

Installing Software AG Products

Version 10.3

October 2018

This document applies to webMethods Product Suite Version 10.3 and to all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 2010-2018 Software AG, Darmstadt, Germany and/or Software AG USA Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

The name Software AG and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA Inc. and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://softwareag.com/licenses>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

Use, reproduction, transfer, publication or disclosure is prohibited except as specifically provided for in your License Agreement with Software AG.

Table of Contents

About this Guide.....	7
Document Conventions.....	7
Online Information and Support.....	8
Data Protection.....	9
Using Command Central or Software AG Installer to Install Products.....	11
Using Command Central to Install Products.....	12
Using Software AG Installer to Install Products.....	12
Hardware Requirements and Operating System, Browser, and RDBMS Support.....	13
Hardware Requirements.....	14
Minimum and Recommended Hardware Requirements.....	14
Additional Hardware Requirements.....	18
Operating System Support.....	20
Browser Support.....	20
RDBMS Support.....	20
Preparing for Product Installation.....	21
Shut Down Products.....	22
Prepare for Installation For All Products (Windows System).....	22
Prepare for Installation in a Microsoft Azure Virtual Machine.....	22
Prepare for Apama Installation (Linux System).....	22
Prepare for API Portal Installation (UNIX System).....	22
Prepare for CentraSite Installation (UNIX System).....	23
Prepare for ApplinX Installation.....	23
Prepare for EntireX Installation.....	24
Prepare for Internal Data Store Installation (Linux System).....	24
Prepare for MashZone NextGen Installation (Linux System).....	25
Prepare for Mobile Designer Installation.....	26
Prepare for My webMethods Server Installation (UNIX System).....	26
Prepare for OneData Installation (Windows System).....	26
Prepare for Optimize Analytic Engine Installation (UNIX System).....	26
Prepare for Zementis Predictive Analytics Installation.....	27
Prepare Database Connection Information.....	27
Preparing for Database Component Creation.....	29
Database Component Descriptions and Installation Requirements.....	30
ActiveTransfer Server Database Components.....	30
API Gateway Database Components.....	30
CloudStreams Database Components.....	30
Software AG Designer Database Components.....	31
Dynamic Business Orchestrator Database Components.....	31

Integration Server or Microservices Runtime Database Components.....	32
External RDBMS versus Embedded Database.....	33
Using the Embedded Database.....	33
Using the External RDBMS.....	34
Mobile Support Database Components.....	34
Monitor Database Components.....	34
My webMethods Server Database Components.....	35
OneData Database Components.....	35
Optimize Database Components.....	36
Process Engine Database Components.....	36
Rules Engine Database Components.....	37
Trading Networks Server Database Components.....	38
Database Drivers for DB2, MySQL Enterprise Edition, Oracle, and SQL Server.....	38
Prepare for Database Component Creation in MySQL.....	39
Install MySQL Community Edition Database Driver.....	39
Edit the MySQL Configuration File.....	39
Database User and Storage, and Database Settings.....	39
Database User.....	39
Database Storage.....	40
Database Administrator Account.....	41
Character Set and Sort Order.....	41
DB2.....	41
MySQL Community Edition and Enterprise Edition.....	42
Oracle.....	42
SQL Server.....	43
Page and Block Size.....	43
Set Database Options.....	44
MySQL.....	44
Oracle.....	44
SQL Server.....	44
Installing Products and Creating Database Components.....	45
Installing Products and Creating Database Components Using Command Central.....	46
Installing Products Using Software AG Installer and Creating Database Components Using Database Component Configurator.....	46
Installing Products Using Software AG Installer.....	46
Recommendations.....	46
Before Installing Using Software AG Installer.....	47
Gather Information.....	47
Save and Verify Product License Files.....	47
Read and Perform Required Actions in Using Software AG Installer.....	47
API Gateway Requirements.....	47
CentraSite Requirements.....	47
Deployer and Asset Build Environment Requirements.....	48
Install Products Using Software AG Installer.....	48

Start the Installer and Provide General Information.....	48
Choose Products to Install.....	49
Choose Language Packs to Install.....	50
Supply Product Configuration Information.....	50
Apama.....	50
API Portal.....	51
EntireX.....	51
Integration Server or Microservices Runtime.....	52
OneData.....	52
Universal Messaging.....	53
After Installing Using Software AG Installer.....	54
Manage Products Using Command Central.....	54
Install Latest Fixes.....	55
Start, Configure, and Customize Products.....	55
Register Daemons and Set File Permissions (UNIX System).....	55
Uninstall Products Using Software AG Installer.....	55
Uninstall EntireX.....	55
Uninstall All Other Products.....	56
Create Database Components Using Database Component Configurator.....	57
Install the Database Component Configurator and Database Scripts.....	57
Database Component Configurator Command.....	57
Main Parameters.....	58
Additional Parameters.....	61
Examples.....	62
Create the Database User and Storage.....	62
Drop the Storage and Revoke Database User Permissions.....	63
Create Database Components.....	63
Drop Database Components.....	64
Connecting Products to Database Components.....	65
Registering Daemons to Automatically Start and Shut Down Products on UNIX Systems.....	67
Overview.....	68
Run the daemon.sh Command Line Tool.....	69
Product Daemons and rc-scripts.....	70
Troubleshooting.....	71
Operating in International Operating Environments.....	73
Overview.....	74
Language Packs.....	74
Language Packs for Software AG Designer.....	74

Extended Character Set.....	75
Configure Browsers and JRE Font.....	75
Configure the Proper Locale.....	76
Running in a Cloud Infrastructure as a Service (IaaS) Environment.....	77
Overview.....	78
Amazon Web Services.....	78
Microsoft Azure.....	79

About this Guide

This guide provides pre-installation, installation, and uninstallation instructions for Software AG products.

This guide does not provide instructions on configuring products after installation. For those instructions, see the product documentation.

The information provided in this guide for Terracotta and Universal Messaging applies 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.

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.

The statement below applies to the Database Component Configurator:

With respect to the EU General Data Protection Regulation (GDPR), our product does not store, collect or process any personal information.

Document Conventions

Convention	Description
Bold	Identifies elements on a screen.
Narrowfont	Identifies service names and locations in the format <i>folder.subfolder.service</i> , APIs, Java classes, methods, properties.
<i>Italic</i>	Identifies: Variables for which you must supply values specific to your own situation or environment. New terms the first time they occur in the text. References to other documentation sources.
Monospace font	Identifies: Text you must type in. Messages displayed by the system.

Convention	Description
	Program code.
{ }	Indicates a set of choices from which you must choose one. Type only the information inside the curly braces. Do not type the { } symbols.
	Separates two mutually exclusive choices in a syntax line. Type one of these choices. Do not type the symbol.
[]	Indicates one or more options. Type only the information inside the square brackets. Do not type the [] symbols.
...	Indicates that you can type multiple options of the same type. Type only the information. Do not type the ellipsis (...).

Online Information and Support

Software AG Documentation Website

You can find documentation on the Software AG Documentation website at "<http://documentation.softwareag.com>". The site requires credentials for Software AG's Product Support site Empower. If you do not have Empower credentials, you must use the TECHcommunity website.

Software AG Empower Product Support Website

If you do not yet have an account for Empower, send an email to "empower@softwareag.com" with your name, company, and company email address and request an account.

Once you have an account, you can open Support Incidents online via the eService section of Empower at "<https://empower.softwareag.com/>".

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](#)".

If you have any questions, you can find a local or toll-free number for your country in our Global Support Contact Directory at "https://empower.softwareag.com/public_directory.asp" and give us a call.

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.

Data Protection

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

1 Using Command Central or Software AG Installer to Install Products

■ Using Command Central to Install Products	12
■ Using Software AG Installer to Install Products	12

Using Command Central to Install Products

Use Command Central in distributed staging and production environments to install, patch, configure, manage, and upgrade Software AG products; to create database components; and to connect products to database components. With Command Central, you can perform these tasks on multiple remote machines from one location.

Using Software AG Installer to Install Products

Use Software AG Installer to install Software AG products locally on standalone development machines.

2 Hardware Requirements and Operating System, Browser, and RDBMS Support

■ Hardware Requirements	14
■ Operating System Support	20
■ Browser Support	20
■ RDBMS Support	20

Hardware Requirements

Minimum and Recommended Hardware Requirements

Important: The hardware requirements in this section 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 18](#). You must determine the overall hardware requirements for each of your machines based on the overall requirements of your operating system and software.

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

Product	Hard Drive Space	RAM	Cores
ActiveTransfer Agent	10MB	512MB	1
ActiveTransfer Server/Gateway	100MB (200MB)	512MB	1
AgileApps*	21GB	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

Product	Hard Drive Space	RAM	Cores
API Gateway**	2GB	1.75GB	1 (2)
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	
CentraSite			
■ Registry Repository	3GB	4GB (8GB)	1 (2)
■ Application Server Tier	3GB	4GB (8GB)	1 (2)
Software AG Designer Eclipse platform	500MB (1.5GB)	1.5GB (2GB)	1 (2)
EntireX			
■ Core Files	100MB	2GB	1
■ Application Monitoring Data Collector	5MB	250MB	1
■ Command Line Scripts	2MB	100MB	1

Product	Hard Drive Space	RAM	Cores
■ Mainframe Broker Monitoring	50MB	500MB	1
■ Broker	20MB Windows, 30MB UNIX	1GB	2
■ 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	2GB	2GB	2
Microservices Runtime	250MB (400MB)	1 GB	1
Mobile Designer	1GB	1GB	1
My webMethods Server	300MB (500MB)	1GB (2GB)	1
OneData			

Product	Hard Drive Space	RAM	Cores
■ 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
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 hard drive space for Integration Server, with which AgileApps is co-hosted.

**Includes 300MB hard drive space and 512MB RAM for Internal Data Store, 130MB hard drive space for Microgateway, and 200MB hard drive space and 256MB RAM for Kibana.

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

Additional Hardware Requirements

The table below lists hardware requirements beyond those listed in the previous section.

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	See the “AgileApps wiki” .
Apama	Hard drive space, RAM, and core 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.
Software AG Designer	For Software AG Designer Eclipse platform, 700MB additional free hard drive space is required in the system temp directory (Windows) or temporary (UNIX) 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 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 cores based on the number of simultaneous users.
Mobile Designer	You will need additional hard drive space to accommodate the SDKs for the mobile devices for which you want to develop applications.

Product	Additional Hardware Requirements
OneData	OneData needs additional hard drive space for temp, import, and export files in the <i>Software AG_directory\OneData\config</i> directory.
Optimize	
<ul style="list-style-type: none"> ■ Analytic Engine 	<p>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.</p> <p>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.</p>
<ul style="list-style-type: none"> ■ Web Service Data Collector 	<p>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.</p>
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 <i>webMethods Trading Networks User's Guide</i> .
Universal Messaging	<p>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.</p> <p>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 cores to handle high-performance cluster event processing.</p>

Product	Additional Hardware Requirements
	If you enable a Universal Messaging realm server to use shared memory, the realm server needs two additional cores 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.

Operating System Support

For information on operating systems that 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

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

3 Preparing for Product Installation

■ Shut Down Products	22
■ Prepare for Installation For All Products (Windows System)	22
■ Prepare for Installation in a Microsoft Azure Virtual Machine	22
■ Prepare for Apama Installation (Linux System)	22
■ Prepare for API Portal Installation (UNIX System)	22
■ Prepare for CentraSite Installation (UNIX System)	23
■ Prepare for ApplinX Installation	23
■ Prepare for EntireX Installation	24
■ Prepare for Internal Data Store Installation (Linux System)	24
■ Prepare for MashZone NextGen Installation (Linux System)	25
■ Prepare for Mobile Designer Installation	26
■ Prepare for My webMethods Server Installation (UNIX System)	26
■ Prepare for OneData Installation (Windows System)	26
■ Prepare for Optimize Analytic Engine Installation (UNIX System)	26
■ Prepare for Zementis Predictive Analytics Installation	27
■ Prepare Database Connection Information	27

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 files that are locked by the operating system can be updated. For instructions, see the product documentation.

Prepare for Installation For All Products (Windows System)

Install the latest version (14.12.25810.0 or later) of the Microsoft Visual C++ Redistributable for Visual Studio 2017.

Note: If you have an older version installed and did not shut down all applications that use it, you will have to reboot your machine and restart the installation process, which will then install this package for you.

Prepare for Installation in a Microsoft Azure Virtual Machine

If you are going to install your products in a Microsoft Azure VM, install the libXrender library in the VM.

Prepare for Apama Installation (Linux System)

The file system to which the Apama 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 for API Portal Installation (UNIX System)

- Check the setting for shared memory (kernel parameter `shmmmax`) by executing the command `sysctl -a | fgrep kernel.shmmmax`. If the value is less than 629145600, log on as root user and increase the value by executing `sysctl -w kernel.shmmmax=629145600` or `echo "kernel.shmmmax=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`.
- API Portal uses Elasticsearch 5.6.5. If you are going to install on a Linux system, the maximum map count checks that the kernel allows a process to have at least 262,144 memory-mapped areas. Go to the `/etc/sysctl.conf` file and set `vm.max_map_count value=262144`. Then restart the machine to make this change to take effect.

Prepare for CentraSite Installation (UNIX System)

- 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`.
- 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 `nfiles` 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.

Prepare for ApplinX Installation

- 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 for EntireX Installation

- 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.
- 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 `nfiles` 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.

Prepare for Internal Data Store Installation (Linux System)

Internal Data Store uses Elasticsearch 5.6.4. Elasticsearch cannot be run as the root user on a Linux system, so you must create an Internal Data Store user and install and run the Internal Data Store as that user.

Elasticsearch does several checks before starting up. Software AG recommends that you review the bootstrap checks (search on "Elasticsearch bootstrap checks") and important system configuration settings (search on "Elasticsearch important system configuration") before starting the Internal Data Store. In particular, you may need to adjust these settings:

1. Check the setting 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 65536, log on as the root user and increase the value to at least 65536 by executing `sysctl -w fs.file-max=65536` or `echo "fs.file-max=65536" >> /etc/sysctl.conf`, then activate the new value by executing `sysctl -p`.
2. Check the Internal Data Store user settings for the maximum number of open file descriptors by executing the commands `ulimit -Hn` and `ulimit -Sn`, where `-Hn` is the hard limit and `-Sn` is the soft limit. If the value is less than 65536, log on as the Internal Data Store user and increase the value to at least 65536 by executing `ulimit -n 65536`. To permanently save this setting for the user, execute:


```
echo "user_name soft nofile 65536" >> /etc/security/limits.conf
echo "user_name hard nofile 65536" >> /etc/security/limits.conf
```
 3. Check the setting for the system-wide maximum map count (kernel parameter `vm.max_map_count`) by executing the command `sysctl -a | fgrep vm.max_map_count`. If the value is less than 262144, log on as the root user and increase the value to at least 262144 by executing `sysctl -w vm.max_map_count=262144` or `echo "vm.max_map_count=262144" >> /etc/sysctl.conf`, then activate the new value by executing `sysctl -p`.
 4. Check the Internal Data Store user settings for the maximum number of processes by executing the command `ulimit -u`. If the value is less than 4096, log on as the Internal Data Store user and increase the value to at least 4096 by executing `ulimit -n 4096`. To permanently save this setting for the user, execute:


```
echo "user_name soft nproc 4096" >> /etc/security/limits.conf
echo "user_name hard nproc 4096" >> /etc/security/limits.conf
```

Prepare for MashZone NextGen Installation (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 for Mobile Designer Installation

- Install an integrated development environment (IDE). Software AG recommends using the Eclipse provided with Software AG Designer.
- 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.

Prepare for My webMethods Server Installation (UNIX System)

Set sufficient user limits for the shell you will 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 `nfiles` to 8192. For more information about setting and displaying the ulimits, read the man page or ask your system administrator.

Prepare for OneData Installation (Windows System)

Install Microsoft Visual C++ Runtime 2005 or higher in any directory on the machine that will host Locate.

Prepare for Optimize Analytic Engine Installation (UNIX System)

Set sufficient user limits for the shell you will 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 `nfiles` to 8192. For more information about setting and displaying the ulimits, read the man page or ask your system administrator.

Prepare for Zementis Predictive Analytics Installation

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).

Prepare Database Connection Information

Some products require you to supply database connection information during installation. 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.
- 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).
- 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
```

- For Integration Server and Microservices Runtime, for MySQL Community Edition 5.7, you must specify connection options for the relaxAutoCommit, useLegacyDatetimeCode, and serverTimezone parameters. For example, you can provide these connection options as follows:

```
jdbc:mysql://host:{3306|port}/databaseName?
relaxAutoCommit=true&useLegacyDatetimeCode=false&serverTimezone=PST
```

- 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.

4 Preparing for Database Component Creation

- Database Component Descriptions and Installation Requirements 30
- Database Drivers for DB2, MySQL Enterprise Edition, Oracle, and SQL Server 38
- Prepare for Database Component Creation in MySQL 39
- Database User and Storage, and Database Settings 39

Database Component Descriptions and Installation Requirements

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 its user interface.

This section lists the products that have database components and describes the database components to create. Each database component is named for the type of data that is written to it.

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

ActiveTransfer Server Database Components

If you are going to install ActiveTransfer Server, 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.

API Gateway Database Components

API Gateway comes with an internal data store 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.

CloudStreams Database Components

If you are going to install CloudStreams, create the CloudStreamsEvents database component. CloudStreams Server writes lifecycle (startup/shutdown), error, policy violation, monitoring, performance metric, and, optionally, 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, create a single CloudStreamsEvents database component for them to share, regardless of whether they are clustered or not clustered.

Software AG Designer 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.

Dynamic Business Orchestrator Database Components

If you are going to install Dynamic Business Orchestrator, create the DynamicBusinessOrchestrator and ProcessAudit database components.

Dynamic Business Orchestrators write process instance operational data for the processes they orchestrate to the DynamicBusinessOrchestrator component. If you are executing dynamic business processes using an Integration Server cluster, create a single DynamicBusinessOrchestrator database component for the cluster to share. If you are not using a cluster, create one DynamicBusinessOrchestrator database component for each Dynamic Business Orchestrator.

Dynamic Business Orchestrator and the following write to the ProcessAudit database component:

- Dynamic Business Orchestrators write process audit data for business processes they orchestrate.
- Third-party products can write process execution data.

The following read from the ProcessAudit database component:

- Optimize Analytic Engines read process data 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 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 so you can monitor process instances in real time.

If you are executing dynamic business processes using an Integration Server cluster, create a single ProcessAudit database component for the cluster to share. If you are not using a cluster, create one ProcessAudit database component for each Dynamic Business Orchestrator.

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

Integration Server or Microservices Runtime Database Components

The table below describes the types of data Integration Server or Microservices Runtime can persist.

Database Component	Types of Data	When Data is Written
ISInternal	Service results, scheduled tasks, client certificate mappings, runtime data for pub.storage services, guaranteed delivery transactions, trigger joins, active OpenID requests, WS-ReliableMessaging runtime data, and configuration and runtime data for OAuth.	You are using the features listed in the Types of Data column
ISCoreAuditLog	<ul style="list-style-type: none"> ■ Error, guaranteed delivery, service, security, and session audit data. 	The audit logger for the type of data is enabled
	<ul style="list-style-type: none"> ■ 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

Database Component	Types of Data	When Data is Written
DistributedLocking	Information that coordinates access to resources across distributed servers and processes.	Executing services in the pub.storage folder.

External RDBMS versus Embedded Database

You must use an external RDBMS with Integration Server or Microservices Runtime 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 data.
- Have services log their input pipelines, or post user-defined progress messages.
- Use triggers for joins.
- Use pub.storage for high-volume storage.
- Write Process Audit Log and Process Engine data for business processes.
- Use Business Rules.
- Use Monitor.

If you are not going to do any of the above, and your database demands are low, your needs might be met by using the embedded database. For example, you can use the embedded database when you are using Integration Server or Microservices Runtime as a runtime for adapters, eStandards Modules, API Gateway, or Deployer.

Using the Embedded Database

Integration Server and Microservices Runtime use Derby, a file-based database, as their embedded database. When you use the embedded database, Integration Server or Microservices Runtime writes IS Internal and Cross Reference data to that database, and writes IS Core Audit Log data to files.

Important: 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.

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 “[Database](#)”).

[Component Descriptions and Installation Requirements” on page 30](#)) and configure Integration Server or Microservices Runtime to write to them (see [“Connecting Products to Database Components” on page 65](#)). You must also complete the steps for switching from the embedded database to an external RDBMS (see *webMethods Integration Server Administrator’s Guide*).

Using the External RDBMS

When you use an external RDBMS, 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). 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 [“ Database Component Descriptions and Installation Requirements” on page 30](#).

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.

Mobile Support Database Components

If you are going to install 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.

Monitor Database Components

If you are going to archive data from the ISCoreAudit and ProcessAudit database components (called the "source" database components in this section), create the Archive database component. The Archive database component includes a copy of the source 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.

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.

My webMethods Server Database Components

If you are going to install My webMethods Server or the Task Engine on Integration Server, create the MywebMethodsServer database component.

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

The Task Engine on My webMethods Server and the Task Engine on Integration Server write task status, task distribution, and business data. The Task Engine on My webMethods Server displays this data in its user interface in My webMethods. The Task Engine on Integration Server displays this data in its gadgets in Business Console.

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.

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.

OneData Database Components

If you are going to install OneData, create the OneData database components.

The table below describes the database components to create for OneData.

Database Component	Description
OneDataMetadata	OneData writes internal configuration data to this database component.
OneDataWorkArea	OneData users create data objects and work-in-progress data values to this database component.
OneDataReleaseArea	Contains the same data objects as the OneDataWorkArea database component, and approved data values. Data

Database Component	Description
	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 Database Components

If you are going to install Optimize, create the Analysis and ProcessTracker database components.

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).

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).

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.

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.

Process Engine Database Components

If you are going to install Process Engine, create the ProcessEngine and ProcessAudit database components.

Process Engines write process execution data for processes they orchestrate to the ProcessEngine 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.

Process Engine and the following write to the ProcessAudit database component:

- Process Engines write process audit data for business processes they orchestrate.
- Task Engines write task audit data for task steps orchestrated by Process Engine.
- Third-party products can write process execution data.

The following read from the ProcessAudit database component:

- Optimize Analytic Engines read process data 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 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 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. 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.

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.

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

Rules Engine Database Components

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

Trading Networks Server Database Components

If you are going to install 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.

If you want to archive Trading Networks data, 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.

Database Drivers for DB2, MySQL Enterprise Edition, Oracle, and SQL Server

Software AG products that support DB2, MySQL Enterprise Edition, 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, MySQL Enterprise Edition, Oracle, and SQL Server RDBMSs unless the documentation for a product states that it supports another driver.

Prepare for Database Component Creation in MySQL

Install MySQL Community Edition Database Driver

Software AG products 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
```

Edit the MySQL Configuration File

1. Open the MySQL configuration file.
2. Remove the values NO_ZERO_IN_DATE and NO_ZERO_DATE from the sql_mode variable.
3. If MySQL is on a Linux system, set the parameter lower_case_table_names=1.

Database User and Storage, and Database Settings

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 table below lists the permissions the database user must have and the file in which to specify those permissions.

RDBMS	Path to File
DB2	<i>Software AG_directory</i> \common\db\scripts\db2\storage\30\create\db2_str_c_ddl_user.sql
MySQL Community Edition	<i>Software AG_directory</i> \common\db\scripts\mysql\storage\30\create\mys_str_c_ddl_user.sql
MySQL Enterprise Edition	<i>Software AG_directory</i> \common\db\scripts\mysql\storage\30\create\mys_str_c_ddl_user.sql
Oracle	<i>Software AG_directory</i> \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	<i>Software AG_directory</i> \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 or filegroups. Do not create them manually.

The table below shows the database storage convention that all database components follow by default.

RDBMS	Data Storage Unit	Index Storage Unit	BLOB Storage Unit
DB2 tablespaces	WEBMDATA	WEBMINDX	WEBMBLOB
MySQL	Default	Default	Default

RDBMS	Data Storage Unit	Index Storage Unit	BLOB Storage Unit
Oracle tablespaces	WEBMDATA	WEBMINDX	WEBMDATA
SQL Server filegroups	Primary	Primary	Primary

You can use a different convention when you create database storage.

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.

The table below lists the character sets and sort order recommended by Software AG.

For	Software AG Recommendation
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.

The table below lists the character set and sort order recommended by Software AG.

For	Software AG Recommendation
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.

The table below lists the character sets and sort order recommended by Software AG.

For	Software AG Recommendation
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.

The table below lists the character sets and sort order recommended by Software AG.

For	Software AG Recommendation
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

The table below provides the page and block sizes specified to use for each type of RDBMS.

RDBMS	Required Page and Block Size	Default
DB2	32K page size	4K

RDBMS	Required Page and Block Size	Default
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 options 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.

5 Installing Products and Creating Database Components

■ Installing Products and Creating Database Components Using Command Central	46
■ Installing Products Using Software AG Installer and Creating Database Components Using Database Component Configurator	46

Installing Products and Creating Database Components Using Command Central

Follow the instructions in *Software AG Command Central Help* to install your new products and fixes and create your database components using Command Central.

Installing Products Using Software AG Installer and Creating Database Components Using Database Component Configurator

Installing Products Using Software AG Installer

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 Software AG Installer*.
- 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.
- 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.

Before Installing Using Software AG Installer

Gather Information

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.

Save and Verify Product License Files

Software AG sends product 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 performs the validations below. You will only be able to install a product if its license file passes these checks. The products will repeat these checks at startup.

- The license file must be the correct one for the product.
- The license file must not have been changed.
- The license file must not have expired.
- The product must be licensed on the operating system on which you are installing.

Read and Perform Required Actions in Using Software AG Installer

Read and perform the required actions in *Using Software AG Installer*. For example, for Windows create a user account for installation that has the proper privileges; for UNIX, create a non-root user account.

API Gateway Requirements

Internal Data Store (Elasticsearch) is not supported on as many operating systems as API Gateway. If you install API Gateway on an operating system that is not supported by Internal Data Store, install Internal Data Store on a supported operating system and then configure API Gateway to use the Internal Data Store. For configuration instructions, see the product documentation.

CentraSite Requirements

- You can install the CentraSite Registry Repository and Application Server Tier on different machines, in different directories on the same machine, or in the same directory but at different times (that is, in two different runs of the installer). Whichever approach you follow, you must always install the Registry Repository first.

If you are going to install multiple instances of CentraSite components on the same machine:

- Each instance of the CentraSite Registry Repository, Application Server Tier, and Platform Manager must use unique ports.

- Software AG recommends that you have all Application Server Tier instances on the machine running before you install a new Application Server Tier instance, so the installer can warn you if a port is already in use.
- All Registry Repository instances on the machine must be shut down before you install a new Registry Repository instance. However, because the instances are not running, the installer will not be able to warn you if the default port is in use; you will have to find and assign a unique port yourself.

Deployer and Asset Build Environment Requirements

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 development 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 *composites* 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.

Install Products Using Software AG Installer

Start the Installer and Provide General Information

For all products except MashZone NextGen, follow the instructions in *Using Software AG Installer* 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.

If you are installing MashZone NextGen 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 CentraSite, 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.

Choose Products to Install

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

Note: Software AG Installer offers a feature that helps you create typical development installations that enable you to perform a certain task (for example, API management). When you select a typical development installation, the installer automatically selects all products that make up that installation. You can select more products or deselect products if desired.

- Some products such as Deployer can run on either Integration Server or Microservices Runtime. The installer will ask which host server you want to use.
- 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.
- The license for API Gateway allows only partial functionality of its host server (that is, Integration Server or Microservices Runtime). Install API Gateway on its own host server. Do not install other products on that host server or those other products might not work properly.
- The license for CloudStreams allows only partial functionality of its host server (that is, Integration Server or Microservices Runtime). Install CloudStreams on its own host server. Do not install other products on that host server because those other products might not work properly.
- If you select packages for Integration Server or Microservices Runtime, 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*.
- 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 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 Runtime. If you install the package on Microservices Runtime, data is provided for unorchestrated business processes only. Data is not provided for orchestrated business processes or events.

Choose Language Packs 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. If you are installing CentraSite, and you select the CentraSite Application Server Tier language pack, you must also select the CentraSite Shared Files language pack.

Supply Product Configuration Information

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

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.

Note: When you install on a UNIX system, the installer might ask whether to use sudo and related scripts. 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.

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 or on providing a license after installation.

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. You can specify the following as the work directory:

- A directory that does not exist (recommended). The installer will create the directory, then populate it with standard files and directories.
- A non-empty directory. The installer will not write standard files and directories to the directory, to avoid overwriting existing user content.

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.

API Portal

The table below shows the system configurations you can select for API Portal to handle your expected usage needs.

System Configuration	Memory Consumption	Maximum Concurrent Users
Small	1 GB	50
Medium	4 GB	100
Large	8 GB	400

EntireX

The table below explains the fields on the EntireX installer panel.

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: <ul style="list-style-type: none"> ■ 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	<ul style="list-style-type: none"> ■ 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.
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.

Integration Server or Microservices Runtime

If you are installing API Gateway, specify the license file for API Gateway, not the license file for the host Integration Server or Microservices Runtime. Select **Embedded database**.

If you are installing CloudStreams, specify the license for CloudStreams, not the license for the host Integration Server or Microservices Runtime.

If you are installing Designer Workstation, specify the license for Designer Workstation.

If you are installing Deployer, select **Embedded database**.

If you are installing Integration Server or Microservices Runtime 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 whether to use an embedded database or external RDBMS in all other cases, see [“Preparing for Database Component Creation” on page 29](#).

OneData

The table below explains the fields on the OneData installer panel.

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, <i>prefix_md</i> , <i>prefix_wa</i> , and <i>prefix_ra</i>). 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 <code>catalogOptions=1</code> so that OneData can retrieve database metadata information for various operations (for example, <code>jdbc:wm:oracle://localhost:1521;serviceName=XE;catalogOptions=1</code>)

Field	Entry
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.

Universal Messaging

The table below explains the fields on the Universal Messaging installer panel when you choose to install a Universal Messaging realm server alone or with Enterprise Manager or Template Applications.

Field	Entry
License file	<p>The contents of this field varies as explained below.</p> <ul style="list-style-type: none"> ■ If you have the standard Universal Messaging license, specify the full path to that license file. ■ 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. <p>If you obtain a Universal Messaging license file after installation, you can add it at that time. For instructions, see the Universal Messaging documentation.</p>
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	<p>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.</p> <div style="background-color: #f0f0f0; padding: 10px;"> <p>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).</p> </div>

Field	Entry
Data directory	<p>Full path to the directory in which to store the data for the realm server instance.</p> <p>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.</p> <p>If you install multiple realm server instances on the same machine, use a different data directory for each instance.</p>
Configuration	<p>Initial configuration for the Universal Messaging realm server. Select the configuration for webMethods suite use cases.</p>

The table below explains the fields on the Universal Messaging installer panel when you choose to install Universal Messaging Enterprise Manager, Template Applications, or both, without also choosing to install a Universal Messaging realm server.

Field	Entry
Realm name/ client profile	<p>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.</p> <p>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).</p>
Realm server URL (RNAME)	<p>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.</p>

After Installing Using Software AG Installer

Manage Products Using Command Central

Software AG Installer automatically installed Command Central's agent, Platform Manager, with your products. This enables you to manage the products remotely after installation from one location using Command Central. For instructions on managing products from Command Central, see *Software AG Command Central Help*.

Install Latest Fixes

Install the latest fixes on the products you installed. You can use Command Central to install fixes on products you installed using Software AG Installer, or you can use Software AG Update Manager. For instructions, see *Software AG Command Central Help* or *Using 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.

For instructions on registering UNIX daemons, see the next section.

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.

If you installed a Zementis Standalone Server, product documentation is in the *Software AG_directory \Zementis \adapa-app \documents*. For instructions on setting up the server, see the *Zementis Deployment Guide*.

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.

Register Daemons and Set File Permissions (UNIX System)

If you choose to register daemons in the installer panels, but you created a script 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 the daemons and set file permissions. Perform these tasks now by running the script *Software AG_directory /bin/ afterInstallAsRoot.sh* with root permission (that is, call the script using the root user account or the sudo utility).

Note: The *afterinstallAsRoot.sh* script only registers the daemons you selected in the installer panels. If you want to change or add daemon registration for any product at a later time, see [“Registering Daemons to Automatically Start and Shut Down Products on UNIX Systems” on page 67](#).

Uninstall Products Using Software AG Installer

Uninstall EntireX

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 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` with root permission (that is, call the script using the root user account or the `sudo` utility).
- 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.

Uninstall All Other Products

Follow the instructions in *Using 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 Runtime, 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 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.

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` with root permission (that is, call the script using the root user account or the `sudo` utility).

- 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.

Create Database Components Using Database Component Configurator

Install the Database Component Configurator and Database Scripts

1. Download the Software AG Installer to any machine on the same network as your database server.
2. Start the Software AG Installer. In the product selection list, select **Database Component Configurator**, then continue through the installation process.
3. The Database Component Configurator has default environment settings. Go to the *Software AG_directory* \common\db\bin directory, open the setEnv.{bat|sh} file in the text editor, and edit the fields as necessary.

The table below shows the default environment settings for Database Component Configurator.

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.	<i>Software AG_directory</i> \common\db\logs

Database Component Configurator Command

Go to the *Software AG_directory* \common\db\bin directory and run the command below.

Each run of the configurator operates on a single schema.

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.

Main Parameters

```
dbConfigurator.{bat|sh} {-a|--action} action
{-d|--dbms}{oracle|sqlserver|db2luw}
{-l|--url} RDBMS_URL
[{-c|--component} db_component [, db_component ...] [{"-v|--version} product_release] |
  {-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]
```

Main Parameters

{-a | --action} *action*

The table below describes the parameters for the **{-a | --action}** parameter.

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. Before dropping any database component, shut down all products that are connected to the database component. After you drop all database components from a schema, you can drop the storage and revoke the database user permissions.
recreate	Performs the drop action and then the create action.
checkEncoding	Indicates whether your RDBMS is Unicode-enabled.

{-d | --dbms} {db2|luw|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.

The table below describes the parameters for the {-l | --url} parameter.

RDBMS	Sample URL Format
DB2	<pre>jdbc:wm:db2://server:{50000 port};databaseName=database [;option=value]...</pre> <p>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:</p> <pre>;AlternateId=schema ;"InitializationString=(SET CURRENT PATH=current_path,schema)"</pre> <p><code>AlternateID</code> is the name of the default schema used to qualify unqualified database objects in dynamically prepared SQL statements.</p> <p>If you are creating Optimize database components (that is, Analysis and ProcessTracker), you must specify <i>schema</i> using all uppercase letters. In addition, you must specify the options <code>CreateDefaultPackage=true</code>, <code>ReplacePackage=true</code>, and <code>DynamicSections=3000</code>. These settings will affect all database components in the same schema or database.</p> <p>If you are using a schema other than the default schema for the specified database user, you must also add the option <code>showSelectableTables=false</code> to the URL. The database driver queries <code>SYSIBM.SYSDBAUTH</code> at connection time to find out whether the current user is an admin or not. The results of this check allows a query to <code>getTables()</code> to return tables for which the user has <code>SELECT</code> privileges. An admin has implicit select privileges on all tables instead of the explicit privileges that other types of users have, so the driver has to alter its native SQL for <code>getTables()</code>. Set <code>showSelectableTables=false</code> and the driver will skip the check for <code>SYSIBM.SYSDBAUTH</code> at connection time.</p>
MySQL Community Edition	<pre>jdbc:mysql://server:{3306 port}/databaseName</pre>
MySQL Enterprise Edition	<pre>jdbc:wm:mysql://server:{3306 port};databaseName=database [;option=value]...</pre>

RDBMS	Sample URL Format
Oracle	<code>jdbc:wm:oracle://server:{1521 port};serviceName=service [;option=value]...</code>
	If you are creating storage and the Data Purge database component, you must specify the <code>sysLoginRole</code> connection option on the URL (for example, <code>;sysLoginRole=sysdba</code>).
SQL Server	<code>jdbc:wm:sqlserver://server:{1433 port};databaseName=database [;option=value]...</code>

{-c|--component} db_component [,db_component ...] | {-pr|--product} product [,product ...]} {-v|--version} latest

One of the following:

- Database user and storage.
- Database components and product release. For this product release you can specify `latest`. `--printComponents` lists database component codes and names; you can use either.
- Product database components and product release. For this product release you can specify `latest`. `--printProducts` lists the database components for specific products.

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

The table below describes the parameters for the `{-u | --user}` and `{-p | -password}` parameters.

Tasks	Values to specify
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 <code>dbo</code> will be assigned to that user. For Oracle, do not use the <code>SYSTEM</code> user to create the database components in the <code>SYSTEM</code> schema.
Creating a database user and storage in DB2	OS user to which to grant permissions.
Performing the <code>checkEncoding</code> action	Existing database user that has create session and create table privileges, and that database user's password.

Tasks	Values to specify
-------	-------------------

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.

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://vmxpd01:50000;databaseName=am01 -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://vmxpd01:50000;databaseName=am01 -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://vmxpd01:50000;databaseName=am01 -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://vmxpd01:50000;databaseName=am01 -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 ALTSHEMA 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=ALTSHEMA;InitializationString=\SET CURRENT
PATH=current_path,ALTSHEMA\" -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 ALTSHEMA 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=ALTSHEMA
;"InitializationString=(SET CURRENT PATH=current_path,ALTSHEMA) "
-u webmuser -p w3bmpass
```


6 Connecting Products to Database Components

This section explains how to connect Integration Server or hosted products to their database components. For information on connecting other products such as Optimize to their database components, see the product documentation.

You connect Integration Server to database components as follows:

- Define *JDBC database connection pools* that specify the connections to database servers that host database components.
- Direct predefined *functions* to write to their database components by pointing each function at the appropriate connection pool. For example, point the ISCoreAudit Log function to the connection pool for the ISCoreAudit database component, point the Xref function to the connection pool for the CrossReference database component, and so on. At run time, Integration Server creates a separate instance of the connection pool for each database component.

Integration Server can connect to the ActiveTransfer, APIGatewayEvents, Archive, BusinessRules, CloudStreamsEvents, CrossReference, DocumentHistory, ISInternal, ISCoreAudit, MywebMethodsServer, MobileSupport, ProcessAudit, ProcessEngine, Staging and Reporting, and TradingNetworks and TradingNetworksArchive database components.

For instructions on creating connection pools and pointing functions at them, see *webMethods Integration Server Administrator's Guide* and the points below.

- If you installed the Task Engine on Integration Server, and there is no MywebMethods Server in the same installation directory, create a connection pool that points to the MywebMethodsServer database component, and then point the CentralUsers function in Integration Server at that connection pool.
- 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.
- If you installed the Integration Server and 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. You can then drop the other, redundant instance.

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

■ Overview	68
■ Run the daemon.sh Command Line Tool	69
■ Product Daemons and rc-scripts	70
■ Troubleshooting	71

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`, where `sag` is a fixed prefix string, and `instance` is a unique identifier for the daemon instance. If you accidentally specify a non-unique `instance` during daemon registration, the `daemon.sh` script automatically adds `number` to make the init-script name unique.

The table below shows the locations of init-scripts generated by the `daemon.sh` script.

System	Location
SLES 11, 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 performed the installation).
- 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 70](#).

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 rc-script names and locations are listed in [“Product Daemons and rc-scripts” on page 70](#).

Important: The listed rc-scripts might offer options to install or uninstall, but these are daemon registration methods implemented by Tanuki Software, Ltd. and are not supported by Software AG. Use only the start or stop options with the rc-scripts, and use only `daemon.sh` to register or un-register daemons.

The table below lists the most commonly used command line options for the `daemon.sh` command line tool; they are the same for every supported UNIX platform.

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

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

When using the `daemon.sh -f` option to register or unregister a daemon, make sure to use the absolute path name. Below is an example that applies to most products.

```
daemon.sh -f Software AG_directory/API_Portal/server/sagyaprelease.sh
```

For products which can have multiple instances, such as Universal Messaging, Integration Server, Integration Agent, and My webMethods Server, the command has an additional option, `-n instance`. The command to use for Universal Messaging, for example, is as follows

```
daemon.sh -n umserver -f Software AG_directory/UniversalMessasging/server/  
umserver/bin/nserverdaemon
```

The table below shows the rc-script for each product daemon.

Product	Daemon	rc-script (under <i>Software AG_directory</i>)
ActiveTransfer Agent	Software AG ActiveTransfer Agent	profiles/MAG/bin/sagmagrelease
API Portal	APICloudAgentrelease	API_Portal/server/bin/ CloudAgentApp.sh
CentraSite Registry Repository	CentraSite Registry/ Repository	CentraSite/bin/saginm
EntireX Broker	EntireX Administration Service	EntireX/bin/sagexx
Integration Agent	Software AG Integration Agent	profiles/ <i>instance</i> /bin/sagintagentrelease
Integration Server, or Microservices Runtime	Software AG Integration Server	profiles/ <i>instance</i> /bin/sagisrelease
My webMethods Server	Software AG My webMethods Server	profiles/ <i>instance</i> /bin/sagmwsrelease_ <i>instance_name</i> <i>instance_number</i>
OneData	Software AG OneData Server	profiles/ODE/bin/sagodererelease

Product	Daemon	rc-script (under <i>Software AG_directory</i>)
Optimize Infrastructure Data Collector	Software AG Optimize Infrastructure Data Collector	profiles/InfraDC/bin/sagwidrelease
Optimize Web Service Data Collector	Software AG Optimize WS Data Collector	optimize/dataCollector/bin/sagoptiDC.sh
Optimize Analytic Engine	webMethods Optimize Analysis Engine	optimize/analysis/bin/sagoptiAE.sh
Platform Manager	Software AG Platform Manager	profiles/SPM/bin/sagspmrelease
Software AG Runtime	Software AG Runtime	profiles/CTP/bin/sagctprelease
Universal Messaging	nserverdaemon	UniversalMessaging/server/instance /bin/nserverdaemon

Troubleshooting

On AIX, the maximum length for an init-script name is limited to 14 characters (see the man-page for the mkitab command). If daemon.sh reports an error with the mkitab command, you can specify a shorter init-script name with the -n option of daemon.sh. Using -n *instance* is just a recommendation; you can use any shorter name.

B Operating in International Operating Environments

■ Overview	74
■ Language Packs	74
■ Extended Character Set	75
■ Configure Browsers and JRE Font	75
■ Configure the Proper Locale	76

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 U.S. English files necessary to run your products is always installed with the 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 on all supported operating systems except Mac OS X.

Language Packs for Software AG Designer

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 Oxygen-compatible language packs 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

Software AG 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

For supported browser releases, see *System Requirements for Software AG Products*.

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.

The table below explains how to indicate the languages to display in your browser 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.

Configure the Proper Locale

The table below provides instructions on how to configure a locale if you are going to install or run your products in a locale other than U.S. English.

System	Action
Windows	<p>These instructions are for Windows 7; instructions for other Windows systems vary slightly.</p> <ol style="list-style-type: none">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	<p>Set your shell to the proper locale by setting the LANG environment variable. Run the appropriate command below:</p> <ul style="list-style-type: none">■ In an sh shell: <code>LANG=<i>locale</i>; export LANG</code>■ In a csh shell: <code>setenv LANG <i>locale</i></code>

C Running in a Cloud Infrastructure as a Service (IaaS) Environment

■ Overview	78
■ Amazon Web Services	78
■ Microsoft Azure	79

Overview

This appendix lists reference material you need to set up Software AG workload to operate in a cloud IaaS environment. Search the Internet using the text provided below and click resulting links that lead to the Amazon or Microsoft web site.

Amazon Web Services

The table below lists reference material for Amazon web services and how to find it on the Internet.

For this reference material	Search using this text
Minimum services	<ul style="list-style-type: none"> ■ AWS Elastic Compute Cloud (EC2) ■ AWS Virtual Private Cloud (VPC)
Network services	<ul style="list-style-type: none"> ■ AWS Direct Connect ■ AWS Manage Cloud DNS - Domain Name System - Amazon Route 53 ■ AWS Cloud Security Resources
AWS services offering	<ul style="list-style-type: none"> ■ AWS Cloud Products and Services
Best practices white papers	<ul style="list-style-type: none"> ■ AWS Architecting for the Cloud ■ AWS Security Best Practices - Cloud Computing Services ■ AWS Trusted Advisor Best Practices (Checks) ■ AWS Shared Responsibility Model
Training and certification	<ul style="list-style-type: none"> ■ AWS Training ■ AWS Certification
Security	<ul style="list-style-type: none"> ■ AWS Compliance Certification

Microsoft Azure

The table below lists reference material for Microsoft Azure and how to find it on the Internet.

For this reference material	Search using this text
Minimum services	<ul style="list-style-type: none"> ■ Azure Virtual Machines ■ Azure Virtual Network - Virtual Private Cloud
Network services	<ul style="list-style-type: none"> ■ Azure Express Route ■ Azure Active Directory - Access & Identity - IDaaS
Microsoft services offering	<ul style="list-style-type: none"> ■ Directory of Azure Services
Best practices white papers	<ul style="list-style-type: none"> ■ Azure Network Security Best Practices ■ Azure Reference Architectures
Training and certification	<ul style="list-style-type: none"> ■ Microsoft Azure Administration for AWS Administrators ■ Azure Certification and Training
Security	<ul style="list-style-type: none"> ■ Microsoft Azure Trust Center