the panels.

> **Important:** Make sure all ports you specify on panels or in response to prompts are not already being used by other products on your machine. The installer cannot detect whether ports are in use when products are shut down, and the shutting down of products is a requirement for running the installer. See the list of default ports used by Software AG products.

Supply Product Configuration Information

When you install webMethods Broker, you install the following:

- A default Broker. Brokers execute client messaging requests.

- Optionally, a Broker Server. Broker Servers host Brokers; they receive client requests, send them to Brokers, and return responses to clients. They also manage memory and disk resources for the Brokers they host.

- A Broker Monitor. Broker Monitors continually check the state of Broker Servers and automatically restarts them if they stop running.

Field	Entry
IP address to bind to	By default, Broker Monitor will bind to all IP addresses on the local machine. If necessary, bind Broker Monitor to a specific IP address instead.
Create a Broker Server Configuration	Every Broker Server has it own data directory, which holds theBroker Server's configuration file and log files, and storage session, which stores configuration (metadata) and run-time data.
	You would *not* create the Broker Server configuration at this time (that is, using the installer) if you want to use a different storage size than is offered by the installer or if you want to create a combined storage session for configuration and run-time data. If you do not need to back up configuration data without shutting down your Broker Server, using a combined session might save you a small amount of disk space. If you choose to not create the Broker Server configuration at this time, you must create it

Field	Entry
	after installation is complete, using instructions in *Administering webMethods Broker*.
	You would create the Broker Server configuration at this time (that is, using the installer) if you want to use a storage size that is offered by the installer and you want to create separate storage sessions for configuration (metadata) and run-time data. Using separate storage sessions minimizes the risk of corruption that goes with a combined storage location and enables you to back up configuration data without having to shut down your Broker Server.
	Software AG recommends creating separate sessions. You cannot later change from a combined session to separate sessions or vice versa.
	To create the Broker Server configuration, select the check box and complete the fields below.
Data directory	Full path to the directory for Broker Server data. If you install multiple Broker Server instances on the same machine, use a different data directory for each instance.
Storage size	Select a pre-configured storage session for the Broker Server that can handle your expected usage needs.
	■ **Small**. Ideal for running development Broker Servers or small number of production integrations, low document volumes, and no document logging. Fastest Broker Server startup time.
	■ **Medium**. Standard deployment size, fits more cases than **Small**; larger maximum transaction size and twice the storage capacity of **Small**. Broker Server startup time two times longer than **Small**.
	■ **Large**. Suitable for production deployments with many integrations running at high document volumes, possibly using document logging as well. Broker Server startup time two times longer than **Medium**, four times longer than **Small**.

When determining the appropriate size for the log file, the factors to balance are Broker Server startup time and the desired maximum transaction size. The smaller the log file, the faster the startup; however, with a larger log file, you can send larger messages (that is, one larger-sized single document or a batch of documents). If necessary, you can remove or replace log files after installation, and you can increase or decrease their size. Startup time does *not* depend on the size of the storage file; additional storage capacity merely prevents the Broker Server from running out of room. After installation, if the amount of storage allocated turns out to be insufficient, you can add storage files and increase their size. You cannot remove storage files or decrease their size. *Administering*

webMethods Broker provides complete information on Broker Server storage sessions and instructions on working with log files and storage files.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Make Sure Broker Server is Running and the Default Broker Exists

After installation, Broker Monitor starts automatically and then starts the Broker Server, which begins running with a default Broker. Make sure the Broker Server is running and the default Broker exists by running this command:

```
broker_status [Broker #1@]Broker Server_host[:Broker Server_port]
```

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

> **Important:** If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Uninstall Products

Follow the instructions in *Using the Software AG Installer*, with the additional guidelines below.

If your Broker Monitor was configured to run as a UNIX daemon, you must un-register the daemon after you uninstall webMethods Broker. If you used the command line tool daemon.sh to register the daemon, use that tool again to un-register the daemon (see "Register Daemons to Automatically Start and Shut Down Products on UNIX Systems" on page 131 for instructions). If you registered the daemon manually, revert your manual registration steps.

14 Installing and Uninstalling ApplinX

Typical Development Installation

The Software AG Installer offers typical development installations of products that enable you to perform a certain task. When you select a typical development installation, the installer automatically selects all products and components that make up that installation. You can select more products or deselect products if desired. The typical development installation for this product is called ApplinX.

To create production environments, work with your administrators, Software AG Global Consulting Services, and best practices documentation.

Software Requirements

■ If you are going to install on a Windows system, and you want to develop .NET ApplinX Web applications in Visual Basic or C#, install Microsoft Visual Studio 2005, or install Microsoft Visual Studio 2008 or 2010 and the Microsoft Visual J# .NET 2.0 Redistributable Package - Second Edition (x64). Microsoft Visual Studio does not have to be running when you install ApplinX.

   If you do not install Microsoft Visual Studio, you will not be able to install the ApplinX Visual Studio Add-in component, and the ApplinX .NET Framework component will have limited functionality; you will not be able to develop composite Web applications.

■ If you are going to install on a Linux x86_64 system, install the rpm package compat-libstdc++ from your operating system distribution using the vendor instructions.

Prepare Your Machine

If you are going to install on a UNIX system, make sure you have set sufficient user limits for the shell you use to start the installation and the product daemons. For example, the daemons for the Software AG Runtime will run out of memory if they are started from the shell with a low user limit for data. If your system policy allows it, Software AG recommends setting the value for coredump, data, file, memory, and threads to unlimited, and the value of nofiles to 8192. For more information about setting and displaying the ulimits, read the man page or ask your system administrator.

Install ApplinX

Start the Installer, Provide General Information, and Choose ApplinX Components

Read the requirements in *Using the Software AG Installer* (for example, the requirement to create a user account that has the proper privileges for Windows and that is non-root for UNIX). Follow the instructions in that guide to start the installer and provide general information such as proxy server, release to install, installation directory, and how to use the product selection tree.

> **Important:** Unless otherwise stated for a specific product, do not install products from this release into a Software AG directory that contains products from any other release. Unless otherwise stated for a specific product, you cannot mix products from different releases in the same installation directory; if you do so, you will experience problems or be unable to access functionality.

On the product tree, choose the ApplinX components to install.

After the product selection tree, the installer displays the language pack selection tree. For information on language packs, see the international operating environments appendix in this guide. The installer then displays panels (GUI mode) or prompts (console mode) that ask for product configuration information. Most are self-explanatory, so this section shows only the panels that require explanation. The information and fields on the prompts is identical to the information and fields on the panels.

> **Important:** Make sure all ports you specify on panels or in response to prompts are not already being used by other products on your machine. The installer cannot detect whether ports are in use when products are shut down, and the shutting down of products is a requirement for running the installer. See the list of default ports used by Software AG products.

Supply Product Configuration Information

If you are installing the ApplinX Server, the installer ask for information such as licence file and ports.

> **Note:** The **Migrate data** field is for use with pre-8.1 release upgrades only. For complete information, see the ApplinX upgrade documentation.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

> **Important:** If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Uninstall ApplinX

Follow the instructions in *Using the Software AG Installer*, with the additional guidelines below.

If you are going to uninstall Software AG Designer Eclipse, and you want to review or preserve Eclipse installation information, do so before starting the uninstaller, because the uninstaller will delete the entire Software AG Designer installation directory. User-created data that is stored in other directories, such as your workspace, will remain untouched.

15 Installing and Uninstalling EntireX

Typical Development Installation

The Software AG Installer offers typical development installations of products that enable you to perform a certain task. When you select a typical development installation, the installer automatically selects all products and components that make up that installation. You can select more products or deselect products if desired. The typical development installation for this product is called EntireX.

To create production environments, work with your administrators, Software AG Global Consulting Services, and best practices documentation.

Software Requirements

- If you are going to install on a Linux system, the library libcrypt.so must be present to support the security infrastructure. If you do not see the libcrypt.so file in the /usr/lib64 directory for Linux x86_64, install the rpm package glibc-devel.

- If you are going to install on a Linux x86_64 system, install the rpm package compat-libstdc++ from your operating system distribution using the vendor instructions.

- If you are going to install on a SUSE Linux Enterprise Server 11 SP1 system, and you have configured the system to use the Blowfish algorithm to encrypt user passwords, download and install patch 18 Aug 2011 - glibc 4944 from Novell Support. This patch fixes an operating system bug relating to user authentication. For detailed information, see Novell's SUSE Security Announcement: SUSE-SA:2011:035.

Prepare Your Machine

- If you are going to install EntireX on a UNIX system, make sure you have set sufficient user limits for the shell you use to start the installation and the product daemons. If your system policy allows it, Software AG recommends setting the value for coredump, data, file, memory, and threads to unlimited, and the value of nofiles to 8192. For more information about setting and displaying the ulimits, read the man page or ask your system administrator.

- If you are going to install EntireX on a UNIX NFS file system and want to do user authentication for EntireX Broker against the local operating system, you must mount the NFS file system with the suid option, which allows set-user-identifier or set-group-identifier bits to take effect.

Install EntireX

Start the Installer, Provide General Information, and Choose EntireX Components

Read the requirements in *Using the Software AG Installer* (for example, the requirement to create a user account that has the proper privileges for Windows and that is non-root for UNIX). Follow the instructions in that guide to start the installer and provide general information such as proxy server, release to install, installation directory, and how to use the product selection tree.

Note:	When you install on a UNIX system, the installer might ask whether to use sudo and related scripts. If you want to have the installer run the script, the user under which you run the installer must be in the sudoers configuration. If you are creating or installing from an installation script, the option to use sudo is not available because the installer does not store the sudo password in scripts, for security reasons. Instead you will perform certain tasks after installation as described later in this chapter.
Important:	Unless otherwise stated for a specific product, do not install products from this release into a Software AG directory that contains products from any other release. Unless otherwise stated for a specific product, you cannot mix products from different releases in the same installation directory; if you do so, you will experience problems or be unable to access functionality.

On the product tree, select the EntireX components to install.

Component	Description
Core Files	Runtime for EntireX clients and servers.
Administration and Monitoring	Enables you to administer and monitor EntireX using scripts and utilities. Also includes the Application Monitoring Data Collector.
Broker	Communication hub for EntireX clients and servers.
Development Tools	C header files and Microsoft Visual Studio Wizard for EntireX .NET Wrapper. Enables you to develop .NET and C-based EntireX applications and work with Microsoft Visual Studio.
Mini Runtime	Runtime for EntireX clients written in .NET, C, and DCOM.

Component	Description
Web Services deployed in Software AG Runtime	XML Listener and XML RPC Server deployed in Software AG Runtime. Enables you to provide and access Web services from EntireX RPC endpoints using Software AG Runtime.

After the product selection tree, the installer displays the language pack selection tree. For information on language packs, see the international operating environments appendix in this guide. The installer then displays panels (GUI mode) or prompts (console mode) that ask for product configuration information. Most are self-explanatory, so this section shows only the panels that require explanation. The information and fields on the prompts is identical to the information and fields on the panels.

> **Important:** Make sure all ports you specify on panels or in response to prompts are not already being used by other products on your machine. The installer cannot detect whether ports are in use when products are shut down, and the shutting down of products is a requirement for running the installer. See the list of default ports used by Software AG products.

Supply Product Configuration Information

EntireX

Field	Entry
License file	Full path to the EntireX license file. **Note:** If you are also installing NaturalONE, EntireX uses the NaturalONE license file.
Default EntireX Broker TCP/IP and SSL ports, and **Autostart**	When you install EntireX Broker: ▪ TCP/IP and SSL ports to use for the Broker. ▪ Whether the installer should start the default Broker after installation and turn on the Autostart option so the Broker will start automatically whenever you start your system.
Application Monitoring Data Collector port and **Autostart**	▪ Port to use for the Application Monitoring Data Collector. ▪ Whether the installer should start the Data Collector automatically after installation and whenever you start your system.

Field	Entry
EntireX Broker Administration Service TCP/IP and SSL ports	When you install EntireX Broker, TCP/IP and SSL ports to use to administer the Broker from Command Central.

Complete the Installation

Install Latest Fixes

Install the latest fixes on the products you installed. For instructions on using the Software AG Update Manager, see *Using the Software AG Update Manager*. Each fix includes a readme file that contains instructions on installing the fix.

Register Daemons and Set File Permissions

If you installed the EntireX Broker on a UNIX system and did not use sudo:

1. The installer was not able to register daemons and set file permissions for EntireX Broker. Perform this task now by executing the script *Software AG_directory*/bin/ afterInstallAsRoot.sh as a user who is in the sudoers configuration.

   Important: You must run this script or some products might not work correctly.

2. Restart the EntireX Broker Admin Service. To do so, run the command `/etc/ init.d/sag`*number*`exx`*release_number* `restart`, where *number* matches the number in the sagexx script in your *Software AG_directory*/EntireX/bin directory.

3. During installation, you might have told the installer to enable autostart so the EntireX Broker would start automatically whenever you start your system. However, since you did not use sudo to install, the installer could not actually enable autostart. If you want to enable autostart, see the section on administering EntireX Broker in the *Software AG Command Central Help*.

Start, Configure, and Customize Products

For instructions on starting, configuring, and customizing products, see the product documentation.

Important: If any product you installed has a default password, you should change that password as soon as possible. For instructions, see the product documentation.

Uninstall EntireX

Follow the instructions in *Using the Software AG Installer*, with the additional guidelines below.

When shutting down products, also shut down all EntireX Brokers, RPC Servers, and customer applications that use EntireX libraries. For instructions, see the product documentation.

If you are going to uninstall Software AG Designer Eclipse, and you want to review or preserve Eclipse installation information, do so before starting the uninstaller, because the uninstaller will delete the entire Software AG Designer installation directory. User-created data that is stored in other directories, such as your workspace, will remain untouched.

If you are going to uninstall from a UNIX system, and you used sudo to register daemons and set file permissions during or after installation, you must also use sudo to unregister daemons and remove file permissions, as follows:

■ If you are not going to use an uninstallation script to uninstall, you can choose the uninstaller option to perform this task, or you can perform this task before starting the uninstaller by executing the script *Software AG_directory*/bin/ beforeUninstallAsRoot.sh as a user who is in the sudoers configuration.

■ If you are going to use an uninstallation script, the uninstaller cannot execute the script because it does not store the sudo password, for security reasons. You must execute the script yourself before starting the uninstaller.

16   Creating and Dropping Database Components

Database Components

If you installed the products below, you must create *database components* for them.

- ActiveTransfer
- API Gateway
- CloudStreams
- Integration Server
- Mediator
- Microservices Container
- Mobile Support
- My webMethods Server
- OneData
- Optimize
- Process Engine
- Rules Engine
- Trading Networks

A database component is a grouping of database objects that can be used by one or more products. For example, Integration Servers write service status data to the ISCoreAudit database component, and Monitor reads the data and displays it in the My webMethods user interface. Each database component is named for the type of data that is written to it.

DB2, MySQL Enterprise Edition, Oracle, and SQL Server Database Driver

Products listed above that support DB2, Oracle, or SQL Server use Java Database Connectivity (JDBC) and the DataDirect Connect JDBC 5.1 database driver to communicate with those RDBMSs. The products come with the client component of this driver; it is a Type 5 JDBC native-protocol driver that does not have a server component. The products use the driver to convert JDBC requests from Java programs directly into the vendor-specific database protocol that each type of RDBMS can understand. The products execute DDL commands at install time and sometimes design time for each RDBMS with which they interact. For information about the DataDirect Connect JDBC driver, see the DataDirect Connect documentation, available on the Software AG Documentation website.

> **Important:** You must use the provided DataDirect Connect JDBC database driver with your products and DB2, Oracle, and SQL Server RDBMSs unless the documentation for a product states that it supports another driver.

MySQL Community Edition Database Driver

Products listed above that support MySQL Community Edition use the MySQL Community Edition database driver to communicate with the RDBMS. For information about the driver, see the vendor documentation.

> **Note:** If you experience any issues with the MySQL native driver, or with fixes, service packs, or patch-bundles for the driver, report them directly to the vendor. Software AG is not responsible for addressing driver-related issues.

1. Download the MySQL Community Edition driver to the *Software AG_directory* \common\lib\ext directory.

2. Go to the *Software AG_directory* \common\db\bin directory, open the setEnv.bat file, and add the MySQL Community Edition driver jar to the classpath, as follows:

```
set CLASSPATH=%CLASSPATH%;%DCI_HOME%\..\lib\ext\driver_name
```

For example:

```
set CLASSPATH=%CLASSPATH%;%DCI_HOME%\..\lib\ext\mysql-connector-java.jar
```

Data Storage

Integration Server or Microservices Container Data Storage

Integration Server or Microservices Container can persist the types of data below.

> **Note:** For detailed information on product concepts mentioned below, see the relevant product documentation.

Database Component	Types of Data	Data written when...
ISInternal	Service results, scheduled tasks, client certificate mappings, run-time data for pub.storage services, guaranteed delivery transactions, trigger joins, active OpenID requests, WS-ReliableMessaging runtime data,	You are using the features listed in the **Types of Data** column

Database Component	Types of Data	Data written when...
	and configuration and runtime data for OAuth.	
ISCoreAuditLog	■ Error, guaranteed delivery, service, security, and session audit data.	The audit logger for the type of data is enabled
	■ Documents that are in doubt, have failed, or have exhausted trigger retries.	You are using triggers
CrossReference	Cross-referencing data for publish-and-subscribe solutions.	You are using publish-and-subscribe solutions
DocumentHistory	Document history data for exactly-once processing in publish-and-subscribe solutions. Integration Server uses the data to detect and reject duplicate documents.	You are using exactly-once processing
DistributedLocking	Information that coordinates access to resources across distributed servers and processes.	Executing services in the pub.storage folder.

Embedded Database versus External RDBMS

When you install Integration Server or Microservices Container, the installer asks whether you want to write data to an external RDBMS or an embedded database. You must use an external RDBMS if you are going to:

■ Cluster server instances.

■ Install in a production environment with high transaction volumes or large datasets.

■ Need a database that is scalable, visible, and highly reliable, and that has robust management and recovery tools.

■ Write document history or cross-referencing data.

■ Have services log their input pipelines, or post user-defined progress messages.

■ Use triggers for joins.

■ Use pub.storage for high-volume storage.

■ In the case of Integration Server:

■ Write Process Audit Log and Process Engine data for business processes.

■ Use Mediator or Monitor.

■ In the case of Integration Server or Microservices Container: Use Business Rules.

If none of the above apply, and your database demands are low, your needs might be met by using the embedded database. For example, in the case of Integration Server or Microservices Container, you can use the embedded database as a runtime for adapters, eStandards Modules, or Deployer. You should use the embedded database with an Integration Server that is an Enterprise Gateway.

Using the Embedded Database

Integration Server and Microservices Container use Derby, a file-based database, as their embedded database. If one database file gets corrupted, the entire database might become unusable. Software AG therefore recommends backing up the *Software AG_directory*\IntegrationServer\db directory periodically so you will be able to return to the last known good state if necessary.

When you choose to use the embedded database, Integration Server and Microservices Container write IS Internal and Cross Reference data to that database, and write IS Core Audit Log data to files.

If you later want to write these types of data to an external RDBMS instead, you will need to create the necessary database components in the RDBMS (see "Product Database Component Descriptions and Installation Requirements" on page 99) and configure Integration Server or Microservices Container to write to them (see "Connect Products to Database Components" on page 126). You must also complete the steps for switching from the embedded database to an external RDBMS as described in the *webMethods Integration Server Administrator's Guide*.

Using the External RDBMS

When you use an external RDBMS, you must create the ISCoreAudit, ISInternal, and DistributedLocking database components in the RDBMS. You must create the ISInternal and DistributedLocking database components in the same schema (Oracle) or database (DB2, MySQL, or SQL Server). You will also create other database components as needed. For example, if you are using publish-and-subscribe solutions and exactly-once processing, you will need the CrossReference and DocumentHistory database components; if you are running business processes, you will need the ProcessAudit and ProcessEngine database components. For complete information about the database components you will need to create, see "Product Database Component Descriptions and Installation Requirements" on page 99.

When you choose to use an external RDBMS, the installer asks you to supply the database connection. From this information, the installer creates a JDBC connection pool and configures Integration Server or Microservices Container to write IS Internal, IS

Core Audit Log, Cross Reference, Distributed Locking, Document History, Process Audit Log, and Process Engine data to the external RDBMS using that pool.

> **Note:** This auto-configuration is for ease of use only. You need not create all the corresponding database components; only create those you will actually use. Also, in a production environment, you might want to create additional JDBC connection pools and reconfigure the product to write different types of data using different pools. "Connect Products to Database Components" on page 126 provides instructions.

If you are clustering server instances, create an ISCoreAudit database component, an ISInternal database component, and a DistributedLocking database component in the same schema (Oracle) or database (DB2, MySQL, or SQL Server) for the cluster to share. Also create a CrossReference database component and a DocumentHistory database component for the cluster to share.

If you are not clustering server instances, create an ISInternal database component and a Distributed Locking database component for each server instance. Also create the ISCoreAudit, CrossReference, and DocumentHistory database components; you can create one of each of those database components for all server instances to share, or you can create each of those database components for each server instance.

Archiving Data

See "Archive Database Component" on page 100.

My webMethods Server Data Storage

My webMethods Server and Task Engine persist the types of data below.

- My webMethods Server writes data about deployment, configuration, security, portal pages, and run-time operations. It uses this data to manage the My webMethods user interfaces (for example, the user interfaces for webMethods Broker, Monitor, Optimize, and Trading Networks) and to support central user management in Integration Server and Optimize.

- Task Engine writes task status, task distribution, and business data. The Task Engine user interface displays the data.

Embedded Database versus External RDBMS

When you install My webMethods Server, the installer asks whether you want My webMethods Server (and Task Engine) to write data to an embedded database or an external RDBMS. You must use an external RDBMS if you are going to:

- Install My webMethods Server in a production environment.

- Cluster My webMethods Servers.

- Use Business Console.

If none of the above apply, you can use the embedded database when you are installing My webMethods Server for the sole purpose of running the Broker Messaging user interface.

> **Important:** If you choose the embedded database during installation, you cannot later switch to write data to an external RDBMS.

Using the Embedded Database

My webMethods Server uses Derby, a file-based database, as its embedded database. If one database file gets corrupted, the entire database might become unusable. Software AG therefore recommends backing up the *Software AG_directory*\MWS\server *server_instance* \data\db folder regularly so you will be able to return to the last known good state if necessary. If you do not take a backup, you can restore the database to an empty state from the *Software AG_directory*\MWS\server\template-derby.zip\data\db directory.

Using an External RDBMS

When you use an external RDBMS, you must create the MywebMethodsServer database components in your external RDBMS. When you choose the external RDBMS in the installer, the installer asks you to supply the database connection.

If you are clustering My webMethods Servers, create a single MywebMethodsServer database component for the cluster to share. If you are not clustering, create a MywebMethodsServer database component for each My webMethods Server.

> **Note:** The CentralConfiguration database component is automatically created when you create the MywebMethodsServer database component. If you are using Optimize, the CentralConfiguration database component stores the Optimize configuration information you enter in the Central Configurator interface in My webMethods. The CentralConfiguration database component is not automatically *dropped* when you drop the MywebMethodsServer database component. If you want to drop the CentralConfiguration database component, you must do so manually.

Product Database Component Descriptions and Installation Requirements

This section briefly describes each database component and its installation requirements.

> **Note:** For detailed information on product concepts mentioned in the sections below, see the relevant product documentation.

APIGatewayEvents Database Component

API Gateway comes with an internal datastore in which it stores events and metrics by default. If you want to store events and metrics in an external RDBMS instead, create the

APIGatewayEvents database component, then configure API Gateway as described in the product documentation.

ActiveTransfer Database Component

If you installed ActiveTransfer Server, you must create the ActiveTransfer database component.

ActiveTransfer Server writes listener (port), server, user configuration, and other metadata to this database component. ActiveTransfer Server also writes run-time data, such as file transfer transaction information and audit data. If you are clustering ActiveTransfer Servers, create a single ActiveTransfer database component for the cluster to share. If you are not clustering, create an ActiveTransfer database component for each ActiveTransfer Server.

Archive Database Component

Archive Database Component

If you want to archive data from the ISCoreAudit and ProcessAudit database components (called the "source" database components in this section), create the Archive database component.

Note:	You might be able to use partitioning for the ProcessAudit database component. Contact Software AG Global Consulting Services for information.

The Archive database component includes a copy of the ISCoreAudit and ProcessAudit database components. You must create the Archive database component in the same type of RDBMS and on the same database server as the source database components, but in a different schema (Oracle) or database (DB2, MySQL, or SQL Server) than the source database components.

OperationManagement Database Component

The OperationManagement database component is automatically created when you create the Archive database component. It provides a common method for configuration, logging, and error handling for the Archive database component.

BPM (Dynamic Applications)

BusinessRules Database Component

If you installed the Rules Engine, create the BusinessRules database component. You must create the database component in the same schema (Oracle) or database (DB2, MySQL, or SQL Server) as the ProcessAudit database component.

As each Rules Engine instance starts, it registers itself in this database component and stores information about deployed projects and the state of business rules instances. When you modify a business rule, the hot deployment functionality in the Business

Rules user interface enables you to deploy changes to all of the instances registered in this database component.

If you install multiple Rules Engine instances, create a single BusinessRules database component for the instances to share.

ProcessAudit Database Component

If you installed the Process Engine, create the ProcessAudit database component.

The following write to this database component:

■   Process Engines write process audit data for business processes they orchestrate.

■   Task Engines write task audit data.

■   Third-party products can write process execution data.

The following read from this database component:

■   Optimize Analytic Engines read process data from this database component so they can analyze capabilities such as transition duration and write data about analysis-enabled processes, then displays this data in the Optimize interface in My webMethods.

■   Monitor reads process data from this database component and displays it in the Monitor interface in My webMethods, where you can track the status of process instances, view logged values, and, in some cases, resubmit process instances.

■   Business Console reads and displays process data from this database component so you can monitor process instances in real time.

If you are distributing business process steps, you cluster the Process Engines that orchestrate the steps. Create a single ProcessAudit database component for the cluster to share. Integration Servers that host these Process Engines register themselves in the shared ProcessAudit database component.

If you are not distributing business process steps, and therefore not clustering Process Engines, you can create either a ProcessAudit database component for each Process Engine or a single ProcessAudit database component for all Process Engines to share.

Create a single ProcessAudit database components for all Task Engines to share.

If you are using Process Engines, Task Engines, Optimize Analytic Engines, or some combination of these, create a single ProcessAudit database component for all to share.

ProcessEngine Database Component

If you installed the Process Engine, create the ProcessEngine database component. Process Engines write process execution data for processes they orchestrate to this database component.

If you are distributing business process steps, you cluster the Process Engines that orchestrate the steps. Create a single ProcessEngine database component for the cluster to share. If you are not clustering, create a ProcessEngine database component for each Process Engine.

Archiving Data

See "Archive Database Component" on page 100.

CloudStreamsEvents Database Component

If you installed CloudStreams, create the CloudStreamsEvents database component.

CloudStreams Server writes lifecycle (startup/shutdown), error, policy violation, monitoring, performance metric, and transaction events to the CloudStreamsEvents database component. CloudStreams Analytics reads the events data and displays it using the MashZone NextGen dashboard.

If you have multiple CloudStreams Server instances, regardless of whether they are clustered or not, create a single CloudStreamsEvents database component for them to share.

Software AG Designer

Staging and Reporting Database Components

If you want to simulate business processes in Software AG Designer using historical data, create the Staging and Reporting database components. Services you run will extract the historical data from the ProcessAudit database component and aggregate and load the data using the Staging and Reporting database components.

You can create the Reporting and Staging database components in the same schema (Oracle) or database (DB2 , MySQL, or SQL Server). However, the services that aggregate and load the historical data consume large amounts of the CPU. For best performance, install the Reporting and Staging database components on a database server that will not adversely affect the ProcessAudit or ProcessEngine database components.

The ProcessAudit, Staging, and Reporting database components have a 1-to-1-to-1 relationship. This means that you cannot gather data from multiple ProcessAudit database components into a single Reporting database component. Simulation can be performed from only one Reporting database component at a time.

Integration Server or Microservices Container

See " Integration Server or Microservices Container Data Storage" on page 95.

MediatorEvents Database Component

If you installed Mediator, create the MediatorEvents database component in the same schema (Oracle) or database (DB2, MySQL, or SQL Server) as the ISCoreAudit database component. Mediator writes events about certain SOA policies to the MediatorEvents database component.

MobileSupport Database Component

If you installed the Mobile Support package on Integration Server, create the MobileSupport database component. Mobile Support writes data used for synchronizing mobile solutions, and metadata about that data, to this database component.

My webMethods Server

See " My webMethods Server Data Storage" on page 98.

OneDataMetadata, OneDataWorkArea, OneDataReleaseArea Database Components

If you installed OneData, create the OneDataMetadata, OneDataWorkArea, and OneDataReleaseArea database components.

■ OneData writes internal configuration data to the OneDataMetadata database component.

■ OneData users create data objects and work-in-progress data values in the OneDataWorkArea database component.

■ The OneDataReleaseArea database component contains the same data objects as the OneDataWorkArea database component, and approved data values. Data values are deployed from the OneDataReleaseArea database component.

You must create each OneData database component in a separate schema (Oracle) or database (DB2, MySQL, or SQL Server), and use a different database user for each database component. You must create the OneDataMetadata database component first, then the OneDataWorkArea database component, then the OneDataReleaseArea database component.

Optimize

If you installed Optimize, create the database components listed below.

Analysis Database Component

Optimize Analytic Engines write computed analytics and process and monitoring data received from Infrastructure Data Collectors and Web Service Data Collectors to the Analysis database component. The Optimize user interface displays the data.

Create a single Analysis database component for all Optimize Analytic Engines to share. If you are going to use root cause analysis for business processes, install the Analysis and ProcessTracker database components in the same schema (Oracle) or database (DB2, MySQL, or SQL Server).

CentralConfiguration Database Component

The CentralConfiguration database component is automatically created when you create the MywebMethodsServer database component. The CentralConfiguration database

component stores the Optimize configuration information you enter in the Central Configurator interface in My webMethods.

The CentralConfiguration database component is not automatically *dropped* when you drop the MywebMethodsServer database component. If you want to drop the CentralConfiguration database component, you must do so manually.

ProcessAudit Database Component

See .

ProcessTracker Database Component

Optimize Analytic Engines write business and process status data received from processes that are not orchestrated by the Process Engine to the ProcessTracker database component. The Optimize user interface displays the data. Monitor reads process status data from this database and displays it in the Monitor interface in My webMethods.

Create a single ProcessTracker database component for all Optimize Analytic Engines to share. If you are going to use root cause analysis for business processes, install the ProcessTracker and Analysis database components in the same schema (Oracle) or database (DB2, MySQL, or SQL Server).

DataPurge and DatabaseManagement Database Components

The DataPurge and DatabaseManagement database components are automatically created when you create the Analysis or ProcessTracker database component. The DataPurge database component provides a common method for purging data from the Analysis and ProcessTracker database components. The DatabaseManagement database component provides core routines for data purge functionality.

OperationManagement Database Component

The OperationManagement database component is automatically created when you create the Analysis, ProcessTracker, or DatabaseManagement database component. It provides a common method for configuration, logging, and error handling for those database components.

Trading Networks

TradingNetworks Database Component

If you installed Trading Networks Server, create the TradingNetworks database component.

Trading Networks Server writes metadata (partner profiles, trading partner agreements, document types, processing rules, and partner profile groups) and run-time data (documents, document content parts, attributes, and delivery information) to the TradingNetworks database component. Trading Networks Server also logs entries about document processing steps. The Trading Networks user interface displays the data.

If you are clustering Trading Networks Servers, create a single TradingNetworks database component for the cluster to share. If you are not clustering, create a TradingNetworks database component for each Trading Networks Server.

TradingNetworksArchive Database Component

If you want to archive Trading Networks data, also create the TradingNetworksArchive database component. Create the TradingNetworks and the TradingNetworksArchive database components in the same schema (Oracle) or database (DB2, MySQL, or SQL Server).

> **Note:** If you are using Oracle, you might be able to use partitioning for the TradingNetworksArchive database component. Contact Software AG Global Consulting Services for information.

Storage and ComponentTracker Database Component

The Storage database component creates default storage structures, such as tablespaces and user credentials, for all database components. The ComponentTracker database component is automatically created when you create the Storage database component, and tracks the database components that are created or dropped.

Preparing for Database Component Creation

Database User

You can create one database user for multiple database components, or you can create a different database user for each database component. Later sections in this chapter explain how to create database users.

The database user must have the permissions listed in the file below.

RDBMS	Path to File
DB2	*Software AG_directory*\common\db\scripts\db2\storage\30\create\db2_str_c_ddl_user.sql
MySQL Community Edition	*Software AG_directory*\common\db\scripts\mysql\storage\30\create\mys_str_c_ddl_user.sql
MySQL Enterprise Edition	*Software AG_directory*\common\db\scripts\mysql\storage\30\create\mys_str_c_ddl_user.sql

RDBMS	Path to File
Oracle	*Software AG_directory*\common\db\scripts\oracle\storage \25\create\ora_str_c_ddl_user.sql
	Software AG scripts require ALTER SESSION permission at the time of creating database tables. You can revoke this permission after the tables have been created.
SQL Server	*Software AG_directory*\common\db\scripts\mssql\storage \30\create\mss_str_c_ddl_user.sql

Database Storage

You can create all database component objects in the same database storage or you can create the objects for each database component in its own database storage. Use the Database Component Configurator and the instructions in this chapter to create database storage.

> **Important:** Always use the Database Component Configurator to create database storage and tablespaces. Do not create them manually.

By default, all database components follow the database storage convention below. You can use a different convention when you create database storage.

RDBMS	Data Storage Unit	Index Storage Unit	BLOB Storage Unit
DB2 tablespaces	WEBMDATA	WEBMINDX	WEBMBLOB
MySQL tablespace	Default	Default	Default
Oracle tablespaces	WEBMDATA	WEBMINDX	WEBMDATA
SQL Server filegroups	Primary	Primary	Primary

Database Administrator Account

You need DBA access only if you want to use the Database Component Configurator to create the necessary database users and storage objects.

Your products are designed to use the database users you create for the database components.

Character Set and Sort Order

Your products are globalized and support Unicode. Software AG strongly recommends choosing a Unicode encoding for your database and the most appropriate sort order for your environment. A database character set determines which languages a database can represent. Database sort order determines collation and comparison behavior.

The sections below list the most appropriate Unicode character encoding and sort order for each RDBMS that your products support. If you want to use a different character set or sort order than recommended below, consult your database administrator and your RDBMS vendor's documentation so you can carefully choose a database character set that supports the languages your data is in.

If you use the Database Component Configurator to create your database components, you can check whether the selected RDBMS is configured for the Unicode character set. If the RDBMS does not support Unicode, the configurator lists the character set the RDBMS does support.

> **Important:** You must set character set and sort order before creating storage.

DB2

Database schemas for DB2 use character data types. DB2 supports UTF-8 for character data types and UTF-16 for graphic data types. Software AG recommends these character sets and sort order:

For...	Software AGrecommends...
Character set	CCSID 1208 (UTF-8)
	My webMethods Server requires this character set.
Graphic Character Set	UTF-16
Sort order	IDENTITY_16BIT
	This sort order ensures the same sorting result for both character and graphic data types.

You can check the database configuration using the GET DATABASE CONFIGURATION command.

MySQL Community Edition and Enterprise Edition

The server character set and collation are used as default values if the database character set and collation are not specified in CREATE DATABASE statements. They have no other purpose.

You can determine the current server character set and collation settings from the values of the character set server and collation server system variables. You can change these variables at runtime.

For...	Software AGrecommends...
Character set	UTF-8
Collation	utf8_general_ci

You can check the database configuration using the SHOW VARIABLES LIKE command.

Oracle

Database schemas for Oracle use character data types. For character data types, Oracle supports the UTF8 and AL32UTF8 Unicode encodings. While UTF8 is CESU-8 compliant and supports the Unicode 3.0 UTF-8 Universal character set, AL32UTF8 conforms to the Unicode 3.1 or higher UTF-8 Universal character set. For nchar data types, Oracle supports the AL32UTF8 and AL16UTF16 Unicode encodings. The supported Unicode version for AL32UTF8 depends on the Oracle database version. Oracle database schemas for your products do not have linguistic indexes. Software AG recommends these character sets and sort order:

For...	Software AGrecommends...
Character set	AL32UTF8
Nchar character set	AL16UTF16
Sort order	Binary

You can check database configuration and session settings by viewing the SYS.NLS_DATABASE_PARAMETERS or V$NLS_PARAMETERS parameter.

SQL Server

Database schemas for SQL Server use nchar data types. SQL Server provides support for UTF-16 through its nchar data types. Since nchar data types are always in UTF-16, you do not have to perform any special database configuration and can choose the most

appropriate code page for your environment as a database character set. Software AG recommends these character sets and sort order:

For...	Software AGrecommends...
Character set	The appropriate encoding for the languages your data is in.
Nchar character set	UTF-16
Sort order	Any case-insensitive collation type. If you do not choose a case-insensitive sort order, you will not be able to create some database components in SQL Server.

You can check the database configuration using the sp_helpdb *database* stored procedure.

Page and Block Size

Use the page and block sizes specified below for each type of RDBMS.

RDBMS	Required Page and Block Size	Default
DB2	32K page size	4K
MySQL	16K page size	16K
Oracle	8K page/block size	8K
SQL Server	8K page/block size	8K

Set Database Options

MySQL

Set the parameter default_time_zone='+00:00' in the MySQL configuration file for the [client] and [mysqld] sections.

Oracle

For your products to function properly, you must set the NLS_LENGTH_SEMANTICS initialization parameter to BYTE.

> **Important:** Your products use BYTE rather than CHAR semantics because BYTE is the default configuration of Oracle database and is used by most Oracle users. If you try to use your products with CHAR semantics, unexpected results might occur. For example, since CHAR semantics allow storage of longer data than BYTE semantics (for Japanese, varchar(12 byte) takes four characters in UTF8 while varchar(12 char) takes 12), using CHAR semantics could cause the buffer in some products to overflow.

SQL Server

The following database settings are required for your products to function properly:

```
ALTER DATABASE database_name SET ALLOW_SNAPSHOT_ISOLATION ON;
ALTER DATABASE database_name SET READ_COMMITTED_SNAPSHOT ON;
```

In addition, enable Named Pipes and TCP/IP protocols using the SQL Server Configuration Manager.

By default, SQL Server uses case-insensitive collations. If you create the My webMethods Server database component, do not change this option from the default; My webMethods Server does not support case-sensitive collations.

Install the Database Component Configurator and Database Scripts

The Database Component Configurator and database scripts are available through the Software AG Installer. For complete instructions on using the installer, see *Using the Software AG Installer*.

1. Download the Software AG Installer as instructed by your installation email from Software AG to the appropriate machine, as follows:

If you are going to create database components...	Download the Software AG Installerto...
Automatically, using the Database Component Configurator	Any machine on the same network as your database server
Manually, using database scripts	A machine equipped with the database client for your type of RDBMS

2. Start the installer.

3. In the product selection list, select **Database Component Configurator**.

4. The Database Component Configurator has the default environment settings shown below. If necessary, you can change them as described below.

    a. Go to the *Software AG_directory* \common\db\bin directory.

    b. Open the setEnv.{bat|sh} file in the text editor and edit the fields as necessary.

Setting	Determines whether the configurator...	Default
-DtermOutOn	Writes execution information to the console.	true
-DspoolOn	Logs execution information.	true
-DlogLevel	Sets the log level to INFO (high-level information) or DEBUG (more detailed information, including internal processing and SQL calls).	INFO
-Dlog.dir	This setting specifies the full path to the directory in which to store the log files. Make sure you have permission to write to this directory.	*Software AG_directory*\common\db\logs

    c. Save and close the file.

Use the Database Component Configurator Graphical User Interface

This section explains how to use the Database Component Configurator graphical user interface (GUI) to create or drop a database user and database storage, and to create or drop database components.

> **Important:** Always use the Database Component Configurator to create database storage and tablespaces. Do not create them manually.

You can use the Database Component Configurator GUI on all operating systems that are supported by your products. UNIX systems must have an X Windows environment. Each run of the configurator operates on a single schema.

The configurator writes execution information to the **Results** tab and to a log file named dcc.log in the *Software AG_directory*\common\db\logs directory. The maximum size of

the file is 10MB. When that size is exceeded, the configurator creates a new file named dcc.*number* .log. If the number of files exceeds 1000, the configurator deletes the oldest file.

Create Database Components, Database User, and Storage

Start the Database Component Configurator GUI

Start the Database Component Configurator GUI as follows:

System	Action
Windows	On the **Start** menu, go to **Programs > Software AG> Tools > Database Component Configurator**.
UNIX	Go to *Software AG_directory* and run the command dbConfigurator.sh.
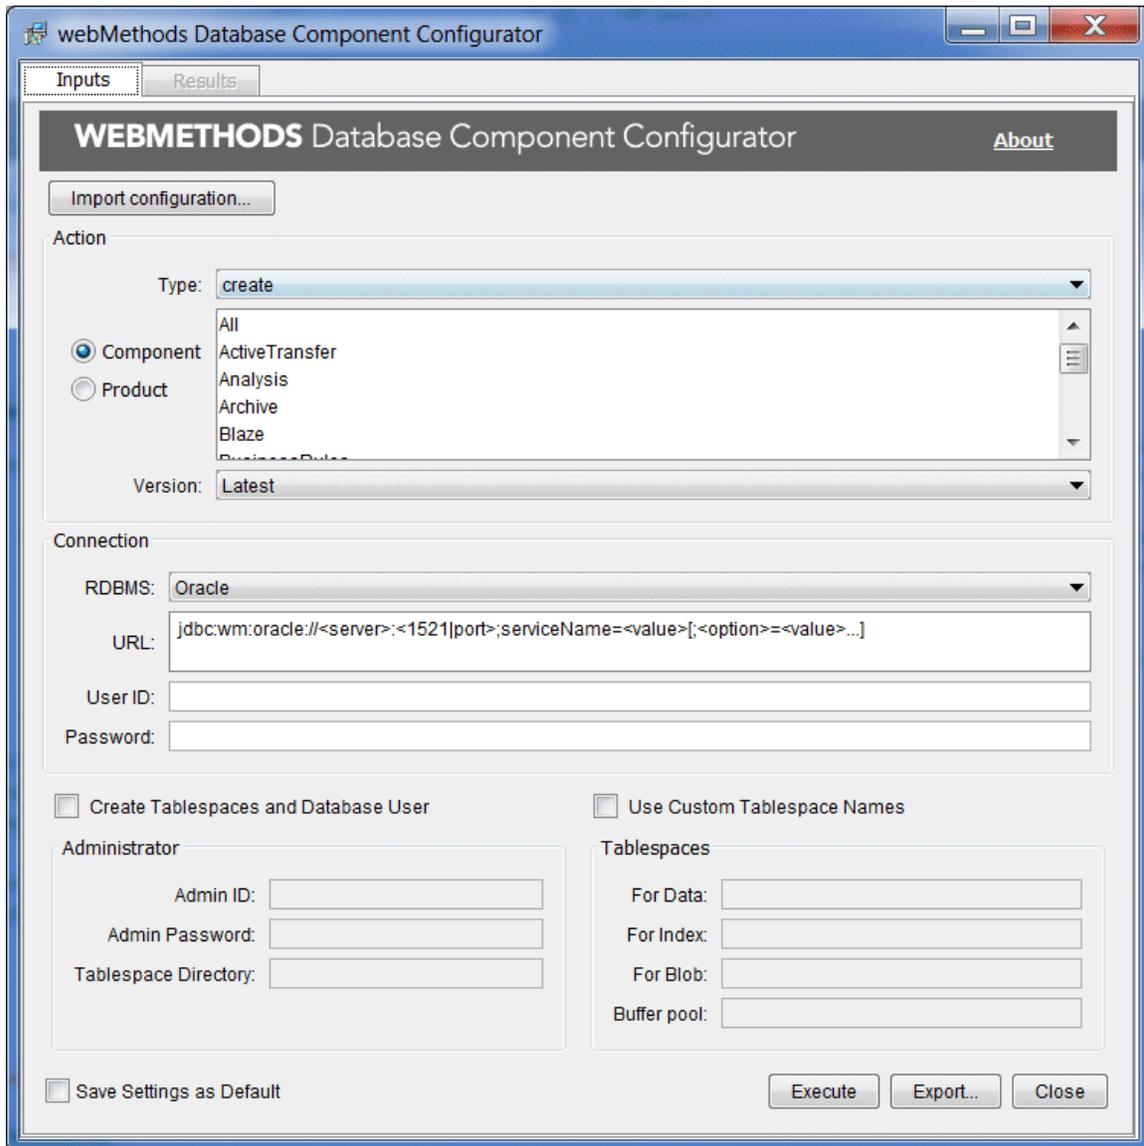
Initially, some of the fields on the **Inputs** tab show values you entered for the same fields in the Software AG Installer during product installation. You can change these values and specify other field values.

Choose the Action to Perform

In the **Action** area, in the **Type** list, click the action to perform, as follows:

Option	Action
create	Creates the database components you select in the **Action** area, and lets you create a database user and storage.

Option	Action
recreate	Drops and then re-creates the database components you select in the **Action** area.
catalog	List existing database components on the **Results** tab.
checkEncoding	Checks whether your RDBMS is Unicode-enabled and displays the answer on the **Results** tab.

Specify the Connection to the RDBMS

Specify the connection for the configurator to use to connect to the RDBMS.

1. In the **RDBMS** list, click the RDBMS in which to create the database components.

2. In the **URL** field, provide the URL for the RDBMS. Sample URL formats for the DataDirect Connect JDBC 5.1 driver are displayed. Below is additional information for completing this field.

   - For DB2, if you are going to create database components in a schema other than the default schema for the specified database user, you must specify these connection options in the URL, where `AlternateID` is the name of the default schema used to qualify unqualified database objects in dynamically prepared SQL statements:

     ```
     ;AlternateId=schema;"InitializationString=(SET CURRENT
     PATH=current_path,schema)"
     ```

     > **Important:** If you are creating Optimize database components (that is, Analysis and ProcessTracker), you must specify *schema* using all uppercase letters. In addition, you must specify the options `CreateDefaultPackage=true,ReplacePackage=true`, and `DynamicSections=3000`. These settings will affect all database components in the same schema or database.

   - For Oracle, if you are going to create storage and the Data Purge database component, you must specify the sysLoginRole connection option on the URL (for example, `;sysLoginRole=sysdba`). You must also specify the admin ID (for example, SYS), which belongs to the sysdba role.

   - For information about options supported by the DataDirect Connect JDBC 5.1 driver used by your products, including options that support clustering, data encryption (SSL), and different authentication methods, see the DataDirect Connect documentation, available on the Software AG Documentation website.

3. In the **User ID** and **Password** fields, your entries depend on the task you are going to perform, as described below.

> **Note:** You can choose to specify this authentication information in the URL, using DataDirect options, rather than in these fields.

If you are...	Specify...
Creating a database user and storage in DB2	The OS user to which to grant permissions, and the password for that user.
Creating a database user and storage in Oracle or SQL Server	The database user and password to create. For SQL Server, the user will be created and a default schema named dbo will be assigned to that user. For Oracle, do not use the SYSTEM user to create the database components in the SYSTEM schema.
Creating database components or performing the **catalog** action	An existing database user and password.
Performing the **checkEncoding** action	An existing database user that has create session and create table privileges, and that database user's password.

Create Database Components, Database User, and Storage

If you selected **create** in the **Action Type** list, do the following In the **Action** area:

1. Select the database components to create. You can use any of the options below.

   ■ Click **Component** and select one or more database components to create. You can select **All** to create all database components.

   ■ Click **Product** and select one or more products whose database components to create. You can see the list of database components for a product by clicking **Product**, clicking **print** in the **Action Type** list, clicking **Execute**, and then clicking the Results tab.

2. In the **Version** list, select **Latest**.

3. If you need to also create the database user and storage, follow the steps below.

   a. Select the check box labeled as follows:

RDBMS	Label
DB2	**Create Tablespaces and Grant Permissions to OS User**
MySQL	**Create Database and Database User**
Oracle	**Create Tablespaces and Database User**
SQL Server	**Create Database and Database User**

For DB2 on Linux systems, tablespaces are created for each DB2 database. If you are creating database components in more than one DB2 database, either the tablespace directory or the tablespace names must be unique for each DB2 database.

b. In the **Admin ID** field, identify the database user or operating system user that has database administrator credentials to create the database user and storage. Supply the password for the user in the **Admin Password** field.

c. The next field and your entry depend on your RDBMS.

RDBMS	Field and Entry
DB2 or Oracle	In the **Tablespace Directory** field, identify the directory in which to create the tablespaces.
MySQL or SQL Server	In the **Database** field, specify the database to create.

4. For DB2 or Oracle, you can select the **Use Custom Tablespace Names** check box and specify custom tablespace names in the fields.

   ▪ For DB2, the custom tablespace names will replace the defaults WEBMDATA, WEBMINDX, and WEBMBLOB. You can also specify a custom name to use for the buffer pool for your products (WEBMBUFF by default).

   > **Note:** Your products support all tablespace configurations deployed by users.

   ▪ For Oracle, the custom tablespace names will replace the defaults WEBMDATA and WEBMINDX.

Execute the Specified Action

1. If you intend to run the configurator more than once, you can set the current field values as the defaults for subsequent runs by clicking **Save Settings as Default**. You can also export field values to .xml files by clicking **Export**, then later import the

values from a file by clicking **Import Configuration**. In each case, the values for the two **Password** fields are not saved.

2. Click **Execute**. The execution information is displayed on the **Results** tab and is written to the log file dcc_*yyyymmddHHMMss* in the *Software AG_directory* \common \db\logs directory.

Drop Database Components

1. Shut down all products that are connected to the database components you want to drop, and back up the database components.

2. Start the Database Component Configurator GUI as follows:

System	Action
Windows	On the **Start** menu, go to **Programs > Software AG> Tools > Database Component Configurator**.
UNIX	Go to *Software AG_directory* and run the command `dbConfigurator.sh`.

3. In the **Action** area, in the **Type** list, click the action to perform, as follows:

Value	Action
drop	Drops the database components you select in the **Action** area.
catalog	Lists existing database components on the **Results** tab.

4. In the **Connection** area, specify the connection for the configurator to use to connect to the RDBMS, as follows:

Field	Do this...
RDBMS	Click the RDBMS from which to drop the database components.
URL	Type the URL for the RDBMS. Sample URL formats are displayed.
User ID and **Password**	Specify the database user and password to use to connect to the RDBMS.

5. In the **Action** area, select the database components to drop, or the products whose database components to drop.

6. In the **Action** area, select from the **Version** list as follows:

If you selected...	Click...
One or more database components	**Latest**. The configurator will create the latest version of the database component or of all database components, respectively.
One or more products	**Latest**

7. Click **Execute**. The execution information is displayed on the **Results** tab and is written to the log file dcc_*yyyymmddHHMMss* in the *Software AG_directory* \common \db\logs directory.

Drop the Storage and Revoke Database User Permissions

If you have dropped all database components in a schema, you can drop the storage and revoke the database user permissions.

1. Start the Database Component Configurator GUI as follows:

System	Action
Windows	On the **Start** menu, go to **Programs > Software AG> Tools > Database Component Configurator**.
UNIX	Go to *Software AG_directory* and run the command dbConfigurator.sh.

2. In the **Action Type** list, click **drop**.

3. In the **Action** area, click **Component** and select **All**.

4. In the **Connection** area, specify the connection for the configurator to use to connect to the RDBMS, as follows:

Field	Do this...
RDBMS	Click the RDBMS from which to drop the database components.
URL	Type the URL for the RDBMS. Sample URL formats are displayed.
User ID and **Password**	Specify the database user and password to use to connect to the RDBMS.

5. Select the **Drop tablespaces and database user** check box.

a. In the **Admin ID** field, identify the database user or operating system user that has database administrator credentials to drop the database user and storage. Supply the password for the database user in the **Admin password** field.

b. The next field depends on your RDBMS.

RDBMS	Fields and Entries
DB2 or Oracle	In the **Tablespace directory** field, identify the directory that contains the tablespaces to drop. If you specified custom tablespace names, supply those names in the **Tablespaces** area.
MySQL or SQL Server	In the **Database** field, specify the database that contains the storage.

6. Click **Execute**. The execution information is displayed on the **Results** tab and is written to the log file dcc_*yyyymmddHHMMss* in the *Software AG_directory* \common \db\logs directory.

Use the Database Component Configurator Command Line Interface

This section explains how to use the Database Component Configurator command line interface to create or drop a database user and storage, and to create or drop database components. You must create database storage before you create database components.

> **Important:** Always use the Database Component Configurator to create database storage and tablespaces. Do not create them manually.

You can use the Database Component Configurator command line interface on all operating systems that are supported by your products. Each run of the configurator operates on a single schema.

Before dropping any database component, shut down all products that are connected to the database component.

If you have dropped all the database components from a schema, you can drop the storage and revoke the database user permissions.

Database Component Configurator Command

Run the command below from the *Software AG_directory*\common\db\bin directory.

Main Parameters

```
dbConfigurator.{bat|sh} {-a|--action} action
{-d|--dbms}{oracle|sqlserver|db2luw}
{-l|--url} RDBMS_URL
```

```
{{-c|--component} db_component [{-v|--version} db_component_version] |
 {-pr|--product} product {-v|--version} product_release }
{-u|--user} db_user {-p|--password} password
[{-au|--admin_user} db_admin_user {-ap|--admin_password} password]
```

Additional Parameters

```
[{-tsdata|--tablespacefordata} data_tspace_name]
[{-tsindex|--tablespaceforindex} index_tspace_name]
[{-tsblob|--tablespaceforblob} BLOB_tspace_name]
[{-b|--bufferpool} buffer_pool_name]
[-t|--tablespacedir directory]
[-n|--dbname database_name]
[{-e|--export|-i|--import} {-dir|--configdir} directory
 {-file|--configfile} file_name]
[-r|--runCatalog]
[-h|--help]
[-pa|--printActions] [-pc|--printComponents] [-pp|--printProducts]
```

Main Parameters

{-a | --action} *action*

Value	Action
catalog	Lists existing database components.
create	Creates a database user and database storage, or database components.
drop	Drops a database user and database storage, or database components.
recreate	Performs the drop action and then the create action.
checkEncoding	Indicates whether your RDBMS is Unicode-enabled.

{-d | --dbms} {db2luw|mysql|oracle|sqlserver|}

Type of RDBMS on which to perform *action*.

{-l | --url} *RDBMS_URL*

URL for the RDBMS. When working with database components on UNIX systems, you must enclose the URL in double quotes. For information about options supported by the DataDirect Connect JDBC 5.1 driver used by your products, see the DataDirect Connect documentation, available on the Software AG Documentation website.

RDBMS	Sample URL Format
DB2	jdbc:wm:db2://*server*:{50000|*port*};databaseName=*database* [;*option*=*value*]...

RDBMS	Sample URL Format

	For DB2, if you are going to create database components in a schema other than the default schema for the specified database user, you must specify these connection options in the URL:
	
	`;AlternateId=schema`
	`;"InitializationString=(SET`
	`CURRENT PATH=current_path,schema)"`
	
	`AlternateID` is the name of the default schema used to qualify unqualified database objects in dynamically prepared SQL statements.
	
	If you are creating Optimize database components (that is, Analysis and ProcessTracker), you must specify *schema* using all uppercase letters. In addition, you must specify the options `CreateDefaultPackage=true`, `ReplacePackage=true`, and `DynamicSections=3000`. These settings will affect all database components in the same schema or database.

MySQL Community Edition	`jdbc:mysql://server:{3306|port}/databaseName`
MySQL Enterprise Edition	`jdbc:wm:mysql://server:{3306|port};databaseName=database` `[;option=value]...`

Oracle	`jdbc:wm:oracle://server:{1521|port};serviceName=service`
	`[;option=value]...`
	
	If you are creating storage and the Data Purge database component, you must specify the sysLoginRole connection option on the URL (for example, `;sysLoginRole=sysdba`).

SQL Server	`jdbc:wm:sqlserver://server:{1433|port};databaseName=database` `[;option=value...`

{{-c | --component} *db_component* [{-v | --version} *db_component_version*] | {-pr | --product} *product* {-v | --version} *product_release* }
One of the following:

- Database user and storage.

- Database components and version. Specify one or more database component codes or names, using commas to separate multiple codes or names, and the version `latest`. `--printComponents` lists database component codes and names.

- Product database components and product release. Specify one or more product codes or names, or `ALL`, using commas to separate multiple codes or names, and

the release `latest.` `--printProducts` lists database components for specific products.

{-u | --user} *db_user* **{-p | - -password}** *password*

If you are...	Specify these values...
Creating a database user and storage in Oracle or SQL Server	Database user and password to create. For SQL Server, the user will be created and a default schema named dbo will be assigned to that user. For Oracle, do not use the SYSTEM user to create the database components in the SYSTEM schema.
Creating a database user and storage in DB2	OS user to which to grant permissions.
Performing the checkEncoding action	Existing database user that has create session and create table privileges, and that database user's password.
Performing any other action	Existing database user and password.

> **Note:** You can choose to specify this authentication information in the URL, using DataDirect options, instead of in these parameters.

[{-au | --admin_user} *db_admin_user* **{-ap | --admin_password}** *password* **]**
If you are going to create a database user and storage, or drop storage and revoke the database user permissions, specify the database user or operating system user and password that has the necessary database administrator credentials.

Additional Parameters

[{-tsdata | --tablespacefordata} *data_tspace_name* **] [{-tsindex | -- tablespaceforindex}** *index_tspace_name* **] [{-tsblob | -- tablespaceforblob}** *BLOB_tspace_name* **] [{-b | --bufferpool}** *buffer_pool_name* **]**
If you are going to create database components in DB2 or Oracle, you can specify custom tablespace names. Your products support all tablespace configurations deployed by users. For Oracle, the custom tablespace names will replace the defaults WEBMDATA and WEBMINDX. For DB2, the custom tablespace names will replace the defaults WEBMDATA, WEBMINDX, and WEBMBLOB. You can also specify a custom name to use for the buffer pool (WEBMBUFF by default). For DB2 or Oracle, if you are going to drop storage and revoke the database user permissions, provide the custom tablespace names.

[{-t | --tablespacedir} *directory* **]**

If you are going to create a database user and storage, and want to create a tablespace directory for DB2 or Oracle, full path to the directory. For DB2 on Linux systems, tablespaces are created for each DB2 database. If you are creating database components in more than one DB2 database, either the tablespace directory or the tablespace names must be unique for each DB2 database. If you are going to drop storage and revoke the database user permissions for Oracle or DB2, identify the directory that contains the tablespaces to drop.

[{-n | --dbname} *database_name* **]**

If you are going to create a database user and storage in SQL Server, name of the target database. If you are going to drop storage and revoke the database user permissions in SQL Server, identify the database that contains the storage.

[{-e | --export} {-dir | --configdir} *directory* **{-file | --configfile}** *file_name* **]**

Exports values for all required parameters except the password parameters to an .xml file.

[{-i | --import} {-dir | --configdir} *directory* **{-file | --configfile}** *file_name* **]**

Imports values for all parameters from an .xml file. Password parameters are not saved, so you must specify them on the command.

[{-r | --runCatalog}]

Runs the catalog action at the end of every execution.

[{-h | --help}]

Lists command line interface help.

[{-pa | --printActions}]

Lists the actions you can perform, like create or drop.

[{-pc | --printComponents}]

Lists all database component codes and names.

[{-pp | --printProducts}]

Lists database component components for specific products.

Examples

> **Important:** On UNIX systems, you must enclose the RDBMS URL in your commands in double quotes.

Create the Database User and Storage

This command creates the database user and the two tablespaces (default names) for Oracle:

```
dbConfigurator.bat -a create -d oracle -c storage -v latest
-l jdbc:wm:oracle://DBserver:1521;serviceName=myservicename -au adminuser
-ap adminpass -u webmuser -p w3bmpass -t tablespace_dir
```

This command creates the database user and the two tablespaces (custom names) for Oracle:

```
dbConfigurator.bat -a create -d oracle -c storage -v latest
-l jdbc:wm:oracle://DBserver:1521;serviceName=myservicename -au adminuser
-ap adminpass -u webmuser -p w3bmpass -t tablespace_dir -tsdata MYDATA
-tsindex MYINDEX
```

This command creates the database user and storage for SQL Server:

```
dbConfigurator.bat -a create -d sqlserver -c storage -v latest
-l jdbc:wm:sqlserver://DBserver:1433;databaseName=master -u webmuser -p w3bmpass
-au sa -ap sa_password -n webmdb
```

This command creates the buffer pool and tablespaces (default names), and grants permissions for DB2:

```
dbConfigurator.bat -a create -d db2luw -c storage -v latest
-l jdbc:wm:db2://vmxpdb01:50000;databaseName=amol -u webmuser -au adminuser
-ap admin_password -t tablespace_dir
```

This command creates the buffer pool and tablespaces (custom names), and grants permissions for DB2:

```
dbConfigurator.bat -a create -d db2luw -c storage -v latest
-l jdbc:wm:db2://vmxpdb01:50000;databaseName=amol -u webmuser -au adminuser
-ap admin_password -t tablespace_dir -tsdata MYDATA -tsindex MYINDEX
-tsblob MYBLOB -b MYBUFF
```

Drop the Storage and Revoke Database User Permissions

If you drop all database components in a schema, you can drop the storage and revoke the database user permissions.

This command drops the two tablespaces (default names) for Oracle:

```
dbConfigurator.bat -a drop -d oracle -c storage -v latest
-l jdbc:wm:oracle://DBserver:1521;serviceName=myservicename -au adminuser
-ap adminpass -u webmuser -p w3bmpass -t tablespace_dir
```

This command drops the two tablespaces (custom names) for Oracle:

```
dbConfigurator.bat -a drop -d oracle -c storage -v latest
-l jdbc:wm:oracle://DBserver:1521;serviceName=myservicename -au adminuser
-ap adminpass -u webmuser -p w3bmpass -t tablespace_dir -tsdata MYDATA
-tsindex MYINDEX
```

This command drops the database for SQL Server:

```
dbConfigurator.bat -a drop -d sqlserver -c storage -v latest
-l jdbc:wm:sqlserver://DBserver:1433;databaseName=master -u webmuser
-p w3bmpass -au sa -ap sa_password -n webmdb
```

This command drops the buffer pool and tablespaces (default names), and revokes permissions for DB2:

```
dbConfigurator.bat -a drop -d db2luw -c storage -v latest
-l jdbc:wm:db2://vmxpdb01:50000;databaseName=amol -u webmuser -au adminuser
-ap admin_password -t tablespace_dir
```

This command drops the buffer pool and tablespaces (custom names), and revokes permissions for DB2:

```
dbConfigurator.bat -a drop -d db2luw -c storage -v latest
```

```
-l jdbc:wm:db2://vmxpdb01:50000;databaseName=amol -u webmuser -au adminuser
-ap admin_password -t tablespace_dir -tsdata MYDATA -tsindex MYINDEX
-tsblob MYBLOB -b MYBUFF
```

Create Database Components

This command creates the latest version of the ProcessAudit and ProcessEngine database components in Oracle:

```
dbConfigurator.bat -a create -d oracle -c processaudit,processengine -v latest
-l jdbc:wm:oracle://DBserver:1521;serviceName=myservicename
-u webmuser -p w3bmpass
```

This command recreates (drops and then creates) the latest versions of all database components in SQL Server:

```
dbConfigurator.bat -a recreate -d sqlserver -c all -v latest
-l jdbc:wm:sqlserver://DBserver:1433;databaseName=webmdb -u webmuser -p w3bmpass
```

This command displays the database components that currently exist in DB2:

```
dbConfigurator.bat -a catalog -d db2luw
-l jdbc:wm:db2://DBserver:50000;databaseName=webmdb -u webmuser -p w3bmpass
```

This command creates the latest database components in DB2 when creating in ALTSCHEMA rather than the default schema for the specified database user:

```
dbConfigurator.bat -a create -d db2luw -c all -v latest
-l "jdbc:wm:db2://DBserver:50000;databaseName=webmdb
;AlternateId=ALTSCHEMA;InitializationString=\SET CURRENT
PATH=current_path,ALTSCHEMA\" -u webmuser -p w3bmpass
```

Drop Database Components

Before dropping, shut down all products that are connected to the database component and back up the database component.

This command drops the latest ProcessAudit and ProcessEngine database components from Oracle:

```
dbConfigurator -a drop -d oracle -c processaudit,processengine -v latest
-l jdbc:wm:oracle://myserver:1521;serviceName=myservicename -u webmuser
-p w3bmpass
```

This command drops the latest versions of all database components from SQL Server:

```
dbConfigurator.bat -a drop -d sqlserver -c core -v latest
-l jdbc:wm:sqlserver://myserver:1433;databaseName=webmdb -u webmuser -p w3bmpass
```

The command drops the latest version of the Analysis database component from DB2:

```
dbConfigurator.bat -a drop -d db2luw -c analysis -v latest
-l jdbc:wm:db2://DBserver:50000;databaseName=webmdb -u webmuser -p w3bmpass
```

This command drops the latest version of the Analysis database component from DB2 when dropping from ALTSCHEMA rather than the default schema for the specified database user:

```
dbConfigurator.bat -a drop -d db2luw -c analysis -v latest
-l jdbc:wm:db2://DBserver:50000;databaseName=webmdb;AlternateId=ALTSCHEMA
;"InitializationString=(SET CURRENT PATH=current_path,ALTSCHEMA)" -u webmuser
-p w3bmpass
```

Connect Products to Database Components

To enable products to interact with a database component, you provide database connection parameters, usually during product installation, and the Software AG Installer uses the parameters to configure JDBC connection pools. After product installation and database component creation is complete, you set up the products to use the appropriate connection pools.

This section explains how to define the database connections, if you did not define them during installation, or if you want to modify the database connections you set during installation.

■ From Integration Server or products it hosts to the ActiveTransfer, APIGatewayEvents, Archive, BusinessRules, CloudStreamsEvents, CrossReference, DocumentHistory, ISInternal, ISCoreAudit, MediatorEvents, MobileSupport, ProcessAudit, ProcessEngine, Staging and Reporting, and TradingNetworks and TradingNetworksArchive database components.

■ From Integration Server to the central user management feature in My webMethods Server. This feature allows My webMethods Server users to access Integration Server, enables Integration Server to delegate user management to My webMethods Server, and enables My webMethods Server's user interfaces to call Integration Server services.

For instructions on defining the database connections for other products, if you did not define them during installation, or if you want to modify the database connections you set during installation, see the following:

Product	Documentation
Software AG Designer	Software AG Designer online help
My webMethods Server	*Administering My webMethods Server*
OneData	*Administering webMethods OneData*
Optimize	*Configuring BAM*
Task Engine	*webMethods Task Engine User's Guide*

Database Connections for Integration Server and Hosted Products or Microservices Container

You use Integration Server Administrator or Microservices Container Administrator to connect Integration Server or Microservices Container to database components, as follows:

■ You define *JDBC database connection pools* in Integration Server Administrator or Microservices Container Administrator. Each pool defines the connection to a database server that hosts database components.

■ Integration Server Administrator and Microservices Container Administrator provide *functions* that correspond to most database components (for example, an ISCoreAudit Log function that corresponds to the ISCoreAudit database component, an Xref function that corresponds to the CrossReference database component, and so on). You direct each function to write to its database components by pointing the function at the appropriate connection pool.

During Integration Server or Microservices Container installation, if you chose an external RDBMS for the product database components and supplied the database connection parameters, the installer automatically did the following:

■ Configured Integration Server or Microservices Container to write to the external RDBMS.

■ Created a default connection pool from the database connection parameters you supplied. The pool is named based on the value you supplied in the Connection name field on the installer panel.

■ Pointed the Xref, DocumentHistory, ISInternal, ISCoreAudit, ProcessAudit, and ProcessEngine functions at that pool.

> **Note:** If you created the BusinessRules database component, you created it in the same schema or database as the ProcessAudit database component. Integration Server therefore writes to the BusinessRules database component using the ProcessAudit function. There is no separate function for the BusinessRules database component.

You can edit the default connection pool using the instructions in this section, or you can create separate pools and redirect individual functions to use them. At run time, Integration Server or Microservices Container creates a separate instance of the appropriate connection pool for each database component.

In the case of Integration Server, you also use Integration Server Administrator to connect other products to database components, as follows:

■ During Trading Networks Server installation, you supplied the database connection parameters for the TradingNetworks database component. The installer created a connection pool from those parameters and pointed the TN function at that

pool. You can edit the connection using Integration Server Administrator and the instructions below. The same is true for ActiveTransfer Server and Mobile Support.

■ If you created the CloudStreamsEvents database component, you must point that function at a connection pool. The same is true for the Archive, Staging, and Reporting database components.

■ You must point the Simulation function at the connection pool for the Reporting database component.

Integration Server or Microservices Container can authenticate clients using internally-defined user and group information, or can use the My webMethods Server central user management feature. To use the latter, and to use the single sign on feature for My webMethods Server, Integration Server or Microservices Container must have a connection pool that points to the My webMethods Server database component, and the CentralUsers function in Integration Server or Microservices Container must point at that connection pool.

If you installed My webMethods Server in the same directory as Integration Server or Microservices Container, the installer created a connection pool from the My webMethods Server database parameters you supplied and pointed the CentralUsers function at that pool. If you installed My webMethods Server in a different directory from Integration Server or Microservices Container, you must create the connection pool and point the CentralUsers function at that pool using Integration Server Administrator or Microservices Container Administrator and the instructions below.

Define an Integration Server or Microservices Container Connection Pool

1. Start Integration Server or Microservices Container, open Integration Server Administrator or Microservices Container Administrator, and go to the **Settings > JDBC Pools** page.

2. On the **Settings > JDBC Pools** page, click **Create a new Pool Alias Definition** and complete the fields as described in the Integration Server online help.

3. Make sure Integration Server or Microservices Container can connect to the database by clicking **Test Connection**. Then click **Save Settings**.

Point Integration Server or Microservices Container Functions at Connection Pools

1. In Integration Server Administrator or Microservices Container Administrator, go to the **Settings > JDBC Pools** page.

2. In the **Functional Alias Definitions** area, click **Edit** in the **Edit Association** column for a function. In the **Associated Pool Alias** list, click the pool you want the function to point to, and then click **Save Settings**.

3. In the **Functional Alias Definitions** area, initialize the pool by clicking **Restart** in the **Restart Function** column for the function. Make sure Integration Server or Microservices Container can connect to the database by clicking ▶ in the **Test** column for the function.

4. Repeat the previous steps for all functions you want to point at a connection pool.

5. If you created a connection pool for the My webMethods Server database component and pointed the CentralUsers function at it, go to the **Settings > Resources** page and make sure the **MWS SAML Resolver URL** field is pointing to your My webMethods Server host and port.

6. Restart Integration Server or Microservices Container.

Drop Redundant Database Component

If you used the Database Component Configurator to install database components, and you installed the Integration Server database components and the Optimize database components in two different schemas, you will have two instances of the Process Audit Log database component. Configure Integration Server and Optimize to point to the same instance using the instructions above for Integration Server and in *Configuring BAM* for Optimize. You can then drop the other, redundant instance.

A Register Daemons to Automatically Start and Shut Down Products on UNIX Systems

Overview

You can register daemons for Software AG products on UNIX systems to make those products start and stop automatically at system start and shutdown time. For many products, you can register daemons at installation time, as described earlier in this guide. After installation, you can use the command line tool daemon.sh to register daemons for more Software AG products. You can also use the tool to unregister daemons for Software AG product.

The daemon.sh script generates an init-script for each daemon. The naming convention for the init-script is `sag[number]instance_name`, where `sag` is a fixed prefix string, and *instance_name* is a unique identifier for the daemon instance. If you accidentally specify a non-unique *instance_name* during daemon registration, the daemon.sh script automatically adds *number* to make the init-script name unique. The table below shows the locations of the generated init-scripts.

System	Location
SLES 11, RHEL 6, Solaris	/etc/init.d
SLES 12, RHEL 7	/usr/lib/systemd/system
HP-UX	/sbin/init.d
Mac OSX	/Library/LaunchDaemons
AIX	/etc

Each product daemon has an rc-script that has these features:

- It is owned and called by the installation user (that is, the non-root UNIX user that ran the installer).

- It accepts the arguments 'start' and 'stop' to start and stop the corresponding product.

At system start and shutdown time, the init-script changes the current user ID from the root user to the user that owns the rc-script and then calls the rc-script with the start or stop option, as appropriate. The rc-script names and locations are listed in "Product Daemons and rc-scripts" on page 133.

> **Note:** Products not listed in this appendix have no daemons to register or have their own method for registering daemons that is covered in product documentation.

Run the daemon.sh Command Line Tool

The daemon.sh command line tool is located in the *Software AG_directory* /common/bin directory. Log on to your system as the root user and call it from any current working directory.

The table below lists the most commonly used command line options, which are the same for every supported UNIX platform. The rc-script names and locations are listed in .

Action	Command with Option
Register daemon with default init-script name	`daemon.sh -f` *`absolute_path_to_rc-script`*
Register daemon and specify instance name	`daemon.sh -f` *`absolute_path_to_rc-script`* `-n` *`instance_name`*
List all registered daemons that have prefix `sag`	`daemon.sh -L`
Unregister daemon	`daemon.sh -r -f` *`absolute_path_to_rc-script`*
Remove daemon registrations for which rc-script no longer exists (clean up)	`daemon.sh -C`
Print information about implementation and other command line options	`daemon.sh -h`

> **Note:** The execution of daemon.sh uses specific run-levels and methods for registering UNIX daemons that have been tested by Software AG. Some systems may offer valid alternative ways for the daemon registration that are not documented here.

Product Daemons and rc-scripts

The table below shows the rc-script for each product daemon. When using the `daemon.sh -f` option to register or unregister a daemon, make sure to use the absolute

path name (for example, `daemon.sh -f *Software AG_directory*/API_Portal/server/sagyap*release*.sh`).

Product	Daemon	rc-script (under *Software AG_directory*)
API Portal	APICloudAgent*release*	API_Portal/server/ sagyap*release_number*
webMethods Broker	webMethods Enterprise Broker subsystem	Broker/aw_broker96
CentraSite Registry Repository	CentraSite Registry/ Repository	CentraSite/bin/saginm
Content Service Platform	CSP Server (Software AG webMethods Content Service Platform)	CSP/wrapper/Unix/csp
EntireX Broker	EntireX Administration Service	EntireX/bin/sagexx
Integration Agent	Software AG Integration Agent	profiles/*instance_name* /bin/ sagintagent*release_number*
Integration Server, or Microservices Container	Software AG Integration Server	profiles/*instance_name* /bin/ sagis*release_number*
My webMethods Server	Software AG My webMethods Server	profiles/*instance_name* /bin/ sagmws*release_number* _ *instance_name _instance_number*
OneData	Software AG OneData Server	profiles/ODE/bin/ sagode*release_number*
OptimizeInfrastructure Data Collector	Software AG Optimize Infrastructure Data Collector	profiles/InfraDC/bin/sagwid *release_number*
Optimize Web Service	Software AG Optimize WS Data Collector	optimize/dataCollector/bin/ sagoptiDC .sh

Product	Daemon	rc-script (under *Software AG_directory*)
Data Collector		
Optimize Analytic Engine	webMethods Optimize Analysis Engine	optimize/analysis/bin/ sagoptiAE .sh
Platform Manager	Software AG Platform Manager	profiles/SPM/bin/ sagspm*release_number*
Software AG Runtime	Software AG Runtime	profiles/CTP/bin/ sagctp*release_number*
System Management Hub	Software AG Instance Manager	InstanceManager/bin/ cim*release_number*
Universal Messaging	nserverdaemon	UniversalMessaging/ server/*instance_name* /bin/nserverdaemon

B   International Operating Environments

Overview

Software AG products are designed for use anywhere in the world, and enable you to exchange data with systems and users who speak languages other than your own.

The products use the Universal Character Set (ISO/IEC 10646-2, or Unicode) as the character set for all internal operations. At minimum, the products support the requirements of Unicode 3.0. HTML content, including administration tools, uses Unicode UTF-8 encoding.

For specific information on the testing, configuration, capabilities, and limitations of any product, see the product's readme.

Language Packs

The Software AG Installer always installs the U.S. English files necessary to run your products. However, for many products, you can also install language packs that provide translated user interfaces, messages, help, and documentation.

Most products support multiple languages at a time; you can install as many language packs as you need. For products that support multiple languages, your client setup (that is, browser language preference, default locale in which the product is started, and locale selected in your user profile) will determine which language you see. If you operate a product in a locale for which language packs are not available or installed, you will see U.S. English.

Language packs are available for webMethods Broker on all supported operating systems. Language packs are available for other products on all supported operating systems except Mac OS X.

Software AG Designer Language Packs

By default, the language that Software AG Designer uses depends on your JRE default locale, which in turn depends on your operating system locale setting. If the operating system locale is not set to the language you want to use, you can override the setting for Software AG Designer by starting Software AG Designer with the -nl option. For example, open a command window and enter `eclipse.exe -nl ja_JP`.

Software AG Designer is based on the Eclipse platform and its projects. Software AG language packs do not include language packs for plug-ins provided by Eclipse projects. If you need language packs for those plug-ins, follow the steps below. For background information about the available translations, their completeness, and more, go to the Eclipse web site.

1. Start Software AG Designer.

2. Go to **Help > Install New Software**, click the **Available Software** tab, and then click **Add...**

3. In the **Location** field, type the URL for Eclipse Neon-compatible language packs (for example, `http://download.eclipse.org/technology/babel/update-site/ R0.14.1/neon`) and click **OK**.

4. Click **Refresh**, select the language pack for the language you need, and click **Install**. Confirm the language pack to install and click **Next**.

5. Read the license agreement. If you accept the terms, select the check box and click **Finish**.

6. Restart Software AG Designer.

Extended Character Set

The Software AG Installer offers an extended character set that extends Java's java.nio.Charset class to support additional installable character encodings and adds more than 870 custom character encodings. The extended character set allows you to use custom character encodings such as EBCDIC and OEM character encodings with Integration Server, adapters, and services.

Configure Browsers and JRE Font

Some HTML interfaces for Software AG products detect the language setting for your browser. If you have installed language packs on your products, those interfaces try to supply the user interface, messages, and data in that language. You can indicate which of the installed languages to display in your browser as explained below. For supported browser releases, see *System Requirements for Software AG Products*.

Browser	Action
Firefox	Go to **Tools > Options > Content**, click **Choose**, and add the language you want to use to the top of the list.
Google Chrome	Go to **Settings > Show Advanced Settings > Languages**, click spell-checker settings, and add the language you want to use to the top of the list.
Internet Explorer	Go to **Tools > Internet Options**, click **Languages**, and add the language you want to use to the top of the list.
Safari	Go to **System Preferences > Language & Text > Language**, and add the language you want to use to the top of the list.

Your JRE might require modification to display non-ASCII data. If you see hollow boxes or black squares where you would expect to see non-ASCII characters, you might

need to modify your JRE's font configuration file to provide an appropriate font for displaying these characters. For example, you might install a full Unicode font, such as Arial Unicode MS, or you can install a font for a particular language. For information on modifying font configuration file, see J2SE at the Oracle Java SE documentation Web site.

Configure the Proper Locale

If you are going to install or run your products in a locale other than U.S. English, follow the instructions below.

System	Action
Windows	These instructions are for Windows 7; instructions for other Windows systems vary slightly. 1. On the **Start** menu, go to **Control Panel > Regional and Language**. 2. Click the **Keyboards and Languages** tab. If the language you want to use is not yet installed, click **Install/Uninstall Languages** and install the language. In the **Display Language** area, select the language and click **OK**. 3. Make the display locale setting take effect by logging off Windows and then logging in again.
UNIX	Set your shell to the proper locale by setting the LANG environment variable. Run the appropriate command below: ▪ In an sh shell: `LANG=locale; export LANG` ▪ In a csh shell: `setenv LANG locale`