


IBM Information Server


Version 8.0.1



IBM Information Server Planning, Installation, and Configuration Guide
Revised March 2007

IBM Information Server

Version 8.0.1



IBM Information Server Planning, Installation, and Configuration Guide
Revised March 2007

Note

Before using this information and the product that it supports, read the information in “Notices” on page 143.

Contents

Chapter 1. Overview of installing IBM

Information Server 1

Architecture and concepts	1
Layers	4
Clients	6
Domain	7
Engine	9
Metadata repository	12
Documentation	12
ODBC drivers	14

Chapter 2. Planning to install IBM

Information Server. 15

Planning for migration to IBM Information Server	15
Licensing requirements for IBM Information Server	15
Required users and groups for installing IBM Information Server (Windows)	15
Required users and groups for installing IBM Information Server (Linux, UNIX)	16
User registries	18
Methods of installation	18
Order of installation	19
Installation planning worksheet	19
System requirements for IBM Information Server	20
WebSphere Federation Server system requirements	20

Chapter 3. Preparing your system for installation 21

Preparing your system for installation (Windows)	21
Preparing to install the metadata repository	23
Preparing to install the domain (Windows)	29
Preparing to install the engine (Windows)	30
Creating users and groups (Windows)	31
Users and groups by layer (Windows)	31
Configuring the analysis database and the analysis engine	32
Preparing your system for installation (Linux, UNIX)	32
Preparing to install the metadata repository	34
Preparing to install the domain (Linux, UNIX)	39
Preparing to install the engine (Linux, UNIX)	40
Creating groups, user names and passwords before you install IBM Information Server (Linux, UNIX)	43
Users and groups by layer (Linux, UNIX)	44
Creating a server profile for existing instances of WebSphere Application Server	44
Configuring the analysis database and the analysis engine	45

Chapter 4. Installing IBM Information Server 47

Installing IBM Information Server (Windows)	47
Installing in graphical mode	47

Installing in console mode	48
Installing in silent mode	48
Stopping and starting the WebSphere DataStage server	49
Installing IBM Information Server (Linux, UNIX)	50
Installing in graphical mode	50
Installing IBM Information Server in console mode	51
Installing in silent mode	51
Stopping and starting the WebSphere DataStage server	52

Chapter 5. Configuring IBM Information Server 53

Creating users and groups for IBM Information Server	53
Setting up permissions and groups for Windows 2003 Server	54
Setting up permissions and groups to access WebSphere DataStage	54
Setting up permissions and groups to access WebSphere DataStage if the server is a domain controller	56
Configuring credential mappings for WebSphere DataStage and WebSphere QualityStage users	57
Assigning WebSphere DataStage and WebSphere QualityStage project roles.	58
WebSphere DataStage and WebSphere QualityStage user roles	59
Configuring WebSphere DataStage and WebSphere QualityStage	59
Tuning the WebSphere DataStage and WebSphere QualityStage server (Linux, UNIX)	60
Configuring the WebSphere DataStage and WebSphere QualityStage server environment (Windows).	61
Configuring a computer to be a processing node for MPP (Windows)	61
Configuring the WebSphere DataStage server environment (Linux, UNIX)	63
Configuring the WebSphere DataStage and WebSphere QualityStage server through PAM	65
Configuring plug-in stages(Linux, UNIX)	65
Configuring NLS for WebSphere DataStage and WebSphere QualityStage (Linux, UNIX)	66
Configuring the parallel engine (Linux, UNIX)	66
Configuring the parallel engine (Windows).	74
Configuring a computer to be a processing node for MPP (Windows)	76
Setting up database connectivity	78
Configuring access to ODBC data sources for WebSphere DataStage and WebSphere QualityStage (Windows)	83
Configuring access to ODBC data sources (Linux, UNIX)	83

Configuring user directories for IBM Information Server	87
Configuring the local operating system user registry	88
Configuring the IBM Information Server user directory	89
Configuring LDAP for user registries	91
Configuring WebSphere Federation Server	92
Checking the setup of the federated server	92
Creating a federated database	97
Confirming and registering the Microsoft SQL Server Client Version 2000 driver (Windows)	101
Configuring security for federated systems	102
Oracle Label Security and federated systems	102
Overview of the user mapping plugin for external repositories	102
LDAP sample plugin	110
Configuring a connection to the analysis database for WebSphere Information Analyzer	120
Validating your installation	121
Tutorials	124

Chapter 6. Installing additional components 125

Updating the license file	125
Enabling the SAS connector (Linux, UNIX)	125
WebSphere MetaBrokers and bridges	126
The Multi-Client Manager	126
MVS components for WebSphere DataStage	126
Installing COBOL copybook files	127
The runtime library for MVS jobs	128
Installation verification process	129
Installing additional components and updates for WebSphere Federation Server	130
Updating WebSphere Information Integration documentation	130
Installing fix packs	130
Adding the wrapper development kit to a system where DB2 is installed (Windows)	130

Adding the wrapper development kit to a system with DB2 installed (UNIX)	131
--	-----

Chapter 7. Troubleshooting installations 133

Installation log files	133
Other log files	133
Troubleshooting the installation process when it appears to be stopped	133
Troubleshooting command-line scripts that fail	134
Runtime error during removal of the ASBAgent component	134
Troubleshooting the user mapping plugin for federated systems	134

Chapter 8. Removing IBM Information Server 137

Removing IBM Information Server from your system (Windows)	137
Removing profiles (Windows)	138
Removing the Multi-Client Manager (Windows)	138
Removing IBM Information Server from your system (Linux, UNIX)	138
Removing profiles (Linux, UNIX)	139
Removing the WebSphere QualityStage migration script from your system	139

Accessing information about the product 141

Providing comments on the documentation	141
---	-----

Notices 143

Trademarks	145
----------------------	-----

Index 147

Chapter 1. Overview of installing IBM Information Server

IBM Information Server is a comprehensive, unified foundation for enterprise information architectures that is capable of meeting any information volume requirement. It combines the capabilities of WebSphere DataStage, WebSphere QualityStage, WebSphere Information Analyzer, WebSphere Business Glossary, WebSphere Information Services Director, and WebSphere Federation Server technologies into a single unified platform, allowing companies to easily understand, cleanse, transform, and deliver trustworthy and context-rich information.

The installation program installs all the IBM® Information Server technologies for which you have a license.

Architecture and concepts

IBM Information Server provides a unified architecture that works with all types of information integration. Common services, unified parallel processing, and unified metadata are at the core of the server architecture.

The architecture is service oriented, enabling IBM Information Server to work within an organization's evolving enterprise service-oriented architectures. A service-oriented architecture also connects the individual suite components of IBM Information Server.

By eliminating duplication of functions, the architecture efficiently uses hardware resources and reduces the amount of development and administrative effort that are required to deploy an integration solution.




Figure 1. IBM Information Server high-level architecture

Figure 1 shows the top levels of the IBM Information Server architecture.

Unified parallel processing engine

Much of the work that IBM Information Server does takes place within the parallel processing engine. The engine handles data processing needs as diverse as performing analysis of large databases for WebSphere® Information Analyzer, data cleansing for WebSphere QualityStage, and complex transformations for WebSphere DataStage®. This parallel processing engine is designed to deliver:

- Parallelism and pipelining to complete increasing volumes of work in decreasing time windows
- Scalability by adding hardware (for example, processors or nodes in a grid) with no changes to the data integration design
- Optimized database, file, and queue processing to handle large files that cannot fit in memory all at once or with large numbers of small files

Common connectivity

IBM Information Server connects to information sources whether they are

structured, unstructured, on the mainframe, or applications. Metadata-driven connectivity is shared across the suite components, and connection objects are reusable across functions.

Connectors provide design-time importing of metadata, data browsing and sampling, run-time dynamic metadata access, error handling, and high functionality and high performance run-time data access. Prebuilt interfaces for packaged applications called *Packs* provide adapters to SAP, Siebel, Oracle, and others, enabling integration with enterprise applications and associated reporting and analytical systems.

Unified metadata

IBM Information Server is built on a unified metadata infrastructure that enables shared understanding between business and technical domains. This infrastructure reduces development time and provides a persistent record that can improve confidence in information. All functions of IBM Information Server share the same metamodel, making it easier for different roles and functions to collaborate.

A common metadata repository provides persistent storage for all IBM Information Server suite components. All of the products depend on the repository to navigate, query, and update metadata. The repository contains two kinds of metadata:

Dynamic

Dynamic metadata includes design-time information.

Operational

Operational metadata includes performance monitoring, audit and log data, and data profiling sample data.

Because the repository is shared by all suite components, profiling information that is created by WebSphere Information Analyzer is instantly available to users of WebSphere DataStage and WebSphere QualityStage, for example.

The repository is a J2EE application that uses a standard relational database such as IBM DB2®, Oracle, or SQL Server for persistence (DB2 is provided with IBM Information Server). These databases provide backup, administration, scalability, parallel access, transactions, and concurrent access.

Common services

IBM Information Server is built entirely on a set of shared services that centralize core tasks across the platform. These include administrative tasks such as security, user administration, logging, and reporting. Shared services allow these tasks to be managed and controlled in one place, regardless of which suite component is being used. The common services also include the metadata services, which provide standard service-oriented access and analysis of metadata across the platform. In addition, the common services layer manages how services are deployed from any of the product functions, allowing cleansing and transformation rules or federated queries to be published as shared services within an SOA, using a consistent and easy-to-use mechanism.

IBM Information Server products can access three general categories of service:

Design

Design services help developers create function-specific services that can also be shared. For example, WebSphere Information

Analyzer calls a column analyzer service that was created for enterprise data analysis but can be integrated with other parts of IBM Information Server because it exhibits common SOA characteristics.

Execution

Execution services include logging, scheduling, monitoring, reporting, security, and Web framework.

Metadata

Using metadata services, metadata is shared “live” across tools so that changes made in one IBM Information Server product are instantly visible across all of the suite components. Metadata services are tightly integrated with the common repository and are packaged in WebSphere Metadata Server. You can also exchange metadata with external tools by using metadata services.

The common services layer is deployed on J2EE-compliant application servers such as IBM WebSphere Application Server, which is included with IBM Information Server.

Unified user interface


The face of IBM Information Server is a common graphical interface and tool framework. Shared interfaces such as the IBM Information Server console and Web console provide a common look and feel, visual controls, and user experience across products. Common functions such as catalog browsing, metadata import, query, and data browsing all expose underlying common services in a uniform way. IBM Information Server provides rich client interfaces for highly detailed development work and thin clients that run in Web browsers for administration.

Application programming interfaces (APIs) support a variety of interface styles that include standard request-reply, service-oriented, event-driven, and scheduled task invocation.

Layers

IBM Information Server is installed in five layers: client, metadata repository, domain, engine, and documentation. Product components are installed in each layer depending on the selections that you make. The layers can be installed on the same computer or on different computers.

The following diagram shows the client, documentation, metadata repository, domain, and engine installed across two computers. Computer A is a Windows[®] XP or Windows Vista[™] system and Computer B is a Windows Server 2003, Linux[®], or UNIX[®] system.



The following table shows the products and components that are installed by the installation wizard by layer. For a complete installation of a product and its components, you must select the product and components as indicated in the table when you install each layer. For example, to install WebSphere DataStage, select WebSphere DataStage and WebSphere QualityStage on the product selection panel of the installation wizard when you install the client, domain, and engine.

Table 1. Product modules and components installed into each layer

Product modules and components	Client	Metadata repository	Domain	Engine	Documentation
IBM WebSphere Business Glossary			X		
IBM WebSphere DataStage	X		X	X	
IBM WebSphere Federation Server				X	
IBM WebSphere Information Analyzer	X	X	X	X	

Table 1. Product modules and components installed into each layer (continued)

Product modules and components	Client	Metadata repository	Domain	Engine	Documentation
IBM WebSphere Information Services Director	X		X	X	
IBM WebSphere Metadata Server			X		
IBM WebSphere QualityStage	X		X	X	
Metadata repository		X			
PDF documentation					X

Clients

The clients for IBM Information Server include IBM WebSphere DataStage and WebSphere QualityStage clients and the IBM Information Server console. IBM WebSphere Information Analyzer and IBM WebSphere Information Services Director are part of the IBM Information Server console. The clients can be installed by the suite installer or the standalone client installer.

WebSphere DataStage and WebSphere QualityStage clients

You can use WebSphere DataStage and WebSphere QualityStage clients to administer, design, and manage WebSphere DataStage and WebSphere QualityStage projects.

WebSphere DataStage and WebSphere QualityStage Administrator

You can use the administrator client to perform administrative tasks for individual WebSphere DataStage and WebSphere QualityStage projects and on a single WebSphere DataStage engine.

WebSphere DataStage and WebSphere QualityStage Designer

For WebSphere DataStage and WebSphere QualityStage projects, you can use the designer client to build the jobs that cleanse, transform, and integrate data as it flows from a source to a target.

WebSphere DataStage and WebSphere QualityStage Director

You can use the director client to validate, schedule, run, and monitor server jobs and parallel jobs.

WebSphere QualityStage components

For WebSphere QualityStage projects, you can use optional components to re-engineer data and to test matching and blocking strategies before you run match jobs. All of the WebSphere QualityStage components are installed with WebSphere DataStage during a typical installation.

IBM Information Server console

The console is a task-oriented user interface that integrates suite components into one unified framework. The console contains workspaces that you use to create and complete information integration tasks such as investigating data, creating job schedules and logs, and deploying applications or Web services.

WebSphere Information Analyzer

WebSphere Information Analyzer is a next-generation data profiling and analysis tool that is designed to help business and data analysts understand the content, quality, and structure of their data sources by automating the data discovery process. WebSphere Information Analyzer helps you to assess the quality of your data by identifying inconsistencies, redundancies, and anomalies in your data at the column, table, and cross-table level. WebSphere Information Analyzer also makes inferences about the best choices regarding data values. Inferences help you to learn more about the structure of your data and what you can do to improve the quality of your data. You can access the WebSphere Information Analyzer client through the IBM Information Server console.

WebSphere Information Services Director

You can use WebSphere Information Services Director to create, deploy, and manage services. You can publish information services with these service interfaces, or bindings: Web services (SOAP over HTTP) and Enterprise JavaBeans™ (EJB). In addition, you can collaborate with other designers on the applications that are defined by a particular project. You can access the WebSphere Information Services Director client through the IBM Information Server console.

Multi-client manager

The multi-client manager is required to run multiple versions of WebSphere DataStage. It is installed on your computer when you install the clients on a computer where DataStage Version 7.5 clients are installed. You can also install it separately using the standalone installer in the MultiClientManager directory on the IBM Information Server CD, DVD, or downloaded installation image. If the multi-client manager is already installed, the installation wizard detects it and registers the new clients with it.

WebSphere MetaBrokers and bridges

WebSphere MetaBrokers and bridges, and the Import/Export client, are installed with the installer in the MetaBrokersAndBridges directory of the suite CD, DVD, or downloaded installation image.

WebSphere MetaBrokers and bridges move metadata in and out of the metadata repository. WebSphere MetaBrokers and bridges can be used by WebSphere DataStage and WebSphere QualityStage for directly importing and exporting metadata between tools.

Domain

The domain includes IBM WebSphere Metadata Server, IBM WebSphere Business Glossary, IBM Information Server Web console, the domain components of the products that you install, and the information center.

WebSphere Metadata Server

The metadata server manages metadata services, shared services, and repository services for the suite components that are installed in a particular domain. The metadata server is hosted by an instance of IBM WebSphere Application Server. WebSphere Application Server is included with the suite. You can choose to use an existing instance of a supported version of the application server. If you use an existing instance of the application server, you must create a profile before you install the suite.

WebSphere Business Glossary

You can use WebSphere Business Glossary to manage, browse, and customize the presentation of metadata that is stored in the metadata repository of WebSphere Metadata Server.

The metadata repository stores metadata for tables, columns, models, schemas, operations, and other components of the data integration process. WebSphere Business Glossary organizes your important business metadata into business categories that contain business terms. You can use business terms to classify other metadata objects in the metadata repository according to the standards and practices of your business.

IBM Information Server Web console

The Web console consists of administration and reporting tools, and the following clients:

- WebSphere Business Glossary
- WebSphere Information Services Director

The Web console has the following tabs.

Administration tab

Using the Administration tab, IBM Information Server administrators perform the following tasks:

- Create users and authorize them to use components and features of IBM Information Server
- Set up an authentication scheme for users of WebSphere DataStage and WebSphere QualityStage.
- Create views of logged events
- Create views of scheduled tasks

Using the Administration tab, authorized users can access logging and scheduling views.

Glossary tab

From the glossary tab, you can work with the WebSphere Business Glossary. Use the Glossary tab to browse the categories, terms, and related metadata in the metadata repository. You can also add notes and submit feedback. Administrators can use categories, terms, and stewards to organize and manage metadata content. If you choose to install WebSphere Business Glossary, it is installed with the domain.

Information Services Catalog tab

The Information Services Catalog tab contains information about deployed information services. If you have the Information Services Catalog

Manager role, the Information Services Director Administrator role, or the Information Services User role, you can browse the services categories tree in the catalog and search categorized, uncategorized deployed services, or all services by custom attributes and design-time metadata. Information about services, such their properties, binding information, custom attributes, and operations, can also be accessed.

If you have the Information Services Catalog Manager role or the Information Services Director role, you can create a hierarchy of services categories and associate deployed services with one or more branches of the hierarchy. All properties of the services categories and some properties of the deployed services can also be modified. You can also create custom attributes, which are user-defined properties that can be associated with all of the deployed services or categories. The values of custom attributes can be modified for each individual service and category in the catalog.

Reporting tab

The type of reports that are available to you vary according to the component that you are working in. Each report corresponds to a task and displays information about the task. Reports can also display general information about the project that you are working in. Reports can display information in a standard template form or in a graphical form such as a diagram or graph. The standard template displays information about the structure of your data, such as the types of data values that make up your data. Diagrams display general information such as the percentage of columns that share a common domain or how much of your data source is analyzed.

Using the Reporting tab, administrators and authorized users can create the following types of reports.

Logging

You can create reports about runtime events that are logged in the metadata repository. The logging reports are organized by the following categories: Severity level, Logging category, and Time frame.

Scheduling

You can create reports of scheduled tasks. The scheduling reports are organized by or describe the following activities: completed tasks, product modules, recurrent tasks, specific users, terminated tasks, time frame, and views of scheduled tasks.

Security

You can create lists of users for each product module in the suite.

Engine

The engine layer consists of the engine components of the products that you install, such as IBM WebSphere DataStage and IBM WebSphere QualityStage.

The engine layer also includes IBM WebSphere Federation Server. The domain layer must be installed on the same computer or on a different computer in the same network before you install the engine layer.

WebSphere DataStage and WebSphere QualityStage

The following engine components are installed with the WebSphere DataStage server.

Engine components

WebSphere DataStage and WebSphere QualityStage use a server component that is installed in the engine layer. The WebSphere DataStage and WebSphere QualityStage server contains runtime components and the engines that run WebSphere DataStage and WebSphere QualityStage jobs to extract, transform, load, and standardize data. The WebSphere DataStage and WebSphere QualityStage server must be installed on the same computer as the metadata server or on a different computer in the same network. WebSphere DataStage and WebSphere QualityStage provide administration tools to manage the server components.

The following components are installed with the WebSphere DataStage and WebSphere QualityStage server:

WebSphere DataStage and WebSphere QualityStage server engine

The server engine runs WebSphere DataStage server jobs and performs some tasks for parallel jobs and WebSphere QualityStage jobs.

Parallel engine

The parallel engine runs parallel jobs and WebSphere QualityStage jobs.

WebSphere Metadata Server agent

The WebSphere Metadata Server agent is a Java™ process that runs in the background on each computer that hosts WebSphere DataStage. When a shared service that is running on the WebSphere Metadata Server receives a service request that requires processing by an IBM Information Server component, the WebSphere Metadata Server agent receives and conveys the request.

On Windows, the WebSphere Metadata Server agent runs as a service. On Linux and UNIX, the WebSphere Metadata Server agent runs as a daemon. The service and daemon name is ASB Agent.

ODBC drivers

A set of ODBC drivers that work with WebSphere DataStage are installed with the engine components.

Resource Tracker

The Resource Tracker for parallel jobs is installed with the engine components for WebSphere DataStage and WebSphere QualityStage on Linux and UNIX systems. The Resource Tracker logs the processor, memory, and I/O usage on each computer that runs parallel jobs. You must do a custom installation to install the Resource Tracker.

WebSphere QualityStage migration utility

The WebSphere QualityStage migration utility is a command line tool that you can use to migrate WebSphere QualityStage 7.0 or 7.5 jobs to WebSphere QualityStage 8.

Component installer

The component installer is a command line tool that you can use to add or update WebSphere DataStage components, such as plug-ins and stages.

WebSphere Federation Server

WebSphere Federation Server consists of a DB2 instance that operates as a federated server, the federated database, one or more data sources, and clients that access the database and data sources. With a federated system, you can send

distributed requests to multiple data sources within a single SQL statement. For example, you can join data that is located in a DB2 table, an Oracle table, and an XML tagged file in a single SQL statement.

Data sources are accessed with the relational and nonrelational wrappers that are installed with WebSphere Federation Server.

Federated systems

A *federated system* is a special type of distributed database management system (DBMS). A federated system consists of a DB2 instance that operates as a federated server, a database that acts as the federated database, one or more data sources, and clients (users and applications) that access the database and data sources.

With a federated system, you can send distributed requests to multiple data sources within a single SQL statement. For example, you can join data that is located in a DB2 table, an Oracle table, and an XML tagged file in a single SQL statement. The following figure shows the components of a federated system and a sample of the data sources you can access.




Figure 2. The components of a federated system

The power of a federated system is in its ability to:

- Correlate data from local tables and remote data sources, as if all the data is stored locally in the federated database
- Update data in relational data sources, as if the data is stored in the federated database

- Move data to and from relational data sources
- Take advantage of the data source processing strengths, by sending requests to the data sources for processing
- Compensate for SQL limitations at the data source by processing parts of a distributed request at the federated server

Metadata repository

The metadata repository contains the metadata for IBM Information Server. The metadata server, hosted by the application server, connects to the metadata repository to store and retrieve data.

The metadata repository can be a DB2 Enterprise Server Edition database, an Oracle database, or a Microsoft® SQL Server database.

If DB2 is the type of database that you want to use for the metadata repository, you can install it with the suite installation wizard, or you can specify an existing instance. If Oracle or Microsoft SQL Server is used, you must run scripts to configure the metadata repository before you install the suite.

The metadata repository must be accessible from the computer where the domain is installed.

Documentation

The documentation for the product modules and components that you install is available in several formats, including PDF, online, and information center.

IBM Information Server Information Center

The IBM Information Server Information Center contains the help for the product components that you install and a subset of the product documentation that is also found in the PDF documents. The information center is installed automatically when you install the domain.

DB2 Information Center

The DB2 Information Center CD contains a subset of the documentation for DB2 and WebSphere Federation Server. You can update the information center to install a complete set of documentation and any updates that have occurred since the CD was created.

The information center is also available at publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp

WebSphere Application Server Information Center

The information center for the application server contains information about configuring and administering the application server that hosts the metadata server. The documentation for the application server is at publib.boulder.ibm.com/infocenter/wasinfo/v6r0/index.jsp?topic=/com.ibm.websphere.base.doc/info/welcome_base.html.

PDF documentation for the suite

By default, when you install the documentation, the PDF documentation for the suite is installed in \IBM\InformationServer\Documentation on Windows and /opt/IBM/InformationServer/Documentation on Linux and UNIX.

You can also find a complete set of the PDF documentation for the suite on the IBM Information Server PDF CD that is included with the installation software.

Accessing documentation

You can install and access the documentation for the suite in several different formats.

About this task

The suite documentation is available in the following formats and locations.

- PDF documents installed locally in the documentation layer
- Information centers installed in the domain layer, locally, or on an intranet server
- Online from the IBM Web site

The DB2 Information Center contains documentation for DB2 and federated systems. You must access it online from the IBM Web site.

Procedure

To access the documentation for the suite, use the following methods.

- To access the PDF documentation for the suite, click **Start** → **IBM Information Server** → **Documentation**. A window opens that describes each document and provides links to open the documents.
- To access help for suite clients and a subset of the suite PDF documentation in online form, access the help from any of the suite clients.
- To access the information center for IBM Information Server, go to publib.boulder.ibm.com/infocenter/iisinfsv/v8r0/index.jsp.
- To access the DB2 Information Center:
 - If you install DB2 or WebSphere Federation Server by using the IBM Information Server installation program, you do not have a local instance of the DB2 information center. Use the DB2 information center on the IBM Web site. Because the information center includes contextual help topics, you need an information center for contextual help to work. In order to configure your installation of DB2 or WebSphere Federation Server to use the information center on the IBM Web site, follow the instructions provided at publib.boulder.ibm.com/infocenter/db2luw/v9/topic/com.ibm.db2.udb.uprun.doc/doc/c0011569.htm.
If you do not have Internet access to use the IBM Web site for your information center, contact your IBM representative to obtain a standalone installation of DB2 or WebSphere Federation Server.
 - If you install DB2 or WebSphere Federation Server separately by using the DB2 or WebSphere Federation Server installation programs, you can use a local information center. Refer to the DB2 or WebSphere Federation Server documentation for information about installing and starting the DB2 information center, which is provided at publib.boulder.ibm.com/infocenter/db2luw/v9/topic/com.ibm.db2.udb.doc/doc/c0008880htm.

- To access documentation for WebSphere Application Server, go to publib.boulder.ibm.com/infocenter/wasinfo/v6r0/index.jsp?topic=/com.ibm.websphere.base.doc/info/welcome_base.html.

ODBC drivers

ODBC drivers for IBM WebSphere DataStage and IBM WebSphere QualityStage are installed with the clients. ODBC drivers for IBM WebSphere Federation Server Version 9.1 must be installed separately.

The DataDirect Technologies Connect ODBC drivers for WebSphere DataStage and WebSphere QualityStage are installed in the ODBCDrivers subdirectory of your IBM Information Server installation. The default path is InformationServer\ODBCDrivers. Documentation for configuring the ODBC drivers is installed in the InformationServer\ODBCDrivers\Documentation directory by default.

The ODBC drivers for WebSphere Federation Server version 9.1 are available from DataDirect Technologies. If you have WebSphere Federation Server Version 9.1 Fix Pack 1, you can use the WebSphere Information Integration ODBC driver.

Chapter 2. Planning to install IBM Information Server

Planning for migration to IBM Information Server

Before you install the IBM Information Server suite, you might need to complete some migration tasks.

Refer to *Migrating to IBM Information Server Version 8* for assistance.

Licensing requirements for IBM Information Server

The product license controls which product modules you can install when you install the suite, the features that are available in the clients, and whether a client can connect to a domain.

When you install the metadata server, you must provide a license file. You can obtain a license file through the Passport Advantage® program. The license must be accessible from the system where you are installing the suite. When you install additional components into an existing domain, the installation wizard reads the license file from the metadata server that is associated with the domain. There is one license file for a domain.

The features that are available to IBM WebSphere DataStage and IBM WebSphere QualityStage clients are controlled by the license for the domain that you connect to when you start the client. You can connect to a domain with IBM WebSphere DataStage and IBM WebSphere QualityStage clients only if the client is licensed for that domain.

If you change your license after you install the suite, you can update your license file in the Web console.

Required users and groups for installing IBM Information Server (Windows)

When you install the suite, certain users and groups are required to complete the installation.

The following table shows the users and groups that are created by the installer when it installs the suite. You can create these users and groups in advance.

Table 2. Users and groups

Layer	User	Sample user name	Group name
Metadata repository	Instance owner	db2admin	Administrators DB2ADMNS
Metadata repository	Repository database owner	xmeta	Not applicable
Domain	Application server ID	wasadmin	Not applicable
Metadata repository	Information analysis database owner	iauser	Not applicable
Domain	IBM Information Server administrator	admin	Not applicable

Metadata repository owner

The metadata repository instance owner is required only if the metadata repository is a DB2 database.

Application server administrator

The application server administrator ID is the ID that you use the first time that you log in to the administrative console for WebSphere Application Server. You can create other users with administrator authority, but the application server ID is authenticated with the user registry when the server starts. If authentication fails, the server does not start. If you use a local operating system registry, the application server administrator ID must be a local operating system user ID. Choose an ID and password that do not expire or change often. See the documentation for the application server for information about changing the application server ID and password.

IBM Information Server administrator

The IBM Information Server administrator ID is used to log in to the Web console for the first time.

User registry types

When you specify the ID and password for the application server ID and the IBM Information Server administrator ID in the installation wizard, you can specify whether you want to use an internal user registry or the local operating system user registry. The installation wizard creates users. If you choose to use an internal user registry, you can create additional suite users in the Web console. If you choose to use a local operating system user registry, you must create additional suite users in the external registry.

You can configure any LDAP-compliant user registry for the IBM Information Server suite. See the documentation for configuring LDAP for WebSphere Application Server.

Required users and groups for installing IBM Information Server (Linux, UNIX)

When you install the suite, certain users and groups are required to complete the installation. The suite installation wizard creates the users and groups when it installs the suite except on systems where NIS, NIS+, or LDAP is used. If NIS, NIS+, or LDAP is used, users and groups must be created on the system before you install the suite.

The following table shows the users and groups that are created by the installer when it installs the suite. You can create these users and groups in advance.

Table 3. Users and groups

Layer	Required user	Sample user name	Group name
Metadata repository	Metadata repository database owner	xmeta	Not applicable
Metadata repository	Instance owner	db2inst1	db2iadm1

Table 3. Users and groups (continued)

Layer	Required user	Sample user name	Group name
Metadata repository	Fenced user	db2fenc1	db2fadm1
Metadata repository	DB2 Administration Server (DAS) user	dasusr1	dasadm1
Domain	Application server ID	wasadmin	Not applicable
Metadata repository	Information analysis database owner	iauser	Not applicable
Engine	WebSphere DataStage administrator	dsadm	dstage
Domain	IBM Information Server administrator	admin	Not applicable

Application server administrator

The application server administrator ID is the ID that you use the first time that you log in to the administrative console for WebSphere Application Server. You can create other users with administrator authority, but the application server ID is authenticated with the user registry when the server starts. If authentication fails, the server does not start. If you use a local operating system registry, the application server administrator ID must be a local operating system user ID. Choose an ID and password that do not expire or change often. See the documentation for the application server for information about changing the application server ID and password.

IBM Information Server administrator

The IBM Information Server administrator ID is used to log in to the Web console for the first time.

Metadata repository owner

The metadata repository instance owner is required only if the metadata repository is a DB2 database.

User registry types

When you specify the ID and password for the application server ID and the IBM Information Server administrator ID in the installation wizard, you can specify whether you want to use an internal user registry or the local operating system user registry. The installation wizard creates users. If you choose to use an internal user registry, you can create additional suite users in the Web console. If you choose to use a local operating system user registry, you must create additional suite users in the external registry.

You can configure any LDAP-compliant user registry for the IBM Information Server suite. See the documentation for configuring LDAP for WebSphere Application Server.

WebSphere DataStage server administrator

WebSphere DataStage is administered on a Linux or UNIX system by a non-root user, dsadm. The dsadm user is created automatically when you install the engine,

unless NIS is active on your system. If NIS is active on the system, you must create this user before you install the engine. All WebSphere DataStage users must belong to the same Linux or UNIX group, which is the dsadm user's primary group. The group is named dstage by default.

User registries

When you install the IBM Information Server suite, you can choose the user registry that you want to use or you can change it after the installation is complete.

The IBM Information Server directory is the internal user directory for the metadata server where users, groups, roles, and credential mappings are stored. The IBM WebSphere Application Server user registry is the user registry for the application server that hosts the metadata server.

The IBM WebSphere Application Server user registry can be configured to access the local operating system user registry or external user registries.

The IBM Information Server directory and the IBM WebSphere Application Server user registry must always be configured to use the same user registry, so that they both access the same set of users and groups. If you choose to use the local operating system user registry, which is accessed through the IBM WebSphere Application Server user registry service, the IBM Information Server directory must be configured when the suite is installed or just after the installation is complete.

Methods of installation

You can install the suite in graphical mode, console mode, or silent mode. Consider which installation method works best for you before you install the suite.

Graphical mode

Your computer must be capable of rendering a graphical user interface to run the graphical installation wizard. The wizard displays a series of screens where you can specify the installation options that you want to use. When all of the options have been specified, you can review a summary of the options that you selected and begin the installation.

Console mode

A console mode installation is done from the command line. The installation program displays text prompts in the console window to gather the installation options that you want to use. When all of the options have been specified, you can review a summary of the options that you selected and begin the installation. This installation method does not require your computer to render a graphical user interface.

Silent mode

You might want to install the suite in silent mode when you are performing identical installations on multiple computers. A silent installation is started from the command line, but does not require you to specify the installation options. Instead, the installation options are read from a response file. You can create a response file by using the graphical installation wizard or the console method. A response file can be created without installing any software or during an installation.

Order of installation

The client and documentation layers can be installed in any order. If you install the metadata repository, domain, and engine layers on the same computer at the same time, the installation wizard controls the order of installation.

However, if you install the layers on different computers or at different times, you must install them in the following order:

1. Metadata repository layer
2. Domain layer
3. Engine layer

If you install the metadata repository, domain, and engine on different computers, the computers must have the same operating system.

Installation planning worksheet

Before you install the suite, create an installation plan. An installation plan describes system requirements, the components to be installed, the installation configuration, and the configuration tasks that must be completed after the installation. Use this worksheet to create and record your plan.

- ___ 1. If you have previous or unsupported versions of DataStage or QualityStage installed, review the migration process and complete the pre-installation migration tasks.
- ___ 2. If you plan to use an existing installation of WebSphere Application Server to host the metadata server, ensure that it is a supported version.
- ___ 3. If you plan to use an existing installation of DB2, ensure that it is a supported version.
- ___ 4. Decide what type of installation you want to do?
 - ___ Graphical installation wizard
 - ___ Console
 - ___ Silent
- ___ 5. Determine which products and components you want to install and the layers that they are installed in.
- ___ 6. What operating systems are you installing the suite on?
 - ___ Windows XP (clients only)
 - ___ Vista (clients only)
 - ___ Windows Server 2003, with these layers:

 - ___ AIX®, with these layers: _____
 - ___ Linux, with these layers: _____
- ___ 7. Ensure that the computers that you are installing on meet the system requirements.
- ___ 8. Ensure that the prerequisite and corequisite software is installed.
 - ___ C++ compiler
- ___ 9. Determine which database server will host the metadata repository.
 - ___ DB2
 - ___ Oracle. Path _____
 - ___ Microsoft SQL Server.
- ___ 10. If you are using the DB2 database system for the metadata repository, and it is already installed, identify the path:

-
- ___ 11. If you are using the DB2 database system for the metadata repository and it is not installed, identify the path where do you want it installed:

 - ___ 12. If WebSphere Application Server is already installed, identify the path:

 - ___ 13. If WebSphere Application Server is not installed, identify the path where do you want it installed:

 - ___ 14. What kind of user registry do you want to use?
 - ___ Internal. If you use the internal user registry, user credentials are stored in the metadata repository.
 - ___ Local operating system. If you use the local operating system user registry, user credentials are verified with the local operating system. If you choose this option, you must create the application server administrator and the suite administrator users on the system where you install the domain.

System requirements for IBM Information Server

Before you install IBM Information Server, ensure that your computer meets the hardware and software requirements for the product modules and components that you want to install.

For information about the system requirements for IBM Information Server, go to www.ibm.com/support/docview/wss?uid+swg270089.

WebSphere Federation Server system requirements

IBM WebSphere Federation Server is an option when you select the engine layer. It can be installed on a different computer from the engine.

For information about system requirements, go to www.ibm.com/software/data/integration/federation_server/requirements.html.

Chapter 3. Preparing your system for installation

Before you install the suite, you must prepare your computers for installation. The preparation that is required depends on the operating system of the computer where you are installing the suite.

Preparing your system for installation (Windows)

Before you install IBM Information Server, you need to prepare your system by verifying system requirements. The steps for preparing your system depend on the operating system and which layers and product modules that you will install.

- You must have administrator access to the installation system to complete this task.
- The clients, metadata repository, and engine must be installed on a computer that is accessible by WebSphere Metadata Server.
- If you install the metadata repository, domain, or engine on different computers, all the computers must run the same operating system.
- If you use SQL Server 2005 as the metadata repository, the user name and database schema must be the same.

Restrictions

You cannot install the metadata repository or WebSphere Federation Server on systems where only DB2 clients are installed.

User input to the installation process is limited to ASCII because the response file for the DB2 installation must be ASCII.

The combined path length of the temporary directory that is used by the installer and the path of files that are extracted into it during installation cannot exceed 256 characters. This limit can be reached on Windows systems if the default temporary directory is defined as C:\Documents and Settings\username\Local Settings\Temp. You can create a temporary directory such as C:\tmp and launch the installer with a command, for example. `install.exe -is:tempdir C:/tmp`.

Avoid installing IBM Information Server on systems that have versions of the MKS toolkit installed. The engine layer of IBM Information Server installs a version of the MKS toolkit.

The installation program queries the local system for its host name and uses it for the configuration of any server applications you choose. If your machine name does not match the name in your local DNS server for your IP address, some of the applications may have name resolution problems that can result in unexpected behavior. For example, accessing the administration console with your DNS resolvable address may work, but accessing it using your machine name may result in the inability of the console to initialize. Any system intended to run as a server needs to have consistent naming and IP resolution to function correctly.

If you are installing the DataStage server on an operating system with a double-byte or multi-byte language such as Simplified Chinese, Traditional Chinese, Japanese or Korean, you must enable NLS during installation. On the National Language Support panel for the WebSphere DataStage server, you must

check the box labeled Install NLS for WebSphere DataStage server. If you do not check the box, your DataStage clients may not be able to communicate with the DataStage server.

About this task

You must be the local administrator on the system on which you are installing IBM Information Server. You must not run the installer as a domain administrator.

You cannot install DB2 9.1 on a system on which Microsoft SQL Server 2005 is installed. You can either install DB2 9.1 on a different system or use Microsoft SQL Server 2005 for the metadata repository. The specific Microsoft library with which DB2 9.1 is not compatible is msvcr80.dll. This library may also be included with other software packages. In this case, the solution is to install DB2 9.1 on a different system.

DB2 9.1 might not install correctly if the Fast User Switching feature of Windows is enabled. If you are installing the version of DB2 that is bundled with IBM Information Server, you must turn off Fast User Switching before starting the installation. You can do this from **Control Panel** → **User Accounts** → **Change the way users log on or off**.

If you install the layers at different times or on different computers, ensure that you follow the order of installation that is indicated by the following steps. If you install the layers at the same time and on the same computer, the installation wizard controls the order of installation.

Procedure

To prepare for installation:

1. Log on to the system as a root administrator.
2. Ensure that your system meets the system requirements for the product modules and components that you want to install.
3. Turn off any firewall software that is installed on the system.
4. Optional: Turn off your antivirus software.
5. If you have an earlier version of IBM Information Server or DataStage installed on your computer, determine whether you need to remove or upgrade the earlier version.

Option	Description
Product or component	Action
IBM Information Server	You must remove IBM Information Server version 8.0.
DataStage	• If you have DataStage 7.0 or later, you can upgrade it. See <i>Migrating to IBM Information Server Version 8</i> for instructions. • If you have a version of DataStage earlier than 7.0, you must upgrade to version 7.5.2 before you can upgrade to version 8.0. See <i>Migrating to IBM Information Server Version 8</i> for instructions.

6. If necessary, create the users and groups for the product modules and components that you want to install. If user names and passwords do not exist, you might need to create them.
7. Prepare the computers for the metadata repository.
8. Prepare the computers for the domain.
9. Prepare the computers for the engine.

After you prepare your systems for installation, you can install IBM Information Server.

Preparing to install the metadata repository

Before you install the suite, you need to prepare your system for installing the metadata repository. The steps that are required to install the metadata repository differ depending on the type of database you use for the metadata repository.

If you use an existing instance of a DB2 database for your metadata repository, ensure that it is installed on a computer that is accessible from the computer where you are installing the metadata repository.

If you use Oracle or Microsoft SQL Server as the metadata repository, you must create and configure the database before you install the domain and engine.

If you use DB2 as your metadata repository, the installation wizard detects whether DB2 is already installed. If DB2 is not installed or if it is installed but you wish to install a new copy, the installation wizard can automatically install a copy of DB2 and create the database to be used as the metadata repository during suite installation. Alternatively, if DB2 is installed, you can manually create and configure a DB2 database to be used as the metadata repository before you install the suite.

To prepare Windows systems for installing the metadata repository:

1. Log on to the system as a root administrator.
2. Ensure that your system meets the system requirements for the product modules and components that you want to install.
3. Turn off any firewall software that is installed on the system.
4. Optional: Turn off your antivirus software.
5. If you plan to manually create a DB2 database to use as the metadata repository:
 - a. Open a command window, and change to the DatabaseSupport\DB2_9 directory on the installation medium.
 - b. Review the settings in the DB2 SQL script file `xmeta_nls_db_creation_config.sql` to ensure that they are adequate for the type of installation that you plan. Review the database path, the location of the tablespace containers, the location for the log files, and the backup and recovery strategy as defined by the LOGRETAIN and USEREXIT1 parameters.
 - c. If you want to change the default names of the database and of the user, edit the `xmeta_nls_db_creation_config.sql` file. The default name of the database is XMETANLS, and the default name of the user is XMETADB.
 - d. Run the `xmeta_nls_db_creation_config.bat` file to create the metadata repository and user.
 - e. Record the following information:

- Database owner and password
- Database name for the metadata repository.
- Database instance
- Database location

Follow the order of installation if you are installing dependent layers on different systems. If you install dependent layers on the same system, the installation wizard manages the order of installation.

After you prepare your system for the metadata repository installation, prepare your system for the domain installation.

Creating the metadata repository database for SQL Server 2005

You can create a Microsoft SQL Server database on the Windows platform to use as the metadata repository. You should create a dedicated database for this purpose.

Configuring the metadata repository on Microsoft SQL Server 2005:

If you want to use Microsoft SQL Server as your metadata repository, you must configure it before you install the suite.

Microsoft SQL Server 2005 must be installed on a computer that is accessible by the domain and the engine.

To create and configure a SQL Server 2005 database to use as the metadata repository on the Windows computer:

1. Open a Windows command prompt and change to the directory where the `xmeta_db_nls_creation_config.bat` file is located. This file can be located in the Database Support\SQLServer2005 directory on the installation media.
2. Review the settings in the `xmeta_db_nls_creation_config.sql` file to ensure that they are adequate for the type of installation being planned. The `xmeta_db_nls_creation_config.sql` file in the Database Support\SQLServer2005 directory is an SQL Server 2005 script file that you can use to create the metadata repository, database, and user. IBM recommends that you place the data and log files on separate physical disks and these disks be striped. By default, the metadata repository that is created is named `xmeta`, and the user that is created is named `xmeta` with a password of `xmeta`. These values can be changed, but the default schema name must be the same as the user name.
3. Create the SQL Server database by running the following command:
`xmeta_db_nls_creation_config.bat` This in turn executes the `xmeta_db_nls_creation_config.sql` file. This command also creates a default login ID and password

After you configure the metadata repository, you can install the suite.

Parameters of the metadata repository in SQL Server:

The parameters used to create the metadata repository in SQL Server may be modified to either meet organizational requirements or for the type of installation being planned. However, the **Authentication mode** parameter should not be changed.

Creating the metadata repository for Oracle 10.2g

You can create an Oracle database instance as the metadata repository by using the Oracle Database Configuration Assistant (DBCA). Configure it before you install the suite.

The Oracle DBCA is available with the Oracle 10.2g installation. After you configure the metadata repository and complete the steps for preparing your computers for the domain and engine, you can install the suite. Advanced users can use methods other than the Database Configuration Assistant to create and configure the metadata repository. One alternative is to use the SQL Plus utility to manually create the database and configure the settings. Only the parameters related to performance tuning or file locations should be modified, other parameters should not be changed.

Installing the Oracle database instance on a Windows computer:

You can use an Oracle database instance as the metadata repository on a Windows computer. You should create a dedicated database for this purpose.

You must install the Oracle 10.2g software before you create the Oracle database instance. You must have administrator privileges with the ability to create databases.

To create an Oracle database instance on a Windows computer:

1. Copy the DatabaseSupport\Oracle10g\xmeta.dbt database template and the DatabaseSupport\Oracle10g\create_xmeta_user.sql file from the installation media to the (ORACLE_HOME)\assistants\dbca\templates directory. These files ensure that the Oracle Database Configuration Assistant recognizes the xmeta.dbt database template.
2. Click **Start** → **All Programs** → **Oracle 102.g application** → **Configuration and Migration Tools** → **Database Configuration Assistant**.
3. From the welcome screen, click **Next**.
4. From the Operations window, select **Create a database** and click **Next**.
5. In the Database Templates window select **xmeta** and click **Next**.
6. In the Database Identification window, specify values in the **Global Database Name** field and the **SID**. If default values are provided then users don't need to enter values in this step. Note the values you supply, because you will use them when you install the Information Server. Click **Next**.
7. In the Management Options window, click **Next**.
8. In the Database Credentials window, select the **Use the Same Password for All Accounts** option and enter a password.
9. In the Storage Options window, specify the type of storage for the metadata repository database. The default storage mechanism is File System. Click **Next**.
10. In the Database File Location window, specify a database file location for the database files. The default option for an xmeta database is **Use Oracle-Managed Files**, which reduces the amount of time required to administer the database. If you select **Use Oracle-Managed Files** or **Use Common Location for All Database Files**, specify a valid location in the **Database Area** field or **Database Files Location** field. Do not select **Use Database File Locations from Template** because the file locations in the template might not be valid in your environment.
11. In the Recovery Configuration window, specify the recovery configuration information. If you select **Enable Archiving**, click the **Edit Archive Mode**

Parameters button to open the **Edit Archive Mode Parameters** window. In the **Archive Log Destination**, enter the disk location where the database logs are to be archived. The database administrator must ensure that these archived logs do not consume all the free disk space and comply with Oracle recommendations for archiving, backup, and recovery of logs.

12. In the Database Content, click **Next**.
13. In the Initialization Parameters window, specify the initialization parameters.
 - a. On the Memory tab, specify the percentage of the total physical memory allocated for the metadata repository database in the **Percentage** option. The default is 25% , which is the recommended setting for a machine that is not a dedicated database server and runs other applications. The default value can be changed, but you must be aware that any decrease in memory allocated to the database can adversely affect the performance, and any increase in this parameter must ensure that there is sufficient real memory allocated in the computer.
 - b. On the Character Sets tab, specify appropriate values for **Database Character Set**, **National Character Set**, **Default Language**, and **Default Date Format** options. Click **All Initialization Parameters** and examine the configuration parameters. You can change the configuration parameters to comply with organizational requirements.
14. In the Database Storage window, specify the storage parameters for the database. By default, the tables and data in the metadata repository are in the Users tablespace. The default storage parameters in the template place all the data, redo logs, and control files in the directory that you specified in the Database File Location window. You can change this location to achieve a uniform distribution of input and output activity across your disk subsystem, and improve performance. Follow guidelines provided by Oracle or the guidelines within your organization for the placement of redo log files.
15. In the Creation Options window, click **Create Database**. To save the information as scripts, click **Generate Database Creation Scripts**.
16. Click **Finish**.
17. In the Confirmation window, select **Save as an HTML file** and click **Save**. This HTML file might be required by Oracle support if you have issues with Oracle 10.2g.
18. Click **OK** and create the metadata repository in Oracle.

The Database Configuration Assistant will now create the metadata repository database.

Related reference

“Database instance setup parameters” on page 27

The Oracle Database Administration Assistant has default values for the initialization parameters that you use with the Information Server.

“Database instance parameters” on page 27

The Oracle Database Administration Assistant has default values for the configuration parameters that you use with the Information Server.

Creating the Oracle metadata repository user:

The xmeta user is used by the IBM Information Server suite of applications to access the metadata repository.

To create the Oracle metadata repository user:

1. Change to the directory where the create_xmeta_user.sql file is located by using the command: `cd {ORACLE_HOME}\assistants\dbca\templates`.
2. At the command prompt enter the command: `sqlplus SYS/systempassword as SYSDBA`.
3. Run the create_xmeta_user.sql script from the sqlplus command line to create the xmeta user and to assign privileges.

This procedure creates the xmeta user with a default password of xmeta. You can change the password by using the Oracle Database Control application.

If you do not want to use xmeta as the metadata repository user, edit the create_xmeta_user.sql script before you run the file.

Database instance setup parameters:

The Oracle Database Administration Assistant has default values for the initialization parameters that you use with the Information Server.

Table 4. Database identification

Name	Value
Global Database Name	xmeta
SID	xmeta

Table 5. Initialization parameters

Name	Value
Database Character Set	AL32UTF8

Table 6. Create user

Name	Value
Name	xmeta
Enter Password	xmeta

Database instance parameters:

The Oracle Database Administration Assistant has default values for the configuration parameters that you use with the Information Server.

Table 7. Common options

Name	Value
Oracle JVM	True
Oracle Data Mining	False
Oracle Text	False
Oracle XML DB	False
Oracle Intermedia	False
Oracle OLAP	False
Oracle Spatial	False
Oracle Ultra Search	False
Oracle Label Security	False

Table 7. Common options (continued)

Name	Value
Sample Schemas	False
Enterprise Manager Repository	True
Oracle Database Extensions for .NET	False

Table 8. Initialization parameters

Name	Value
Audit_file_dest	{ORACLE_BASE}\admin\{DB_UNIQUE_NAME}\adump
background_dump_dest	{ORACLE_BASE}\admin\{DB_UNIQUE_NAME}\bdump
compatible	10.2.0.1.0
Core_dump_dest	{ORACLE_BASE}\admin\{DB_UNIQUE_NAME}\cdump
db_block_size	8192
db_create_file_dest	e:\oradata
db_domain	
db_file_multiblock_read_count	16
db_name	xmeta
db_recovery_file_dest	{ORACLE_BASE}\flash_recovery_area
db_recovery_file_dest_size	2048MB
job_queue_processes	10
log_archive_dest_1	'LOCATION={ORACLE_BASE}\ArhivedLogs'
log_archive_format	ARC%S_%R.%T
Open_cursors	1500
pga_aggregate_target	117MB
processes	150
remote_login_passwordfile	EXCLUSIVE
sga_target	353MB
Undo_management	AUTO
Undo_tablespace	UNDOTBS1
User_dump_dest	{ORACLE_BASE}\admin\{DB_UNIQUE_NAME}\udump

Table 9. Character sets

Name	Value
Database Character Set	AL32UTF8
National Character Set	AL16UTF16

Table 10. Tablespaces

Name	Type	Extent Management
SYSAUX	PERMANENT	LOCAL

Table 10. Tablespaces (continued)

Name	Type	Extent Management
SYSTEM	PERMANENT	LOCAL
TEMP	TEMPORARY	LOCAL
UNDOTBS1	PERMANENT, UNDO	LOCAL
USERS	PERMANENT	LOCAL

Table 11. Data files

Name	Tablespace	Size(M)
<OMF_SYSAUX_DATAFILE_0>	SYSAUX	240
<OMF_SYSAUX_DATAFILE_0>	SYSTEM	480
<OMF_TEMP_DATAFILE_0>	TEMP	20
<OMF_UNDOTABS1_DATAFILE_0>	UNDOTBS1	25
<OMF_USERS_DATAFILE_0>	USERS	1000

Table 12. Redo log files

Group	Size(K)
1	71200
2	71200
3	71200

Preparing to install the domain (Windows)

Before you install the IBM Information Server domain, you need to prepare your system for the product modules and components that you want to install.

Before you begin

The clients, metadata repository, and engine must be installed on a computer that is accessible by WebSphere Metadata Server.

When you create the user names and passwords for the application server and the IBM Information Server administrator, you choose to create them in an internal user registry or in the local operating system user registry.

Procedure

To prepare your system to install the domain:

1. If you install the metadata repository on a different system, you must supply the following information about the metadata repository during the domain installation:
 - Database type
 - Host name
 - Database port
 - Database name
 - User name and password
2. If you have an earlier version of IBM Information Server or DataStage installed on your computer, determine whether you need to remove or upgrade the

earlier version.

Option	Description
Product or component	Action
IBM Information Server	You must remove IBM Information Server version 8.0.
DataStage	• If you have DataStage 7.0 or later, you can upgrade it. See <i>Migrating to IBM Information Server Version 8</i> for instructions. • If you have a version of DataStage earlier than 7.0, you must upgrade to version 7.5.2 before you can upgrade to version 8.0. See <i>Migrating to IBM Information Server Version 8</i> for instructions.

3. Prepare an instance of WebSphere Application Server:

To...	Do this...
Use an existing instance of WebSphere Application Server	Create a dedicated profile by using the WebSphere Application Server Profile Creation wizard.
Create a new installation of WebSphere Application Server	Complete the steps for preparing your operating system for WebSphere Application Server as described in the WebSphere Application Server information center.

After you prepare your system or systems for the domain installation, prepare your system for the engine installation.

Creating a server profile for existing instances of WebSphere Application Server

To use an existing instance of WebSphere Application Server, you need to create a server profile before you begin the installation.

To create a server profile, use the Profile creation wizard of WebSphere Application Server.

Preparing to install the engine (Windows)

Before you install IBM Information Server, you need to prepare your system by verifying system requirements. The steps for preparing your system depend on the operating system and which layers and product modules that you will install.

Before you begin

The clients, metadata repository, and engine must be installed on a computer that is accessible by WebSphere Metadata Server.

Procedure

To prepare to install the engine:

1. If the domain is installed on another system, record the following information about WebSphere Metadata Server:

- Host name
 - TCP/IP port. This is the HTTP port of the WebSphere Application Server profile, which is 9080 by default.
2. If you want to access federated data sources, ensure that you meet the prerequisites for installing relational and nonrelational wrappers before you install WebSphere Federation Server.
 3. Unless you are installing WebSphere Metadata Server at the same time that you are installing the engine, start WebSphere Metadata Server.

After you prepare your system or systems the metadata repository, domain, and engine installations, you can install IBM Information Server.

Creating users and groups (Windows)

If user names and passwords do not exist, you might need to create them.

About this task

When you create the user names and passwords for the application server and the IBM Information Server administrator, you choose to create them in an internal user registry or in the local operating system user registry.

Procedure

To create users and groups:

1. Click **Start** → **Control Panel**.
2. Double-click **User Accounts**.
3. In the User Accounts window, turn off fast user switching by clicking **Change the way users log on or off**.
4. Click **Create a new account**.
5. Type a name for the new account and click **Next**.
6. Select the **Computer administrator** radio button and click **Create Account**. The account is created.
7. In the User Accounts window, double-click the account that you just created.
8. In the window that opens, click **Change the password**. Follow the instructions in the window to set the password for the account.
9. Repeat the process for each user ID and password that you need to create.

Users and groups by layer (Windows)

The metadata repository layer and the domain layer require users and might require groups.

This table lists the type of user required for each layer, a sample user name, and a group name, if required.

Table 13. Users and groups

Layer	User	Sample user name	Group name
Metadata repository	Instance owner	db2admin	Administrators DB2ADMNS
Metadata repository	Repository database owner	xmeta	Not applicable
Domain	Application server ID	wasadmin	Not applicable

Table 13. Users and groups (continued)

Layer	User	Sample user name	Group name
Metadata repository	Information analysis database owner	iauser	Not applicable
Domain	IBM Information Server administrator	admin	Not applicable

The user ID for the metadata repository instance owner is required only if you use a DB2 database as your metadata repository.

Configuring the analysis database and the analysis engine

If you are installing WebSphere Information Analyzer and you are using DB2 as your analysis database, that database is automatically created by default during installation. If you want to use Microsoft SQL Server or Oracle as your analysis database, you must configure the analysis database before you install the suite. Additionally, if you install the analysis engine on a local computer or remote computer, you must manually create the analysis database DSN.

Configuring Microsoft SQL Server as the analysis database

To use Microsoft SQL Server as your analysis database:

See the instructions in the \Database Support\SQLServer2005\AnalysisDatabase directory. By default, this directory is located in the root directory of the IBM Information Server CD, DVD, or download image.

Configuring Oracle as the analysis database

To use Oracle as your analysis database:

See the instructions in the \Database Support\Oracle10g\AnalysisDatabase directory. By default, this directory is located in the root directory of the IBM Information Server CD, DVD, or download image.

Configuring the analysis engine

If you install the analysis engine on a local computer or a remote computer, you must manually create a DSN for the analysis database on that computer. This DSN must be a system DSN. Additionally, you must create a DSN for each source database or file that you want to profile on this computer.

Related tasks

“Configuring a connection to the analysis database for WebSphere Information Analyzer” on page 120

When you analyze data, the results of your data analysis are stored in the analysis database. The connection to the analysis database must be configured after you install the suite.

Preparing your system for installation (Linux, UNIX)

Before you install IBM Information Server, you need to prepare your system by verifying system requirements and creating users and groups. The steps for preparing your system depend on the operating system and which layers and product modules that you will install.

Restrictions

You cannot install the metadata repository or WebSphere Federation Server on systems where only DB2 clients are installed.

The installation program queries the local system for its host name and uses it for the configuration of any server applications you choose. If your machine name does not match the name in your local DNS server for your IP address, some of the applications may have name resolution problems that can result in unexpected behavior. For example, accessing the administration console with your DNS resolvable address may work, but accessing it using your machine name may result in the inability of the console to initialize. Any system intended to run as a server needs to have consistent naming and IP resolution to function correctly.

If you are installing the DataStage server on an operating system with a double-byte or multi-byte language such as Simplified Chinese, Traditional Chinese, Japanese or Korean, you must enable NLS during installation. On the National Language Support panel for the WebSphere DataStage server, you must check the box labeled Install NLS for WebSphere DataStage server. If you do not check the box, your DataStage clients may not be able to communicate with the DataStage server.

About this task

The following task describes the steps required to prepare Linux and UNIX systems for installing the metadata repository, domain, and engine. If you install the layers at different times or on different computers, ensure that you follow the order of installation that is indicated by the steps. If you install the layers at the same time and on the same computer, the installation wizard controls the order of installation.

The clients, metadata repository, and engine must be installed on a computer that is accessible by WebSphere Metadata Server.

The clients are supported on Windows XP. The clients and documentation require no additional preparation other than that your system must meet the installation requirements for IBM Information Server.

Procedure

To prepare Linux and UNIX systems for installing the suite:

1. Log on to the system as a root user.
2. Ensure that your system meets the system requirements for the product modules and components that you want to install.
3. Turn off any firewall software that is installed on the system.
4. Optional: Turn off your antivirus software.
5. If you have an earlier version of IBM Information Server or DataStage installed on your computer, determine whether you need to remove or upgrade the earlier version.

Option	Description
Product or component	Action

Option	Description
IBM Information Server	- To preserve version 8.0 of IBM Information Server, you must install version 8.0.1 using the multiple server method, which is described in <i>Migrating to IBM Information Server Version 8</i>. - If you do not need to preserve Version 8.0 of IBM Information Server, you must remove it.
WebSphere DataStage	- If you have DataStage 7.0 or later, you can upgrade it. See <i>Migrating to IBM Information Server Version 8</i> for instructions. - If you have DataStage 7.0 or later, you can preserve it by using the multiple server method, which is described in <i>Migrating to IBM Information Server Version 8</i>. - If you have a version of DataStage earlier than 7.0, you must upgrade to version 7.5.2 before you can upgrade to the current version. See <i>Migrating to IBM Information Server Version 8</i> for instructions.

- If user names and passwords do not exist, you might need to create them. If your system uses NIS, NIS+, or other security programs that might restrict the installation wizard's ability to create the credentials for the components that you install, create the users and groups that you need before you install IBM Information Server.
- Create the users and groups that you need for the layers that you are going to install.
- Complete the steps for preparing Linux and UNIX systems for installing the domain layer
- Complete the steps for preparing Linux and UNIX systems for installing the engine layer

After you have prepared your systems for installation, you can install IBM Information Server.

Preparing to install the metadata repository

Before you install the suite, you need to prepare your system for installing the metadata repository. The steps that are required to install the metadata repository differ depending on the type of database that you use for the metadata repository.

If you use an existing instance of a DB2 database for your metadata repository, ensure that it is installed on a computer that is accessible from the computer where you are installing the metadata repository.

If you use Oracle as the metadata repository, you must create and configure the database before you install the domain and engine.

If you use DB2 as your metadata repository, the installation wizard detects whether DB2 is already installed. If DB2 is not installed or if it is installed but you wish to install a new copy, the installation wizard can automatically install a copy of DB2

and create the database to use as the metadata repository during suite installation. Alternatively, if DB2 is installed, you can manually create and configure a DB2 database to be used as the metadata repository before you install the suite.

To prepare your system for installing the metadata repository:

1. Log on to the system as a root administrator.
2. Ensure that your system meets the system requirements for the product modules and components that you want to install.
3. Turn off any firewall software that is installed on the system.
4. Optional: Turn off your antivirus software.
5. If you plan to manually create a DB2 database to use as the metadata repository:
 - a. Open a command window, and change to the DatabaseSupport\DB2_9 directory on the installation medium.
 - b. Review the settings in the DB2 SQL script file `xmeta_nls_db_creation_config.sql` to ensure that they are adequate for the type of installation that you plan. Review the database path, the location of the tablespace containers, the location for the log files, and the backup and recovery strategy as defined by the LOGRETAIN and USEREXIT1 parameters.
 - c. If you want to change the default names of the database and of the user, edit the `xmeta_nls_db_creation_config.sql` file. The default name of the database is XMETANLS, and the default name of the user is XMETADB.
 - d. Run the `xmeta_nls_db_creation_config.sh` to create the metadata repository and user.
 - e. Record the following information:
 - Database owner and password
 - Database name for the metadata repository.
 - Database instance
 - Database location

Follow the order of installation if you are installing dependent layers on different systems. If you install dependent layers on the same system, the installation wizard manages the order of installation.

After you prepare your system for the metadata repository installation, prepare your system for the domain installation.

Creating the metadata repository for Oracle 10.2g

You can create an Oracle database instance as the metadata repository by using the Oracle Database Configuration Assistant (DBCA). Configure it before you install the suite.

The Oracle DBCA is available with the Oracle 10.2g installation. After you configure the metadata repository and complete the steps for preparing your computers for the domain and engine, you can install the suite. Advanced users can use methods other than the Database Configuration Assistant to create and configure the metadata repository. One alternative is to use the SQL Plus utility to manually create the database and configure the settings. Only the parameters related to performance tuning or file locations should be modified, other parameters should not be changed.

Installing the Oracle database instance on a Unix platform:

You can create an Oracle Database instance on a Unix system by using the Oracle Database Assistant.

X-Win32 is required to run the graphical mode installer from a Windows computer.

This procedure uses the Xterminal window from a Windows computer but other terminal emulation software can also be used.

1. At the windows command prompt run the `txterm` command to log in to the Unix system as a member of the administrative group that is authorized to install Oracle software and create and run the database.
2. Change to the `cd $ORACLE_HOME/bin.` directory.
3. Ensure that the `DISPLAY` environment variable is set to `DISPLAY=localhost:10.0`
4. Run the command `./dbca` to start the Oracle Database Configuration Assistant.
5. Follow the instructions in the Oracle Database Configuration Assistant to create an Oracle database instance.

Related tasks

“Installing the Oracle database instance on a Windows computer” on page 25
You can use an Oracle database instance as the metadata repository on a Windows computer. You should create a dedicated database for this purpose.

Creating the Oracle metadata repository user:

The `xmeta` user is used by the IBM Information Server suite of applications to access the metadata repository.

To create the Oracle metadata repository user:

1. Change to the directory where the `create_xmeta_user.sql` file is located by using the command: `cd {ORACLE_HOME}\assistants\dbca\templates.`
2. At the command prompt enter the command: `sqlplus SYS/systempassword as SYSDBA.`
3. Run the `create_xmeta_user.sql` script from the `sqlplus` command line to create the `xmeta` user and to assign privileges.

This procedure creates the `xmeta` user with a default password of `xmeta`. You can change the password by using the Oracle Database Control application.

If you do not want to use `xmeta` as the metadata repository user, edit the `create_xmeta_user.sql` script before you run the file.

Database instance setup parameters:

The Oracle Database Administration Assistant has default values for the initialization parameters that you use with the Information Server.

Table 14. Database identification

Name	Value
Global Database Name	xmeta
SID	xmeta

Table 15. Initialization parameters

Name	Value
Database Character Set	AL32UTF8

Table 16. Create user

Name	Value
Name	xmeta
Enter Password	xmeta

Database instance parameters:

The Oracle Database Administration Assistant has default values for the configuration parameters that you use with the Information Server.

Table 17. Common options

Name	Value
Oracle JVM	True
Oracle Data Mining	False
Oracle Text	False
Oracle XML DB	False
Oracle Intermedia	False
Oracle OLAP	False
Oracle Spatial	False
Oracle Ultra Search	False
Oracle Label Security	False
Sample Schemas	False
Enterprise Manager Repository	True
Oracle Database Extensions for .NET	False

Table 18. Initialization parameters

Name	Value
Audit_file_dest	{ORACLE_BASE}\admin\{DB_UNIQUE_NAME}\adump
background_dump_dest	{ORACLE_BASE}\admin\{DB_UNIQUE_NAME}\bdump
compatible	10.2.0.1.0
Core_dump_dest	{ORACLE_BASE}\admin\{DB_UNIQUE_NAME}\cdump
db_block_size	8192
db_create_file_dest	e:\oradata
db_domain	
db_file_multiblock_read_count	16
db_name	xmeta
db_recovery_file_dest	{ORACLE_BASE}\flash_recovery_area
db_recovery_file_dest_size	2048MB
job_queue_processes	10
log_archive_dest_1	'LOCATION={ORACLE_BASE}\ArhivedLogs'

Table 18. Initialization parameters (continued)

Name	Value
log_archive_format	ARC%S_%R.%T
Open_cursors	1500
pga_aggregate_target	117MB
processes	150
remote_login_passwordfile	EXCLUSIVE
sga_target	353MB
Undo_management	AUTO
Undo_tablespace	UNDOTBS1
User_dump_dest	{ORACLE_BASE}\admin\{DB_UNIQUE_NAME}\udump

Table 19. Character sets

Name	Value
Database Character Set	AL32UTF8
National Character Set	AL16UTF16

Table 20. Tablespaces

Name	Type	Extent Management
SYSAUX	PERMANENT	LOCAL
SYSTEM	PERMANENT	LOCAL
TEMP	TEMPORARY	LOCAL
UNDOTBS1	PERMANENT, UNDO	LOCAL
USERS	PERMANENT	LOCAL

Table 21. Data files

Name	Tablespace	Size(M)
<OMF_SYSAUX_DATAFILE_0>	SYSAUX	240
<OMF_SYSAUX_DATAFILE_0>	SYSTEM	480
<OMF_TEMP_DATAFILE_0>	TEMP	20
<OMF_UNDOTABS1_DATAFILE_0>	UNDOTBS1	25
<OMF_USERS_DATAFILE_0>	USERS	1000

Table 22. Redo log files

Group	Size(K)
1	71200
2	71200
3	71200

Preparing to install the domain (Linux, UNIX)

Before you install IBM Information Server, you need to prepare your system by verifying system requirements. The steps for preparing your system depend on the operating system and which layers and product modules that you will install.

Ensure that your system meets the system requirements for the product modules and components that you want to install.

Procedure

To prepare to install the domain:

1. Log on to the system as a root user.
2. Turn off any firewall software that is installed on the system.
3. Optional: Turn off your antivirus software.
4. If you have an earlier version of IBM Information Server or DataStage installed on your computer, determine whether you need to remove or upgrade the earlier version.

Option	Description
Product or component	Action
IBM Information Server	• To preserve version 8.0 of IBM Information Server, you must install version 8.0.1 using the multiple server method, which is described in <i>Migrating to IBM Information Server Version 8</i>. • If you do not need to preserve Version 8.0 of IBM Information Server, you must remove it.
WebSphere DataStage	• If you have DataStage 7.0 or later, you can upgrade it. See <i>Migrating to IBM Information Server Version 8</i> for instructions. • If you have DataStage 7.0 or later, you can preserve it by using the multiple server method, which is described in <i>Migrating to IBM Information Server Version 8</i>. • If you have a version of DataStage earlier than 7.0, you must upgrade to version 7.5.2 before you can upgrade to the current version. See <i>Migrating to IBM Information Server Version 8</i> for instructions.

5. If user names and passwords do not exist, you might need to create them.
If your system uses NIS, NIS+, or other security programs that might restrict the installation wizard's ability to create the credentials for the components that you install, create the users and groups that you need before you install IBM Information Server.
6. If you install the metadata repository on a different system, you must supply the following information about the metadata repository during the domain installation:
 - Host name
 - Database port

- Database name
- User name and password
- Instance name. Ensure that the instance is started.

If you use an existing instance of a DB2 database for your metadata repository, ensure that it is installed on a computer that is accessible from the computer where you are installing the metadata repository.

7. Record the following information:

- Database owner and password
- Database name for the metadata repository.
- Database instance
- Database location

8. Prepare an instance of WebSphere Application Server:

To...	Do this...
Use an existing instance of WebSphere Application Server	Create a dedicated profile by using the WebSphere Application Server Profile Creation wizard.
Install a new instance of WebSphere Application Server	Complete the steps for preparing your operating system for WebSphere Application Server as described in the WebSphere Application Server information center.

After you prepare your system for installing the domain layer, prepare your system or systems for installing the engine layer.

Creating a server profile for existing instances of WebSphere Application Server

To use an existing instance of WebSphere Application Server, you need to create a server profile before you begin the installation.

To create a server profile, use the Profile creation wizard of WebSphere Application Server.

Preparing to install the engine (Linux, UNIX)

Before you install IBM Information Server, you need to prepare your system by verifying system requirements. The steps for preparing your system depend on the operating system and which layers and product modules that you will install.

Ensure that your system meets the system requirements for the product modules and components that you want to install.

Procedure

To prepare to install the engine:

1. Log on to the system as a root administrator.
2. Turn off any firewall software that is installed on the system.
3. Optional: Turn off your antivirus software.
4. To ensure the best possible performance of the installation wizard, copy the contents of the IBM Information Server DVD or CDs to the system where you are performing the installation instead of installing the product from the DVD or CDs. Do not run the installation from a cross-mounted DVD or CD drive.

5. If user names and passwords do not exist, you might need to create them. If your system uses NIS, NIS+, or other security programs that might restrict the installation wizard's ability to create the credentials for the components that you install, create the users and groups that you need before you install IBM Information Server.
6. If your system uses NIS or NIS+, update the NIS master `/etc/system` files to include an entry for `dsrpc` as port `31538/tcp`.
7. If the domain is installed on another system, record the following information about WebSphere Metadata Server:
 - Host name
 - TCP/IP port. This is the HTTP port of the WebSphere Application Server profile, which is 9080 by default.
8. If you want to access federated data sources, ensure that you meet the prerequisites for installing relational and nonrelational wrappers before you install WebSphere Federation Server.
9. Set the kernel parameters.

After you prepare your system or systems to install the metadata repository, domain, and engine, you can install WebSphere Information Server.

Kernel parameters (Linux, UNIX)

WebSphere DataStage and WebSphere QualityStage have requirements for setting kernel parameters.

When you run the installation wizard, WebSphere DataStage and WebSphere QualityStage check the settings. The installation wizard warns you if any are too low, and gives you the opportunity of quitting the install so you can reconfigure the kernel.

The following tables give minimum requirements for the `MSGMAX`, `MSGMNB`, `MSGSEG`, `SEMMNS`, and `SEMMSL` parameters.

Table 23. Minimum requirements for the MSGMAX, MSGMNB, MSGSEG, SEMMNS, and SEMMSL parameters

Platform	MSGMAX parameter	MSGMNB parameter	MSGSEG parameter	SEMMNS parameter	SEMMSL parameter
AIX	N/A	N/A	N/A	N/A	N/A
Linux	8192	16384	N/A	128000	250

The following table gives minimum requirements for the `SEMMNI`, `SEMOPM`, `MAXUPROC`, and `NOFILES` parameters.

Table 24. Minimum requirements for the SEMMNI, SEMOPM, MAXUPROC, and NOFILES parameters

Platform	SEMMNI parameter	SEMOPM parameter	MAXUPROC parameter	NOFILES parameter
AIX	N/A	N/A	100 per processor	1000
Linux	1024	32	100 per processor	1000

The following guidelines are the minimum settings for shared memory kernel parameters. These settings apply to all Linux or UNIX platforms:

- SHMMAX - maximum shared memory segment size.
Minimum of 307200000. If disk caching turned is on and the DISKCACHE variable is set larger than 512, then the value of the SHMMAX parameter must be greater.
- SHMMNI - shared memory identifiers.
The value for this parameter depends on how many client connections, work processes, and IPC links are active at a time. 2000 is an adequate minimum value for all but the most complex systems.
- SHMSEG - maximum attached shared memory segments per process.
A minimum setting of 200 is adequate for this parameter.

These settings give the generic names for the kernel parameters. The actual name and the case vary for each Linux or UNIX platform.

If necessary, rebuild the kernel with changes to these parameters.

Kernel configuration (AIX):

For WebSphere DataStage and WebSphere QualityStage, you must set the per-user parameters to ensure that users who run WebSphere DataStage and WebSphere QualityStage jobs have sufficient resources.

The following table lists the user parameters, their description, and the values that ensure sufficient resources for running WebSphere DataStage and WebSphere QualityStage on AIX.

Table 25. User parameters and settings for AIX systems

User Parameter	Descriptions	Setting
fsize	Largest file a user may create	2 GB - 512 bytes (4194303 512-byte blocks)
data	Largest data segment, which is called a heap, that a program can have	128 MB minimum (262144 512-byte blocks)
stack	Largest stack size a program can have	32 MB minimum (65536 512-byte blocks)
fsize	Maximum amount of physical memory that a user process may use	64 MB minimum (131072 512-byte blocks)

Preparing to install WebSphere Federation Server

Editing the Oracle 9i gencntsh script and creating the libcntsh file before installation (Linux and Solaris):

If your system uses the Oracle 9i client, the Oracle NET8 wrapper, and a federated server that runs the Linux or Solaris operating systems, your federated instance will fail when you attempt remote operations that include Oracle and the NET8 wrapper unless you complete this task.

Before you begin

- Install the Oracle 9i client if it is not already installed.

- Back up the following files:
 - Linux 32-bit: \$ORACLE_HOME/bin/genclntsh, \$ORACLE_HOME/lib/libclntsh.so.9.0, \$ORACLE_HOME/lib/libclntst9.a
 - Linux 64-bit: \$ORACLE_HOME/lib32/libclntsh.so.9.0, \$ORACLE_HOME/lib32/libclntst9.a
 - Solaris 32-bit: \$ORACLE_HOME/bin/genclntsh, \$ORACLE_HOME/lib/libclntsh.so.9.0, \$ORACLE_HOME/lib/libclntst9.a
 - Solaris 64-bit: \$ORACLE_HOME/lib32/libclntsh.so.9.0, \$ORACLE_HOME/lib32/libclntst9.a

Procedure

To edit the genclntsh script and create the libclntsh file before you install:

1. In a text editor, open the \$ORACLE_HOME/bin/genclntsh script.
 On Linux, add -Wl,-Bsymbolic to the link line. For example:

```
LD="gcc -shared -Wl,-relax -L${ORACLE_HOME}/lib -Wl,-Bsymbolic"
# shared library link command
```

 On Solaris, add -Bsymbolic to the link line. For example:

```
LD="ld -m -i -G -z text -L${ORACLE_HOME}/${LIB} -Bsymbolic"
# shared library link command
```
2. From a command prompt, run the genclntsh script to create the libclntsh file.
3. When the WebSphere Federated Server or WebSphere Replication Server is installed, the changes that were made to the genclntsh script are recorded in the Oracle wrapper library.

Creating groups, user names and passwords before you install IBM Information Server (Linux, UNIX)

If user names and passwords do not exist, you might need to create them. If your system uses NIS, NIS+, or other security programs that might restrict the installation wizard's ability to create the credentials for the components that you install, create the users and groups that you need before you install IBM Information Server.

Before you begin

- You must have root authority on the system to complete this task.
- Review the user account requirements and naming rules for DB2 before you create users and groups.

Restrictions

You cannot install the suite on systems where DB2 administration tools are installed.

About this task

When you create the user names and passwords for the application server and the IBM Information Server administrator, you choose to create them in an internal user registry or in the local operating system user registry.

Procedure

To create groups, user IDs, and passwords before you install the suite on Linux and UNIX:

1. Log on to the system as a root user.
2. For the metadata repository, follow the instructions for creating the required user names, groups, and passwords for DB2 that are documented at publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp?topic=/com.ibm.db2.udb.uprun.doc/doc/t0006742.htm. You can use the same commands to create the other credentials that you need. Record the user names and passwords that you create, so that you can provide them during the installation process.

After you create the user names and passwords, complete the process for preparing your system for installing the suite on Linux and UNIX.

Users and groups by layer (Linux, UNIX)

The metadata repository layer and the domain layer require users and might require groups.

This table lists the type of user required for each layer, a sample user name, and a group name, if required.

Table 26. Users and groups

Layer	Required user	Sample user name	Group name
Metadata repository	Metadata repository database owner	xmeta	Not applicable
Metadata repository	Instance owner	db2inst1	db2iadm1
Metadata repository	Fenced user	db2fenc1	db2fadm1
Metadata repository	DB2 Administration Server (DAS) user	dasusr1	dasadm1
Domain	Application server ID	wasadmin	Not applicable
Metadata repository	Information analysis database owner	iauser	Not applicable
Engine	WebSphere DataStage administrator	dsadm	dstage
Domain	IBM Information Server administrator	admin	Not applicable

The DB2 instance owner, DAS user, and fenced user are required only if you use DB2 as the metadata repository.

If NIS is enabled, go to publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp?topic=/com.ibm.db2.udb.uprun.doc/doc/r0007059.htm for additional requirements for configuring users and groups.

Creating a server profile for existing instances of WebSphere Application Server

To use an existing instance of WebSphere Application Server, you need to create a server profile before you begin the installation.

To create a server profile, use the Profile creation wizard of WebSphere Application Server.

Configuring the analysis database and the analysis engine

If you are installing WebSphere Information Analyzer and you are using DB2 as your analysis database, that database is automatically created by default during installation. If you want to use Microsoft SQL Server or Oracle as your analysis database, you must configure the analysis database before you install the suite. Additionally, if you install the analysis engine on a local computer or remote computer, you must manually create the analysis database DSN.

Configuring Microsoft SQL Server as the analysis database

To use Microsoft SQL Server as your analysis database:

See the instructions in the \Database Support\SQLServer2005\AnalysisDatabase directory. By default, this directory is located in the root directory of the IBM Information Server CD, DVD, or download image.

Configuring Oracle as the analysis database

To use Oracle as your analysis database:

See the instructions in the \Database Support\Oracle10g\AnalysisDatabase directory. By default, this directory is located in the root directory of the IBM Information Server CD, DVD, or download image.

Configuring the analysis engine

If you install the analysis engine on a local computer or a remote computer, you must manually create a DSN for the analysis database on that computer. This DSN must be a system DSN. Additionally, you must create a DSN for each source database or file that you want to profile on this computer.

Related tasks

“Configuring a connection to the analysis database for WebSphere Information Analyzer” on page 120

When you analyze data, the results of your data analysis are stored in the analysis database. The connection to the analysis database must be configured after you install the suite.

Chapter 4. Installing IBM Information Server

You can install the suite with the installation wizard, from the console, or in silent mode. A full installation of all the components can take 2-3 hours to complete.

Installing IBM Information Server (Windows)

IBM Information Server includes a set of collaborative components that can be distributed across multiple computers.

Installing in graphical mode

The suite can be installed using the installer in graphical mode, console mode, or silent mode. In graphical mode, you use a wizard to install the suite.

Complete the task, “Preparing your system for installation (Windows)” on page 21.

About this task

When you install WebSphere Application Server with IBM Information Server, the installation wizard creates a profile names default and an application server instance named server1. The application server hosts WebSphere Metadata Server.

To install IBM Information Server on Windows:

1. Log on to the system as a root administrator.
2. Turn off any firewall software that is installed on the system.
3. Optional: Turn off your antivirus software.
4. Insert the IBM Information Server DVD or CD, or go to the root directory of the downloaded installation image.
5. In the root directory of the CD, DVD, or downloaded installation image, double-click **install.exe**.
6. Follow the instructions in the installation wizard to complete the installation for the components that you want to install.

The metadata repository layer option in the installation wizard configures DB2 as the metadata repository. If you use Microsoft SQL Server or Oracle as the metadata repository, you do not need to install the metadata repository layer.

If you plan to install WebSphere Information Services Director, WebSphere DataStage and WebSphere QualityStage, or WebSphere Information Analyzer clients, ensure that you install the required components by selecting these components when you install the client, domain, and engine.

7. After the installation is complete, restart the system.
8. Turn on the firewall software that runs on the computer if necessary.
9. If you turned off the antivirus software on your system, turn it on.
10. If you are using an existing instance of the application server and it is not configured to start automatically after the computer is restarted, start the application server.
11. Configure credential mappings between the WebSphere DataStage server users and WebSphere Metadata Server.

Installing in console mode

The suite can be installed using the installer in graphical mode, console mode, or silent mode. You can install the suite in console mode from a command prompt.

Complete the task, “Preparing your system for installation (Windows)” on page 21.

To install the suite in console mode:

1. From the root directory of the DVD, CD, or image on your hard drive, run the following command to start the installer in console mode. You can optionally provide a log file and use a different temporary directory.

```
install -console
```

2. Follow the prompts to install the suite.
3. After the installation is complete, restart the system.

install command syntax for console installations

The install command syntax is used to install the IBM Information Server suite from a command prompt.

```
►►—install—┐└─console—-is:log—fileName—-is:tempdir—directory—┘◄◄
```

Option	Description
-console	Runs the installation in console mode, which uses the command line. You might want to use this installation method if your computer does not support graphical user interfaces or if you have accessibility requirements.
-is:log <i>fileName</i>	Writes detailed information about the installation process to the specified file name. The output includes all of the std out and std err messages from the Java process.
-is:tempdir <i>directory</i>	Sets the path to the directory where the installer writes temporary files. If the specified directory does not exist or is not a directory, the installer uses the system temporary directory instead, and no error message is provided.

Examples: Console installation

The following example shows how to run the install command in console mode on Windows. The log file is stored in the c:\logs\log.txt file. The temporary directory is c:\temp.

```
install.exe -console -is:log c:\logs\log.txt -is:tempdir c:\temp
```

Installing in silent mode

The suite can be installed using the installer in graphical mode, console mode, or silent mode. You can install the suite in silent mode using a response file. A response file is useful for performing multiple installations with the same settings or for performing an unattended install.

Complete the task, “Preparing your system for installation (Windows)” on page 21.

About this task

You can create the response file manually, or you can use the installation wizard to create it. The installation wizard can create the response file without installing the suite.

To install IBM Information Server silently on Windows:

1. Create a response file. You can run the installation wizard in graphical mode and capture your settings. You can create a response file without installing any software. The default name and location of the response file is C:\Documents and Settings*user*\is_install.rsp. *user* is the home directory of the user who created the response file.
2. If necessary, move the response file to the computer where you want to install the suite.
3. In a text editor, edit the response file to add the passwords for the user credentials that are required for the installation. The installation wizard inserts placeholders instead of the passwords into the response file. You must replace the placeholders with the passwords before you run the installation.
4. Open a command prompt and change to the root directory of the CD, DVD, or downloaded installation image.
5. Run the installer in silent mode from the command prompt:
`install.exe -silent -options response_file_name`
6. Restart your computer.

install command syntax for silent installations

The install command syntax is used to install the suite silently.

►►—install—-silent—-options—*absoluteResponseFilePath*—◄◄

Option	Description
-silent	Runs the installer in silent mode with no user interaction. Installation settings are provided in a response file which can be created by running the installation wizard in graphical or console mode. You can create a response file without installing any software.
-options <i>absoluteResponseFilePath</i>	Specifies the absolute path to the response file that contains your installation options.

Examples: Silent installation

The following example shows how to run the install command in silent mode on Windows. The installer reads the installation options from the is_install.rsp file.

```
install.exe -silent -options c:\temp\is_install.rsp
```

Stopping and starting the WebSphere DataStage server

You might need to start the WebSphere DataStage server after you complete the suite installation. You must restart the server after you change any of the WebSphere DataStage environment variables.

Before you begin

You must have administrator access to the installation system to complete this task.

Procedure

To stop and start the WebSphere DataStage server on Windows:

1. Click **Start** → **Control Panel** → **WebSphere DataStage**.
2. Click **Stop All Services**, and click **Yes**.
3. Click **Start All Services**, and click **Yes**.
4. Click **OK** to close the window.

Installing IBM Information Server (Linux, UNIX)

IBM Information Server includes a set of collaborative components that can be distributed across multiple computers.

Installing in graphical mode

The suite can be installed using the installer in graphical mode, console mode, or silent mode. In graphical mode, you use a wizard to install the suite.

Prepare your computers for installation.

The clients, metadata repository, and engine must be installed on a computer that is accessible by WebSphere Metadata Server.

When you install WebSphere Application Server with IBM Information Server, the installation wizard creates a profile names default and an application server instance named server1. The application server hosts WebSphere Metadata Server.

Procedure

To install IBM Information Server on Linux or UNIX:

1. Log on to the system as a root user.
2. Turn off any firewall software that is installed on the system.
3. Optional: Turn off your antivirus software.
4. If necessary, mount the IBM Information Server DVD.
5. To improve the performance of the installer, copy the contents of the DVD to the system. If your medium is CD, copy each CD into a directory named `diskn`.
6. In the directory where you copied the contents of the DVD, double-click **install** or run the following command from a shell prompt:
`./install`
7. If you want to install the Resource Tracker, specify the **Custom** installation type in the installation wizard when you install the domain.
8. If you plan to install WebSphere Information Services Director, WebSphere DataStage and WebSphere QualityStage, or WebSphere Information Analyzer clients, ensure that you install the required components by selecting these components when you install the client, domain, and engine.
9. Follow the instructions in the installation wizard to complete the installation for the components that you want to install.
10. Turn on the firewall software that runs on the computer if necessary.
11. If you turned off the antivirus software on your system, turn it on.

12. Stop and start the application server.
13. Stop and start the WebSphere DataStage server.
14. Configure credential mappings between the WebSphere DataStage server and WebSphere Metadata Server.

After the installation is complete, you can configure the suite.

Installing IBM Information Server in console mode

If your computer does not support graphical user interfaces, or if you choose not to use the graphical installation wizard, you can install IBM Information Server from a command line.

- You must have root authority on the system to complete this task.
- If the domain, metadata repository, or engine are installed on different computers, they must all run the same operating system.

To install the suite in console mode:

1. From the root directory of the DVD, CD, or image on your hard drive, run the following command to start the installer in console mode. You can optionally provide a log file and use a different temporary directory.

```
./install -console -is:tempdir directory
```
2. Follow the prompts to install the suite.
3. After the installation is complete, restart the system.

Installing in silent mode

You can install the suite in silent mode by using a response file. A response file is useful for performing multiple installations with the same settings or for performing an unattended install.

Before you begin

- Ensure that your system meets all the prerequisites for installation.
- If the domain, metadata repository, or engine are installed on different computers, they must all run the same operating system.
- Ensure that your all data is saved and other programs are stopped before you begin the installation.

About this task

You can create the response file manually, or you can use the installation wizard to create it. The installation wizard can create the response file with or without installing the suite.

Procedure

To install IBM Information Server silently on Linux and UNIX:

1. Create a response file. You can run the installation wizard in graphical mode and capture your settings. You can create a response file without installing any software. On Linux, the default name and location of the response file is `/root/is_install.rsp`. On AIX, the default name and location is `/is_install.rsp`.
2. If necessary, move the response file to the computer where you want to install the suite.
3. In a text editor, edit the response file to add the passwords for the user credentials that are required for the installation. The installation wizard inserts

placeholders instead of the passwords into the response file. You must replace the placeholders with the passwords before you run the installation.

4. From a command line, change to the root directory of the CD, DVD, or downloaded installation image.
5. Run the installer in silent mode from the command line:
`./install -silent -options response_file_name`

install command syntax for silent installations

The install command syntax is used to install the IBM Information Server suite from a command prompt.

►►—install—-silent—-options—*absoluteResponseFilePath*—◀◀

Option	Description
-silent	Runs the installer in silent mode with no user interaction. Installation settings are provided in a response file, which can be created by running the installation wizard in graphical mode or console mode. You can create a response file without installing any software.
-options <i>absoluteResponseFilePath</i>	Specifies the absolute path to the response file that contains your installation options.

Example: Silent installation

The following example shows how to run the install command in silent mode on Linux or UNIX.

```
./install -silent -options /opt/temp/is_install.rsp
```

Figure 3. Silent installation on Linux or UNIX

Stopping and starting the WebSphere DataStage server

You need to stop and start the WebSphere DataStage after you complete certain tasks. For example, you must stop and restart the server after you change or add environment variables or after you complete the suite installation.

Before you begin

You must have root authority on the system to complete this task.

Procedure

To stop and start the WebSphere DataStage server:

Run the following commands.

```
cd $DSHOME
. ./dsenv
bin/uv -admin -stop
bin/uv -admin -start
```

`$DSHOME` is the name of the installation directory for the runtime engine. The default directory is `/opt/IBM/InformationServer/Server/DSEngine`.

Chapter 5. Configuring IBM Information Server

Depending on which products you installed or the tasks that you want to complete, you might need to complete some configuration tasks.

IBM WebSphere DataStage and WebSphere QualityStage

If you installed WebSphere DataStage and WebSphere QualityStage, you need to complete the following tasks before you use the products:

- Create users and groups in the IBM Information Server Web Console.
- If the WebSphere DataStage server and the IBM Information Server do not share the same user registry, create credential mappings between the metadata server and the WebSphere DataStage server. See the *IBM Information Server Administration Guide* for information about sharing user registries.
- If you want to run parallel jobs, configure the parallel engine.
- If you have DataStage 7.0 or 7.5 jobs or QualityStage 7.0 or 7.5 jobs that you want to use, you must migrate them to WebSphere DataStage and WebSphere QualityStage 8.0. For information about migrating your jobs, see *Migrating to IBM Information Server Version 8*.

WebSphere Information Analyzer

If you installed IBM WebSphere Information Analyzer, you must complete the following tasks before you can analyze data:

- Create users and groups in the IBM Information Server Web Console.
- Verify the WebSphere DataStage credentials that are used to connect to the analysis engine.
- Configure the connection to the analysis database for the WebSphere DataStage project that you want to use.
- Locate and define the host computers and the data stores in your organization. See the *IBM WebSphere Information Analyzer User Guide* for information about configuring your system resources.

WebSphere Federation Server

If you installed IBM WebSphere Federation Server, you must configure the federated system to access your data sources.

Creating users and groups for IBM Information Server

You must create users and groups before you can use the suite.

If you use the internal user registry, you create users and groups in the Web console. If you use an external user registry, you create users and groups using the administration functions of the external user registry.

You must be a suite administrator to complete this task.

Procedure

To create users and groups for IBM Information Server:

1. On the computer where you installed the clients, log in to the Web console, which is on the computer where the domain is installed.
 - Windows: Click **Start** → **IBM Information Server** → **IBM Information Server Web Console** and log in to the console with your suite administrator credentials.
 - Linux or UNIX: Open Microsoft Internet Explorer on a Windows computer and go to the URL of the Web console. Log in to the Web console with your suite administrator credentials. The URL takes the form `http://host_server:port`. *host_server* is the name or IP address of the computer where the domain is installed. *port* is the port number for the Web console. The default port number is 9080.
2. Click the **Administration** tab and create users and groups.

You must assign at least one product role to users before they can log in to the clients.

After you create users and groups, you can complete the following tasks:

- Create credential mappings for suite users with WebSphere DataStage and WebSphere QualityStage privileges.
- Start WebSphere Information Analyzer.
- Start WebSphere Information Services Director.
- Configure the suite.

Setting up permissions and groups for Windows 2003 Server

Microsoft Windows 2003 Server uses security policies that are different from other Windows platforms. To access a WebSphere DataStage server on Windows 2003 Server, you need to perform additional steps.

Setting up permissions and groups to access WebSphere DataStage

If Windows 2003 Server is not a domain controller, you must follow these additional steps to access the WebSphere DataStage server.

You must be logged in as an administrator.

To set up permissions and groups:

1. Configure the server to allow local users to log in.
 - a. From the Start menu, click **Control Panel** → **Administrative Tools** → **Local Security**.
 - b. From the main directory of the Local Security window, expand **Local Policies** → **User Rights Assignment** to display the policies.
 - c. In the Local Security window click the policy **Allow Log on Locally**, and click **Actions** → **Properties**.
 - d. In the Allow Log on Locally Properties window, click the **Add User or Group** button.
 - e. In the Select Users, Computers, or Groups window, click the **Locations** button, click the name of your local computer, and click **OK**.
 - f. In the Select Users, Computers, or Groups window, click the **Advanced** button, and then click the **Find Now** button.
 - g. From the results of the search, click **Authenticated Users**, and then click **OK** three times to save the results and to return to the Local Security window.

- h. Close the Local Security Policy window.
2. Create a group of users to access WebSphere DataStage.
 - a. From the Start menu, click **Control Panel** → **Administrative Tools** → **Computer Management**.
 - b. From the main directory of the Computer Management window, expand **System tools** → **Local Users and Groups** → **Groups**.
 - c. In the Computer Management window, click **Action** → **New Group**.
 - d. In the New Group window, type the name for the group, click **Create**, and click **Close**.
3. Add users to the group that you created.
 - a. From the Computer Management window, click **Groups**.
 - b. Click the name of the group that you want to add users to.
 - c. Click **Actions** → **Add to Group**.
 - d. In the User Properties window, click the **Add** button.
 - e. In the Select Users, Computers, or Groups window, click the **Location** button.
 - f. Click the name of your local computer, and then click **OK**.
 - g. In the Select Users window, click the **Advanced** button.
 - h. In the window that opens, click the **Find Now** button.
 - i. Click the names of users that you want to include in the group, and click **OK**. At a minimum, include all authenticated users.
 - j. Click **OK** three times to return to the Computer Management window.
 - k. Close the Computer Management window.
4. Set permissions for the WebSphere DataStage folder.
 - a. In Windows Explorer, locate the WebSphere DataStage server folder. The default location is c:\IBM\Information Server\Server.
 - b. Click **File** → **Properties**.
 - c. In the Properties window, click the **Security** tab, and click **Add**.
 - d. In the Select Users, Computers, or Groups window, click the **Locations** button, click the name of your local computer, and click **OK**.
 - e. In the Select Users window, click the **Advanced** button.
 - f. In the window that opens, click the **Find Now** button.
 - g. Click the name of the group that you want to set permissions for.
 - h. Click **OK**, and then click **OK** again.
 - i. In the Permissions list, click **Modify**, click **Write** in the Allow column, and click **OK**.
 - j. If you receive a message to confirm your changes, confirm by clicking **Apply changes to this folder, subfolders and files**.
5. Set permissions for the registered-server.xml file.
 - a. In Windows Explorer, locate the registered-server.xml file. The default location is c:\IBM\Information Server\ASBNode\conf.
 - b. Click **File** → **Properties**.
 - c. In the Properties window, click the **Security** tab, and click **Add**.
 - d. In the Select Users, Computers, or Groups window, click the **Locations** button, click the name of your local computer, and click **OK**.
 - e. In the Select Users window, click the **Advanced** button.
 - f. In the window that opens, click the **Find Now** button.

- g. Click the name of the group that you want to set permissions for.
- h. Click **OK**, and then click **OK** again.
- i. In the Permissions list, click **Modify**, click **Write** in the Allow column, and click **OK**.
- j. If you receive a message to confirm your changes, confirm by clicking **Apply changes to this folder, subfolders and files**.

Setting up permissions and groups to access WebSphere DataStage if the server is a domain controller

If Windows 2003 Server is a domain controller, you must follow these additional steps to access the WebSphere DataStage server.

You must be logged in as an administrator.

Because you cannot add the built-in authenticated users group to a group that you create in steps 2 and 3, you might prefer to skip steps 2 and 3 and use the authenticated users group directly.

1. Configure the server to allow local users to log in.
 - a. From the Start menu, click **Control Panel** → **Administrative Tools** → **Domain Security Policy**.
 - b. From the main directory of the Domain Security Policy window, expand **Local Policies** → **User Rights Assignment** to display the policies.
 - c. In the Domain Security window, click the policy **Allow Log on Locally** and click **Actions** → **Properties**.
 - d. In the Allow Log on Locally Properties window, click the **Add User or Group** button.
 - e. Click **Browse**.
 - f. In the Select Users, Computers, or Groups window, click **Advanced** and then click the **Find Now** button.
 - g. From the results of the search, click **Authenticated Users**, and then click **OK** three times to return to the Domain Security Policy window.
 - h. Close the Domain Security Policy window.
2. Create a group of users to access WebSphere DataStage.
 - a. From the Start menu, click **Control Panel** → **Administrative Tools** → **Active Directory and Computers**.
 - b. In the Active Directory and Computers window, click **Users** in the current domain.
 - c. In the window that opens, click **Action** → **New Group**.
 - d. In the New Group window, type the name for the group.
 - e. Leave **Group scope** as **Global** and **Group type** as **Security**, and click **OK**.
3. Add users to the group that you created.
 - a. From the Users in the current domain window, click the name of the group that you want to add users to, and click **OK**. Authenticated users are not available.
 - b. Click **Action** → **Properties**.
 - c. In the Properties window, click the **Members** tab, and then click **Add**.
 - d. In the window that opens, click **Advanced**, and then click **Find Now**.
 - e. Click the names of users that you want to add to the group, and then click **OK**. Authenticated users are not available.

- f. Click **OK** two times to save your results and to return to the Active Directory and Computers window.
 - g. Close the Active Directory and Computers window.
4. Set permissions on WebSphere DataStage server folder.
 - a. In Windows Explorer, locate the WebSphere DataStage server folder. The default location is c:\IBM\Information Server\Server.
 - b. Click **File** → **Properties**.
 - c. In the Properties window, click the **Security** tab, and click **Add**.
 - d. In the Select Users, Computers, or Groups window, click the **Locations** button.
 - e. In the window that opens, click the **Advanced** button, and then click the **Find Now** button.
 - f. Click the name of the group that you want to set permissions for.
 - g. Click **OK**, and then click **OK** again.
 - h. Click the name of the group that you want to set permissions for.
 - i. In the Permissions list, click **Modify**, click **Write** in the Allow column, and click **OK**.
 - j. If you receive a message to confirm your changes, confirm by clicking **Apply changes to this folder, subfolders and files**.
5. Set permissions for the registered-server.xml file.
 - a. In Windows Explorer, locate the registered-server.xml file. The default location is c:\IBM\Information Server\ASBNode\conf.
 - b. Click **File** → **Properties**.
 - c. In the Properties window, click the **Security** tab, and click **Add**.
 - d. In the Select Users, Computers, or Groups window, click the **Locations** button, click the name of your local computer, and click **OK**.
 - e. In the Select Users window, click the **Advanced** button.
 - f. In the window that opens, click the **Find Now** button.
 - g. Click the name of the group that you want to set permissions for.
 - h. Click **OK**, and then click **OK** again.
 - i. In the Permissions list, click **Modify**, click **Write** in the Allow column, and click **OK**.
 - j. If you receive a message to confirm your changes, confirm by clicking **Apply changes to this folder, subfolders and files**.

Configuring credential mappings for WebSphere DataStage and WebSphere QualityStage users

You can create credential mappings for suite users with WebSphere DataStage and WebSphere QualityStage privileges.

Prerequisites

- You must be a suite administrator to complete this task.
- Create users and groups in the Web console before you begin this task.

If you use the IBM Information Server user registry, you must create credential mappings before you can use WebSphere DataStage and WebSphere QualityStage clients.

If you want the suite users to configure their own credential mappings, you can give them instructions on how to log into the Web console and set their own credential mappings.

To configure credential mappings:

1. On the computer where you installed the clients, log in to the Web console, which is on the computer where the domain is installed.
 - Windows: Click **Start** → **IBM Information Server** → **IBM Information Server Web Console** and log in to the console with your suite administrator credentials.
 - Linux or UNIX: Open Microsoft Internet Explorer on a Windows computer and go to the URL of the Web console. Log in to the Web console with your suite administrator credentials. The URL takes the form `http://host_server:port`. *host_server* is the name or IP address of the computer where the domain is installed. *port* is the port number for the Web console. The default port number is 9080.
2. On the Administration tab, expand the **Domain Management** section and click **DataStage Credentials**.
3. Select the server where the domain is installed.
4. Click **Open User Credentials**.
5. Click **Browse** to search for suite users.
6. Optional: Specify additional search criteria, and click **Search** to provide a list of users.
7. From the search results, select the suite users that you want to map to the WebSphere DataStage and WebSphere QualityStage server and click **OK**.
8. On the DataStage Credentials tab, select the user or users that you want to map to the metadata server. If you want to map some suite users to one WebSphere DataStage and WebSphere QualityStage user and other suite users to a different WebSphere DataStage and WebSphere QualityStage user, select one subset of users and continue.
9. In the Map User Credentials window, specify the WebSphere DataStage and WebSphere QualityStage users that you want to map to the metadata server.
10. In the Assign User Credentials window, specify the user credentials for the WebSphere DataStage and WebSphere QualityStage server. This user must exist on the local operating system where the WebSphere DataStage and WebSphere QualityStage server is installed.
11. Click **Apply**.
12. To map additional suite users to a user on the WebSphere DataStage and WebSphere QualityStage server, repeat steps 5 through 11 as necessary.

After you map the credentials and the suite user or group is added to a WebSphere DataStage or a WebSphere QualityStage User role in the WebSphere DataStage Administrator client, you can log in to the Designer client.

Assigning WebSphere DataStage and WebSphere QualityStage project roles

You can create suite users by using the IBM Information Server Web console. You assign WebSphere DataStage and WebSphere QualityStage project roles with the Administrator client.

To assign WebSphere DataStage and WebSphere QualityStage user roles:

1. On the computer where you installed the clients, click **Start → All Programs → IBM Information Server → IBM WebSphere DataStage and QualityStage Administrator**.
2. Click the **Projects** tab in the Administrator window.
3. Select the project that you want to assign user roles to.
4. Click **Properties** to open the Project Properties window.
5. Click the **Permissions** tab. By default this shows all the users or groups that are defined as DataStage Administrators in the Administration tab of the Web console. Such users automatically have access to the DataStage Administrator role for this project. You cannot change the authorization of users who have the DataStage Administrator role.
6. Click **Add User or Group**. The Add Users and Groups dialog box appears.
7. In the Add Users and Groups window, click the user or users that you want to assign a role to, and click **Add**. Click **Add All** to assign all users to the same role.
8. In the **User Role** list, select the role that you want to assign the user or users to, and click **OK**.

WebSphere DataStage and WebSphere QualityStage user roles

To prevent unauthorized access to WebSphere DataStage or WebSphere QualityStage projects, you must assign the appropriate user role to the WebSphere DataStage and WebSphere QualityStage users or groups on your system.

There are five WebSphere DataStage and WebSphere QualityStage user roles:

DataStage and QualityStage Developer

This role has access to all areas of a DataStage or QualityStage project (except for protected projects) and can use the Director client and the Designer client.

DataStage and QualityStage Production Manager

This role has access to all areas of a DataStage or QualityStage project and can also create and manipulate protected projects. The role can use the Director client and the Designer client. (On UNIX systems, the Production Manager must be root or the administrative user to protect or remove protection from projects.)

DataStage and QualityStage Operator

This role has permission to run and manage DataStage jobs from the Director client within projects that are assigned to the role.

DataStage and QualityStage Super Operator

This role has permission to run and manage DataStage jobs and QualityStage jobs and to view items in the Designer client. This role does not have permission to edit items.

DataStage and QualityStage Administrator

This role has full access to all areas of all DataStage and QualityStage projects.

Configuring WebSphere DataStage and WebSphere QualityStage

After you install IBM WebSphere DataStage and WebSphere QualityStage, you must configure them.

Tuning the WebSphere DataStage and WebSphere QualityStage server (Linux, UNIX)

The WebSphere DataStage and WebSphere QualityStage server can be tuned to maximize performance.

The hardware configurations of Linux and UNIX systems can vary widely. WebSphere DataStage contains a generically configured WebSphere DataStage and WebSphere QualityStage server (named DSEngine) that can be installed in most environments. However, larger hardware configurations might not use all of the available resources, such as shared memory.

About this task

Use the shmtest utility to obtain the best server values for tuning a particular hardware configuration.

Use the shmtest utility to obtain valid shared memory segment attach offset values. If the WebSphere DataStage server engine is running and no options are specified, the tunable parameters in the current \$DSHOME/.uvconfig binary file are used. Regardless of whether the server is running, the -f option can be used to specify a text file that contains the engine tunable parameters to be tested. The results of this process are valid shared memory attach offset values that are stored in the \$DSHOME/uvconfig text file. You can make the new configuration active.

If you run the shmtest utility on a text file with the server stopped, the results reflect the best offset locations where segments can be attached. With the server running, the shared memory in use by the running engine is available to the shmtest utility and offset locations bigger than necessary are returned in the result.

Procedure

To tune the WebSphere DataStage and WebSphere QualityStage server:

1. Rename your original \$DSHOME/uvconfig file to \$DSHOME/uvconfig.bak.
2. If the WebSphere DataStage server is running, run the following command to stop it.
`$DSHOME/bin/uv -admin -stop`
3. Run the following command to test the current configuration in the \$DSHOME/uvconfig text file:
`$DSHOME/bin/shmtest -f $DSHOME/uvconfig`
When the utility completes, it returns recommend attach values for the DMEMOFF, PMEMOFF, CMEMOFF, and NMEMOFF parameters.
4. Change the tunable values DMEMOFF, PMEMOFF, CMEMOFF, and NMEMOFF in \$DSHOME/uvconfig to the recommended values.
5. Save the changes to the uvconfig file.
6. As a user with root privileges or as the WebSphere DataStage administrator, run the following command to generate a new \$DSHOME/.uvconfig binary file:
`$DSHOME/bin/uv -admin -regen`
7. Start the WebSphere DataStage server by running the command:
`$DSHOME/bin/uv -admin -start`

shmtest utility syntax and options

The shmtest utility obtains the best WebSphere DataStage server values for tuning a particular hardware configuration.

Syntax

The shmtest utility has the following syntax.

```
$DSHOME/bin/shmtest [options]
```

Options

- c *size*. Specify the size of the Catalog shared memory segment to override the default size of 10485760.
- f *path*. Specify the path of the text file of parameters to override the current \$DSHOME/.uvconfig parameters.
- h Display help details.
- n *size* specify the size of the NLS shared memory segment size to override the default size of 6291456.
- s Display the shmat utility detail during location determination (for debugging).
- v Display processing details.

Configuring the WebSphere DataStage and WebSphere QualityStage server environment (Windows)

After you install the WebSphere DataStage and WebSphere QualityStage server on Windows, you must complete some configuration tasks.

Setting the environment variables for the Microsoft Visual Studio .NET 2003 C++ compiler version 7.1 or later

If you are using Microsoft Visual Studio .NET 2003 C++ for compiling jobs, you must set the APT_COMPILEOPT and APT_LINKOPT environment variables after you install the suite.

Procedure

1. Click **Start** → **IBM Information Server** → **IBM WebSphere DataStage and QualityStage Administrator**.
2. Log in to the Administrator interface.
3. Click **Projects** → **Properties** → **Environment**
4. Select **Parallel** → **Compiler** → **APT_COMPILEOPT** and delete the string -W/Zc:wchar_t- from the end of the existing setting in the Value field.
5. Select **Parallel** → **Compiler** → **APT_LINKOPT** and delete the string -W/Zc:wchar_t- from the end of the existing setting in the Value field.

After you complete this task, the compiler is ready to use.

Configuring a computer to be a processing node for MPP (Windows)

WebSphere DataStage can take advantage of massively parallel processing (MPP). With a user ID and a secure shell connection, you can configure a WebSphere

DataStage server as a parallel processing node, which can share the processing load with other WebSphere DataStage servers.

You must have administrator access to the installation system to complete this task.

1. Create a generic WebSphere DataStage user account.
 - a. Click **Start** → **Run** → **cmd** to open a Window command prompt. *cmd* is the text to enter in the **Open:** field.
 - b. From the command window, create the generic dstage user using the following format:

```
NET USER dstage password /ADD/COMMENT:"DataStage user account"
/FULLNAME:"DataStage User"/ACTIVE:YES/EXPIRES:NEVER/PASSWORDCHG:NO
```

password is the password assigned to the dstage user. All machines participating in the MPP configuration should use the same user and password.
 - c. Log off via the Start menu and log in as the dstage user.
 - d. Launch the Control Panel from **Start** → **Control Panel** → **System**, click the **Advanced** tab, and click **Environment variables**.
 - e. Under **User variables for dstage**, select **HOME** and click the **Edit** button. If the entry contains other environment variables (for example, %HOMEPATH%) replace the entire contents with the literal value shown in the User variables list for dstage. The default location is typically C:\Documents and Settings\dstage. MKS does not have the ability to resolve these environment variables when connecting remotely.
 - f. Click **OK** three times to save your changes and to exit from the System Applet.
 - g. Log off as the dstage user and log back in as yourself.
2. Set user permissions. With the new account created, you need to set permissions on some critical PXEngine directories so that the new user will have sufficient access rights to execute jobs. The following table lists the minimum directories that you must set and the appropriate permissions.

Directory	Permissions
C:\tmp	Write, Read, Modify
C:\IBM\InformationServer\Server\Scratch	Write, Read, Modify
C:\IBM\InformationServer\Server\Datasets	Write, Read, Modify

- a. Using Explorer, browse to one of the directories, right click **Sharing and Security**.
 - b. From the Security tab and under Group or user names, click **Add**.
 - c. Enter dstage and click **OK**. You should see dstage under the list of users with permissions to this directory.
 - d. Select dstage and select the permission check boxes for the permissions Write, Read, and Modify.
 - e. Click **Apply**, and then click **OK** to close the window.

Repeat these step for each directory that you want to set permissions for.
3. Set remote shell permissions
 - a. Click **Start** → **Run** → **ksh**. *ksh* is the text to enter in the **Open:** field.

- b. From the shell, run the ssh command to connect to your machine. The format is `ssh -l dstage host name`. *host name* is the name of your computer. You will most likely have to accept an rsa key. You must enter the password for the dstage user.
 - c. Once you are connected, run the rsetup command to add the dstage user and password to the Windows XP/2003 LSA database for use by the rshd service. Enter the same password when prompted.
 - d. Edit the .rhosts file in the dstage user's home directory. You might need to create the .rhosts file if it does not already exist.
 - e. Add the line + dstage.
 - f. Exit from the ssh session.
 - g. Edit the hosts.equiv file located in the \$ROOTDIR/etc folder. The path name is shown in UNIX format, which is required by MKS. The default location is C:\Program Files\MKS Toolkit\etc\hosts.equiv.
 - h. For testing purposes, add the line + + You might need to create this file if it does not already exist.
 - i. Test your configuration by running `rsh -l dstage host name ls`. *host name* is the name of your computer. This gives you a listing of the dstage user's home directory.
4. Configure PXEngine for RSH/SSH. To run PXEngine parallel jobs on remote hosts, you need to use a remote shell script. A sample script named remsh.example is provided for this purpose. The script is located in the PXEngine\etc folder.
 - a. Copy the script file, and rename it as remsh.
 - b. Edit the script file and change the rsh line. Save your changes.

Original script file	Changed script file
#!/bin/sh # Example \$APT_ORCHHOME/etc/remsh rsh "\$@"	#!/bin/sh # Example \$APT_ORCHHOME/etc/remsh rsh -l dstage "\$@"

5. Configure the distribute-component script file, which is located in the PXEngine\etc directory. This script is responsible for managing the installation and removal of PXEngine component binaries on remote computing nodes. This script is used when jobs containing compiled PX transforms need to have their transformer binaries copied to the remote nodes.
 - a. Make a backup copy of the distribute-component script file.
 - b. Edit the script file and change the theUser line.

Original script file	Changed script file
# The user (optional) # theUser=	# The user (optional) # theUser=dstage

- c. Save your changes.

You can now run PXEngine parallel jobs in an MPP environment.

Configuring the WebSphere DataStage server environment (Linux, UNIX)

Any environment variable that is required by the WebSphere DataStage server must be included in the dsenv file. The dsenv file contains a series of Bourne shell

arguments that are referenced when the WebSphere DataStage server starts. The file can be referenced by interactive users, other programs or scripts.

Before you begin

You must have root authority on the system to complete this task.

Shared libraries must be loaded for some plug-ins. In addition, you must include the library path in an environment variable.

Procedure

To configure the environment variables for the WebSphere DataStage server environment:

1. In a Bourne shell, run the following commands from the \$DSHOME/DSEngine directory:

```
. ./dsenv  
bin/uv -admin -stop
```
2. Wait at least 30 seconds for the server to stop, then run the following command.

```
bin/uv -admin -start
```

Library path environment variables

The WebSphere DataStage server has a centralized file, named dsenv, for storing environment variables in the WebSphere DataStage installation directory, which is named \$DSHOME. The default directory is /opt/IBM/InformationServer/Server/DSEngine.

Table 27. Library path environment variables for WebSphere DataStage

Operating system	Environment variables
AIX	LIBPATH
Linux	LD_LIBRARY_PATH

Stopping and starting the WebSphere DataStage server

You need to stop and start the WebSphere DataStage after you complete certain tasks. For example, you must stop and restart the server after you change or add environment variables or after you complete the suite installation.

Before you begin

You must have root authority on the system to complete this task.

Procedure

To stop and start the WebSphere DataStage server:

Run the following commands.

```
cd $DSHOME  
. ./dsenv  
bin/uv -admin -stop  
bin/uv -admin -start
```

\$DSHOME is the name of the installation directory for the runtime engine. The default directory is /opt/IBM/InformationServer/Server/DSEngine.

Configuring the WebSphere DataStage and WebSphere QualityStage server through PAM

The WebSphere DataStage and WebSphere QualityStage server can be configured to use the local operating system authentication or an external authentication mechanism (such as LDAP) through support for Pluggable Authentication Module (PAM).

You can configure the WebSphere DataStage and WebSphere QualityStage server to use the same user registry as the user registry used by IBM Information Server. If you share the same user registry and that user registry is not using the local operating system authentication services, then you must use the WebSphere DataStage and WebSphere QualityStage server PAM support. When the WebSphere DataStage and WebSphere QualityStage server and IBM Information Server are using the same user registry, no credential mapping is required between the suite users and the WebSphere DataStage and WebSphere QualityStage server users.

Configuring WebSphere DataStage and WebSphere QualityStage for use with PAM (Linux, UNIX)

PAM is currently supported on AIX and Linux platforms.

Before you begin

To complete this task, you must have a working knowledge of PAM and the authentication modules and strategies.

Procedure

To configure PAM on WebSphere DataStage and WebSphere QualityStage:

1. Add to or create the PAM configuration file on your platform.
2. Stop the DataStage server engine:

```
$DSHOME/bin/uv -admin -stop
```

Example PAM configuration files:

On a Linux system, you must create a file named `dsepam` in the `/etc/pam.d` directory. The following example shows the possible contents of the `dsepam` file on a Linux system:

```
##PAM-1.0
auth      required /lib/security/pam_stack.so service=system-auth
password  required /lib/security/pam_stack.so service=system-auth
account   required /lib/security/pam_stack.so service=system-auth
```

Configuring plug-in stages(Linux, UNIX)

WebSphere DataStage provides plug-in stages that are used to access databases and perform other special functions. These plug-in stages are used in server jobs and parallel jobs.

The DataStage installation directory is identified by the `DSHOME` environment variable.

The following table lists the location of the plug-in stage package for each platform.

Table 28. Plugin directories

Operating system	Path
AIX	/cdrom/packages/aix
Linux	/cdrom/PACKAGES/LINUX

Most plug-in stages require you to specify some environment variables in the dsenv file.

BCPLoad stage

The BCPLoad stage enables you to use the bulk copy program (BCP) on Microsoft SQL Server and Sybase databases.

The BCPLoad stage is installed with WebSphere DataStage by default.

The BCPLoad stage uses the BCP API in client libraries. The BCPLoad plug-in does not support MSDBLIB as a client library on any UNIX platform.

If you use a BCPLoad stage with a UNIX server, you must specify SYBCTLIB as client library by setting the Client Library property on the Stage editor Properties page.

Note that the Sybase DBLIB library is no longer supported on AIX platforms. Use the CTLIB library instead.

Configuring NLS for WebSphere DataStage and WebSphere QualityStage (Linux, UNIX)

If you installed NLS but you have not configured it, you need to edit the \$DSHOME/dsenv file to specify the locale that the server uses. This locale is different from the default locales for projects or jobs. Server jobs depend on the server locale to compile correctly.

Procedure

To configure NLS for WebSphere DataStage and WebSphere QualityStage:

1. Locate the following section in the dsenv file.

```
#LANG="<langdef>";export LANG
#LC_ALL="<langdef>";export LC_ALL
#LC_CTYPE="<langdef>";export LC_CTYPE
#LC_COLLATE="<langdef>";export LC_COLLATE
#LC_MONETARY="<langdef>";export LC_MONETARY
#LC_NUMERIC="<langdef>";export LC_NUMERIC
#LC_TIME="<langdef>";export LC_TIME
#LC_MESSAGES="<langdef>"; export LC_MESSAGES
```
2. Replace all occurrences of <langdef> with the locale that the server uses. (the locale must be one that is listed when you use the locale -a command).
3. Remove the number signs (#) at the start of the lines.
4. Stop and restart the WebSphere DataStage and WebSphere QualityStage server.

Configuring the parallel engine (Linux, UNIX)

After you install the parallel engine, you must configure it.

Remote Shells

The parallel engine must be able to run the remote shell command `rsh`, or the secure shell command `ssh`, without a password on all processing nodes running parallel jobs.

Configuring the parallel engine to use ssh:

You can use Secure Shell (`ssh`) with the parallel engine instead of `rsh`. The `ssh` program provides strong authentication and secure communications between computers on a network.

Ensure that `ssh` servers are installed on all of the computers where the parallel engine runs.

About this task

When you configure `ssh` to work with the parallel engine, it connects from the primary computer to all of the other computers using public key encryption for authentication. This task must be completed for each user that runs parallel jobs.

Procedure

To configure the parallel engine to use `ssh`:

1. On the primary computer, run the following command to create public and private RSA keys. Press enter twice when you run the command. The keys are saved in the `~/.ssh/id_rsa` directory.

```
primaryComputer% ssh-keygen -b 1024 -t rsa -f ~/.ssh/id_rsa
```

primaryComputer is the name of the primary computer.

2. This step is required only if the user home directories are not automatically mounted. However, you must still log in to each computer once before you can run parallel jobs. Run the following commands to send the public key to the other computers. The `scp` command is the secure version of the `rcp` command.

```
primaryComputer% cd .ssh
```

```
primaryComputer% scp id_rsa.pub user@secondaryComputer:~/.ssh
```

secondaryComputer is the name of one of the other computers where the parallel engine runs.

3. Log in to each of the secondary computers and add the public key to the list of authorized keys.

```
secondaryComputer% cd .ssh
```

```
secondaryComputer% cat id_rsa.pub >> authorized_keys2
```

```
secondaryComputer% chmod 640 authorized_keys2
```

```
secondaryComputer% rm -f id_rsa.pub
```

The file that contains the authorized keys might be called `authorized_keys` or `authorized_keys2` depending on the version of `ssh` that is installed on your systems.

4. Run the following command to test the `ssh` configuration.

```
primaryComputer% ssh secondaryComputer ls
```

5. On the primary computer, create the `remsh` file in the `/Server/PXEngine/etc/` directory with the following content.

```
#!/bin/sh
exec /usr/bin/ssh "$@"
```

The `/Server/PXEngine/etc/` directory is located in the directory where you installed the engine. The default directory is `/opt/IBM/InformationServer/Server/PXEngine/`.

Specifying the Location of rsh:

To find `rsh` on a processing node, the parallel engine searches the following paths in the following order.

```
install_dir/etc/remsh (if it exists) /user/lpp/ssp/rcmd/bin/rsh (AIX only)
/usr/ucb/rsh /usr/bin/remsh /bin/remsh /usr/bin/rsh
```

where `install_dir` is the top-level directory of your parallel engine installation.

If the parallel engine does not find your `rsh` command, you must specify its location. To do so, copy or rename the supplied file `install_dir/etc/remsh.example` to `install_dir/etc/remsh`. This file contains the following shell script:

```
#!/bin/sh
# Example apt/etc/remsh
exec /usr/bin/rsh "$@"
```

As written, this shell script invokes `/usr/bin/rsh`. Edit the last line of this script to invoke your specific remote shell command. The script should be executable by all users. Use `chmod` to ensure that it is:

```
# chmod 755 script-filename
```

Test this by running `rsh` on each node:

```
# rsh nodename uptime
```

Allowing users to run rsh without a password:

You perform this process differently according to the type of system you are running. For example, you can edit `/etc/hosts.equiv` or create a `.rhosts` file for each user. In both cases, add the host name of each parallel processing node to `/etc/hosts.equiv` or `.rhosts`, one host name per line.

About this task

The host name that is included in this file must correspond to the setting of the node's **fastname** parameter in the parallel configuration file. For information on the **fastname** configuration option, see the *IBM WebSphere DataStage Parallel Job Developer Guide*.

Procedure

To allow users to run `rsh` without a password, use one of the following methods:

- If you choose to edit `/etc/hosts.equiv`, the file must be owned by root and must grant read/write access to root and no access to any other user (file mode of 600).
- If you choose to create an `.rhosts` file for each user, it must:
 - Be located in the home directory of each Parallel user.
 - Be owned by the user.
 - Grant read/write access to the user and no access to any other user (file mode of 600).

To check that users can use rsh without a password, issue the following command on each node:

```
$ rsh hostname uptime
```

hostname is the name of a processing node that you use with the parallel engine. If *hostname* is accessible, this command prints a message displaying the time it has been up.

Copying the parallel engine to multiple computers

To use the parallel engine on multiple computers in a network, you must copy the parallel engine components to each computer and configure the systems.

The user who runs jobs must have rsh or ssh access to each computer where the parallel engine is copied.

Procedure

To copy the parallel engine to multiple computers:

1. After you install the engine on your primary computer, copy the contents of the PXEngine directory to all the other computers. For example, if you installed the engine in /opt/IBM/InformationServer/, the PXEngine directory is in /opt/IBM/InformationServer/Server/PXEngine. The engine must be in the same directory on each computer. You can use a symbolic link.
2. Add entries in the parallel engine configuration file for each of the new computers.

Setting the maximum number of processes for parallel jobs

You must set the maximum number of processes for parallel jobs. A user who runs parallel jobs should be able to run at least 100 processes per node and possibly more on large systems or symmetric multiprocessing (SMP) servers.

Procedure

To set the number of processes for parallel jobs:

1. Determine the current setting for the number of processes, as defined by the setting of the **maxuproc** parameter, by issuing the following command:
node: \$ lsattr -E -l sys0 | grep maxuproc
2. If the value of the **maxuproc** parameter is less than 100, increase it:
 - Set the maximum number of processes on a specific node or workstation, by running the command:
node: # chdev -l sys0 -a maxuproc=100
This command must be issued on the node by a user with root-access privileges.
 - On all nodes of an AIX RS/6000® SP™ system, issue the following dsh command to set the maximum number of processes per user:
control: # dsh -a chdev -l sys0 -a maxuproc=100
This command must be executed from the control workstation with root privileges. On AIX, you can also issue the smit command to perform this action.

Setting the maximum amount of memory for parallel jobs (AIX)

You can set different parameters to control memory resources that improve the performance of parallel jobs. The parallel engine processes consume large amounts of memory. You must have permission to use sufficient resources to run your parallel jobs.

About this task

The following table describes the parameters that control memory resources and recommends a value for each parameter.

Parameter	Description	Recommended value
fsize	The size of the largest file that a user can create	2 GB minus 512 bytes (4,194,303 512-byte blocks)
data	The size of the largest data segment (sometimes called heap) that a program can have	128 MB (262,144 512-byte blocks) or more, or -1 to specify no limit
stack	The size of the largest stack that a program can have	32 MB (65,536 512-byte blocks) or more
rss	The maximum amount of physical memory that a user's process can use	At least 64 MB (131,072 512-byte blocks)

Procedure

To set the maximum amount of memory for parallel jobs:

1. Run the following command to determine the value of the memory size parameters for a user:

```
lsuser -f user
```

This command displays various parameters, including:

```
fsize=2097151 data=262144 stack=65536 rss=131072
```

These numbers represent the number of 512-byte blocks.

2. Set the parameters for a user as defined by the recommended values.

To set **fsize** to 4,194,303 512-byte blocks (2 GB minus 512 bytes), run the following command on a single node or workstation using root access privileges.

```
node: # chuser fsize=4194303 user1
```

3. Run the following command to set the parameter on all nodes of an AIX RS/6000 SP. This command must be issued from the control workstation using root access privileges.

```
control: # dsh -a chuser fsize=4194303 user1
```

After this command runs, users must log out and log in again for it to take effect.

To assign values to the other parameters, use the syntax `control: # dsh -a chuser param=quantity username`, where *param* is the parameter, *quantity* is the size or amount, and *username* is the name of the user.

System paging space for the parallel engine

For parallel jobs that process large amounts of data, your system must have the minimum settings for the paging space.

The minimum amount of paging space to run parallel jobs is 500 MB per processor for systems with fewer than 20 processors and 1 GB per processor for systems with at least 20 processors.

Depending on your parallel job design, you might find your requirements are greater. The more stages an individual job has, the greater the swap space requirement is likely to be. Similarly, you are likely to need more swap space if you intend to run more than one job at one time.

Attention: Only a system administrator with root privileges should reset the system swap space. The procedure is potentially destructive to your processing environment.

Setting the paging space on an AIX system:

For parallel jobs that process large amounts of data, your system must have the minimum settings for the paging space.

Attention: Only a system administrator with root privileges should reset the system paging space.

Procedure

To set the paging space on an AIX system:

1. Run the following command on each node or workstation in your system to determine the current paging size for a node or workstation.

```
node: $ lsps -a
```

The `lsps` command displays information in the form shown in this table.

Page Space	Physical Volume	Volume	Group Size	% Used	Active	Auto	Type
paging01	hdisk1	pagevg	2000MB	1	yes	yes	lv
paging00	hdisk0	rootvg	192MB	5	yes	yes	lv
hd6	hdisk0	rootvg	64MB	23	yes	yes	lv

The Group Size column lists the paging size for each disk drive. In this case, the total for the node is $2000 + 192 + 64$ MB = 2256 MB (approximately 2.2 GB).

You can also run the following `dsh` command from the control workstation of an AIX RS/6000 SP to display paging size information about all nodes:

```
control: $ dsh -a lsps -a
```

2. Change the paging space as necessary for each node or workstation.

Setting the paging space on Linux systems:

For parallel jobs that process large amounts of data, your system must have the minimum settings for the paging space.

Attention: Only a system administrator with root privileges should reset the system paging space. The procedure is potentially destructive.

See the documentation for Red Hat Enterprise Linux for information about setting the paging space for your system.

Editing the /etc/magic file

The file `install_dir/etc/magic` contains additions to `/etc/magic` that are specific to the parallel engine. `install_dir` is the top-level directory of your parallel engine installation. The default path is `/opt/IBM/InformationServer/Server/PXEngine/`. If you choose, you can edit your system's `/etc/magic` by appending the parallel engine descriptors to this file.

You must be logged in as a root user to edit `/etc/magic`.

The UNIX `file` command examines a file to classify it by type. The command reads the file `/etc/magic` to identify files that have a *magic* number. A *magic* number is a constant indicating the file type.

Here is a sample `file` command and sample system output after the edits described above have been made:

```
$ file t1.ds
t1.ds: IBM Information Server parallel engine dataset descriptor
```

Setting environment variables for the parallel engine

You must set various environment variables to ensure smooth operation of your parallel engine. Environment variables are set on a per-project basis from the WebSphere DataStage and WebSphere QualityStage Administrator client.

Procedure

1. Click **Start** → **All Programs** → **IBM Information Server** → **WebSphere DataStage and QualityStage Administrator** and log in to the administrator client.
2. Click the **Project** tab and select a project.
3. Click **Properties**.
4. On the **General** tab, click **Environment...**
5. Set the values for the environment variables as necessary.

Environment variables for the parallel engine

The following environment variables can be set to configure your parallel engine.

Network settings

APT_IO_MAXIMUM_OUTSTANDING

If your system connects to multiple processing nodes through a network, you need to set the `APT_IO_MAXIMUM_OUTSTANDING` environment variable to specify the amount of memory, in bytes, to reserve for the parallel engine on every node for TCP/IP communications. The default value is 2 MB.

If TCP/IP throughput at that setting is so low that there is idle processor time, increment it by doubling it until performance improves. If the system is paging, however, or if your job fails with messages about broken pipes or broken TCP connections, the setting is probably too high.

APT_RECVBUFSIZE

The value of this environment variable specifies the TCP/IP buffer space that is allocated for each connection. If any of the stages within a job has a large number of communication links between nodes, you might need to set this environment variable. The value of this environment variable is specified in bytes.

The `APT_SENDBUFSIZE` and `APT_RECVBUFSIZE` values are the same. If you set one of these environment variables, the other is automatically set

to the same value. These environment variables override the `APT_IO_MAXIMUM_OUTSTANDING` environment variable that sets the total amount of TCP/IP buffer space that is used by one partition of a stage. The `APT_IO_MAXIMUM_OUTSTANDING` environment variable has a default setting of 2 MB.

APT_SENDBUFSIZE

The value of this environment variables specifies the TCP/IP buffer space that is allocated for each connection. If any of the stages within a job has a large number of communication links between nodes, you might need to set this environment variable. The value of this environment variable is specified in bytes.

The `APT_SENDBUFSIZE` and `APT_RECVBUFSIZE` values are the same. If you set one of these environment variables, the other is automatically set to the same value. These environment variables override the `APT_IO_MAXIMUM_OUTSTANDING` environment variable that sets the total amount of TCP/IP buffer space that is used by one partition of a stage. The `APT_IO_MAXIMUM_OUTSTANDING` environment variable has a default setting of 2 MB.

Transform library

If you are working on a non-NFS MPP system you must set the `APT_COPY_TRANSFORM_OPERATOR` environment variable to true to enable Transformer stages to work in this environment. DataStage users must have the appropriate privileges to create project directory paths on all the remote nodes at runtime. This environment variable is set to false by default.

Job monitoring

By default, the job monitor uses time-based monitoring in the WebSphere DataStage and QualityStage Director. The job monitor window is updated every five seconds. You can also specify that the monitoring is based on size. For example, the job monitor window is updated based on the number of new entries. To base monitoring on the number of new entries, set a value for the `APT_MONITOR_SIZE` environment variable. Overriding the default setting for the `APT_MONITOR_TIME` overrides the setting of the `APT_MONITOR_SIZE` environment variable.

Setting transport block size

You can set the `APT_AUTO_TRANSPORT_BLOCK_SIZE` to True to automatically calculate the block size for transferring data internally as jobs run.

To define a block size, use the `APT_DEFAULT_TRANSPORT_BLOCK_SIZE` environment variable. By default, this environment variable is set to 32768.

Producing detailed information about jobs

Set the `APT_DUMP_SCORE` value to True to produce detailed information about jobs as they run. By default, this environment variable is set to False.

C++ compiler

The environment variables `APT_COMPILER` and `APT_LINKER` are set at installation time to point to the default locations of the supported compilers. If

your compiler is installed on a different computer from the parallel engine, you must change the default environment variables for every project by using the Administrator client.

Temporary directory

By default, the parallel engine uses the directory `/tmp` for some temporary file storage. If you do not want to use this directory, assign the path name to a different directory by using the environment variable `TMPDIR`.

User privileges for running parallel jobs

A user who runs a parallel job must have the necessary privileges on all nodes that the application uses.

The following privileges are required for running parallel jobs:

- Login access
- Read, write, and modify access to:
 - `/ishome/Server/Scratch`
 - `/ishome/Server/Datasets`
 - `/ishome/Server/DSEngine/uvtemp`
 - `/tmp`
- Run permissions on local copies of programs and scripts
- Read access to `$APT_ORCHHOME`

Configuring the parallel environment

To develop and run parallel jobs, your development directory must be globally visible to all of the other nodes in the system. The directory must be globally cross-mounted.

Procedure

To configure the parallel environment:

1. Make your development directory visible to all other nodes in the system with the same path name. Typically, you use NFS to make the directory globally visible.
2. If you have a network-cluster system, you are using the Bourne or C shell, and your development directory is not statically mounted or automounted, set the environment variable `PWD` to the globally visible path of the directory. The directory must be visible with the same path name on all processing nodes on which parallel jobs are run. Update the `PWD` environment variable whenever you change to a new directory.

For example, if you use C shell and the directory `/usr/me/parallel-extender/app1` as your development directory, set `PWD` as shown here:

```
$ setenv PWD /usr/me/parallel-extender/app1
```

After you complete this configuration, you can copy the sample applications from the directory `install_dir/examples` to your directory and run the examples

Configuring the parallel engine (Windows)

After you install the suite, you can configure the parallel engine.

Environment variables for the parallel engine

The following environment variables can be set to configure your parallel engine.

Network settings

APT_IO_MAXIMUM_OUTSTANDING

If your system connects to multiple processing nodes through a network, you need to set the APT_IO_MAXIMUM_OUTSTANDING environment variable to specify the amount of memory, in bytes, to reserve for the parallel engine on every node for TCP/IP communications. The default value is 2 MB.

If TCP/IP throughput at that setting is so low that there is idle processor time, increment it by doubling it until performance improves. If the system is paging, however, or if your job fails with messages about broken pipes or broken TCP connections, the setting is probably too high.

APT_RECVBUFSIZE

The value of this environment variable specifies the TCP/IP buffer space that is allocated for each connection. If any of the stages within a job has a large number of communication links between nodes, you might need to set this environment variable. The value of this environment variable is specified in bytes.

The APT_SENDBUFSIZE and APT_RECVBUFSIZE values are the same. If you set one of these environment variables, the other is automatically set to the same value. These environment variables override the APT_IO_MAXIMUM_OUTSTANDING environment variable that sets the total amount of TCP/IP buffer space that is used by one partition of a stage. The APT_IO_MAXIMUM_OUTSTANDING environment variable has a default setting of 2 MB.

APT_SENDBUFSIZE

The value of this environment variable specifies the TCP/IP buffer space that is allocated for each connection. If any of the stages within a job has a large number of communication links between nodes, you might need to set this environment variable. The value of this environment variable is specified in bytes.

The APT_SENDBUFSIZE and APT_RECVBUFSIZE values are the same. If you set one of these environment variables, the other is automatically set to the same value. These environment variables override the APT_IO_MAXIMUM_OUTSTANDING environment variable that sets the total amount of TCP/IP buffer space that is used by one partition of a stage. The APT_IO_MAXIMUM_OUTSTANDING environment variable has a default setting of 2 MB.

Transform library

If you are working on a non-NFS MPP system you must set the APT_COPY_TRANSFORM_OPERATOR environment variable to true to enable Transformer stages to work in this environment. DataStage users must have the appropriate privileges to create project directory paths on all the remote nodes at runtime. This environment variable is set to false by default.

Job monitoring

By default, the job monitor uses time-based monitoring in the WebSphere DataStage and QualityStage Director. The job monitor window is updated every five seconds. You can also specify that the monitoring is based on size. For example, the job monitor window is updated based on the number of new entries. To base monitoring on the number of new entries, set a value for the

APT_MONITOR_SIZE environment variable. Overriding the default setting for the APT_MONITOR_TIME overrides the setting of the APT_MONITOR_SIZE environment variable.

Setting transport block size

You can set the APT_AUTO_TRANSPORT_BLOCK_SIZE to True to automatically calculate the block size for transferring data internally as jobs run.

To define a block size, use the APT_DEFAULT_TRANSPORT_BLOCK_SIZE environment variable. By default, this environment variable is set to 32768.

Producing detailed information about jobs

Set the APT_DUMP_SCORE value to True to produce detailed information about jobs as they run. By default, this environment variable is set to False.

C++ compiler

The environment variables APT_COMPILER and APT_LINKER are set at installation time to point to the default locations of the supported compilers. If your compiler is installed on a different computer from the parallel engine, you must change the default environment variables for every project by using the Administrator client.

Temporary directory

By default, the parallel engine uses the directory /tmp for some temporary file storage. If you do not want to use this directory, assign the path name to a different directory by using the environment variable TMPDIR.

User privileges for running parallel jobs

A user who runs a parallel job must have the necessary privileges on all nodes that the application uses.

The following privileges are required for running parallel jobs:

- Login access
- Read, write, and modify access to:
 - /ishome/Server/Scratch
 - /ishome/Server/Datasets
 - /ishome/Server/DSEngine/uvtemp
 - /tmp
- Run permissions on local copies of programs and scripts
- Read access to \$APT_ORCHHOME

Configuring a computer to be a processing node for MPP (Windows)

WebSphere DataStage can take advantage of massively parallel processing (MPP). With a user ID and a secure shell connection, you can configure a WebSphere DataStage server as a parallel processing node, which can share the processing load with other WebSphere DataStage servers.

You must have administrator access to the installation system to complete this task.

1. Create a generic WebSphere DataStage user account.
 - a. Click **Start** → **Run** → **cmd** to open a Window command prompt. *cmd* is the text to enter in the **Open:** field.
 - b. From the command window, create the generic dstage user using the following format:


```
NET USER dstage password /ADD/COMMENT:"DataStage user account"
/FULLNAME:"DataStage User"/ACTIVE:YES/EXPIRES:NEVER/PASSWORDCHG:NO
```

password is the password assigned to the dstage user. All machines participating in the MPP configuration should use the same user and password.
 - c. Log off via the Start menu and log in as the dstage user.
 - d. Launch the Control Panel from **Start** → **Control Panel** → **System**, click the **Advanced** tab, and click **Environment variables**.
 - e. Under **User variables for dstage**, select **HOME** and click the **Edit** button. If the entry contains other environment variables (for example, %HOMEPATH%) replace the entire contents with the literal value shown in the User variables list for dstage. The default location is typically C:\Documents and Settings\dstage. MKS does not have the ability to resolve these environment variables when connecting remotely.
 - f. Click **OK** three times to save your changes and to exit from the System Applet.
 - g. Log off as the dstage user and log back in as yourself.
2. Set user permissions. With the new account created, you need to set permissions on some critical PXEngine directories so that the new user will have sufficient access rights to execute jobs. The following table lists the minimum directories that you must set and the appropriate permissions.

Directory	Permissions
C:\tmp	Write, Read, Modify
C:\IBM\InformationServer\Server\Scratch	Write, Read, Modify
C:\IBM\InformationServer\Server\Datasets	Write, Read, Modify

- a. Using Explorer, browse to one of the directories, right click **Sharing and Security**.
 - b. From the Security tab and under Group or user names, click **Add**.
 - c. Enter dstage and click **OK**. You should see dstage under the list of users with permissions to this directory.
 - d. Select dstage and select the permission check boxes for the permissions Write, Read, and Modify.
 - e. Click **Apply**, and then click **OK** to close the window.

Repeat these step for each directory that you want to set permissions for.
3. Set remote shell permissions
 - a. Click **Start** → **Run** → **ksh**. *ksh* is the text to enter in the **Open:** field.
 - b. From the shell, run the ssh command to connect to your machine. The format is `ssh -l dstage host name`. *host name* is the name of your computer. You will most likely have to accept an rsa key. You must enter the password for the dstage user.
 - c. Once you are connected, run the rsetup command to add the dstage user and password to the Windows XP/2003 LSA database for use by the rshd service. Enter the same password when prompted.

- d. Edit the .rhosts file in the dstage user's home directory. You might need to create the .rhosts file if it does not already exist.
 - e. Add the line + dstage.
 - f. Exit from the ssh session.
 - g. Edit the hosts.equiv file located in the \$ROOTDIR/etc folder. The path name is shown in UNIX format, which is required by MKS. The default location is C:\Program Files\MKS Toolkit\etc\hosts.equiv.
 - h. For testing purposes, add the line + + You might need to create this file if it does not already exist.
 - i. Test your configuration by running `rsh -l dstage host name ls`. *host name* is the name of your computer. This gives you a listing of the dstage user's home directory.
4. Configure PXEngine for RSH/SSH. To run PXEngine parallel jobs on remote hosts, you need to use a remote shell script. A sample script named `remsh.example` is provided for this purpose. The script is located in the PXEngine\etc folder.
 - a. Copy the script file, and rename it as `remsh`.
 - b. Edit the script file and change the `rsh` line. Save your changes.

Original script file	Changed script file
#!/bin/sh # Example \$APT_ORCHHOME/etc/remsh rsh "\$@"	#!/bin/sh # Example \$APT_ORCHHOME/etc/remsh rsh -l dstage "\$@"

5. Configure the distribute-component script file, which is located in the PXEngine\etc directory. This script is responsible for managing the installation and removal of PXEngine component binaries on remote computing nodes. This script is used when jobs containing compiled PX transforms need to have their transformer binaries copied to the remote nodes.
 - a. Make a backup copy of the distribute-component script file.
 - b. Edit the script file and change the `theUser` line.

Original script file	Changed script file
# The user (optional) # theUser=	# The user (optional) # theUser=dstage

- c. Save your changes.

You can now run PXEngine parallel jobs in an MPP environment.

Setting up database connectivity

The following topics describe the steps you need to take to connect to various databases from a parallel job.

Configuring WebSphere DataStage to run on DB2 servers

WebSphere DataStage should be configured to run on all processing nodes functioning as DB2 servers.

Procedure

To configure WebSphere DataStage to run on DB2 servers:

Include node definitions for your DB2 server nodes in your configuration file.

See the *IBM WebSphere DataStage Parallel Job Developer Guide* for details about the configuration file and the resource DB2 option.

Running the DB2 Setup program:

You must run the DB2 Setup program from the Linux or UNIX command line to configure WebSphere DataStage to access a DB2 database.

Procedure

To run the DB2 Setup program:

For example, the following command calls `db2setup.sh` to configure DataStage to access the database `db2_8`:

```
$ db2setup.sh db2_8
```

Run the setup program from a Linux or UNIX command line. Run the setup program once for each DB2 database to be accessed by WebSphere DataStage users. Pass the database name as an argument.

```
$APT_ORCHHOME/bin/db2setup.sh
```

Granting user privileges to run DB2 Enterprise stage jobs:

You must grant privileges to each user who runs jobs containing the DB2 Enterprise stages.

Before you begin

You need DBADM privileges to complete this task.

Procedure

To grant user privileges for running DB2 Enterprise stage jobs:

Run the `$APT_ORCHHOME/bin/db2grant.sh` script once for each user. Use the following syntax.

```
$ db2grant.sh database_name user_name
```

Accessing a remote DB2 server from DB2 Enterprise stages:

You connect from a DB2 Enterprise stage to a remote DB2 server. The connection is made through a DB2 client.

About this task

To remotely connect from a DB2 client to a DB2 server, the DB2 client must be located on the same machine as the WebSphere DataStage server. Both DB2 client and DB2 server need to be configured for remote connection communication. Ask your database administrator for assistance.

Procedure

The WebSphere DataStage configuration file must contain the node on which WebSphere DataStage and the DB2 client are installed and the nodes of the remote computer where the DB2 server are installed.

On the DB2 Enterprise stage in your parallel job, you need to set the following properties:

- **Client Instance Name.** Set this to the DB2 client instance name. If you set this property, WebSphere DataStage assumes you require remote connection.
- **Server.** Optionally set this to the instance name of the DB2 server. Otherwise use the DB2 environment variable, DB2INSTANCE, to identify the instance name of the DB2 server.
- **Client Alias DB Name.** Set this to the DB2 client's alias database name for the remote DB2 server database. This is required only if the client's alias is different from the actual name of the remote server database.
- **Database.** Optionally set this to the remote server database name. Otherwise use the environment variables APT_DBNAME or APT_DB2DBDFT to identify the database.
- **User.** Enter the user name for connecting to DB2. This is required for a remote connection.
- **Password.** Enter the password for connecting to DB2. This is required for a remote connection.

You can use WebSphere DataStage's remote connection facilities to connect to a different DB2 server within the same job. You could, for example, read from a DB2 database on one server, use this data to access a lookup table on another DB2 server, then write any rejected rows to a third DB2 server. Each database would be accessed by a different stage in the job with the Client Instance Name and Server properties set appropriately.

Performing System Configuration of DB2:

DB2 can be configured in many ways on your system. The DB2 administrator typically follows this procedure to configure DB2 and WebSphere DataStage to communicate with one another.

1. Install DB2 and make sure it is working properly.
2. Grant the WebSphere DataStage users SELECT privileges on the system tables syscat.nodegroupdef, syscat.tablespace, and syscat.tables.
3. Make the file *db2nodes.cfg* readable by the DataStage administrative user.
4. WebSphere DataStage runs many processes for each job. This can require the system administrator to modify DB2 resources, configuration parameters, and manager configuration parameters of your system. See the DB2 administration manuals for more information.
5. The DB2 Enterprise stage requires that users invoking it in load mode have DBADM privilege on the DB2 database written to by the stage. Among the ways to grant this privilege is to start DB2, connect to a database, and grant privileges as follows:

```
db2> CONNECT TO database_name
db2> GRANT DBADM ON DATABASE TO USER user_name
```

Performing User Configuration of DB2:

To access WebSphere DataStage, WebSphere DataStage parallel jobs connect to DB2 using an AIX user name. Therefore, WebSphere DataStage users who invoke DB2 read operators must have SELECT privilege on the table they access, and users who invoke DB2 write operators must have INSERT privileges on the tables they access.

The following procedure configures an Orchestra® user to access DB2. This procedure can be performed by either a system administrator or a WebSphere DataStage user.

1. Set the environment variable DB2INSTANCE to the user name of the owner of the DB2 instance. DB2INSTANCE determines the location of db2nodes.cfg, as follows:

```
~owner_user_name/sqlllib/db2nodes.cfg
```

For example, if you set DB2INSTANCE to Jan, the location of db2nodes.cfg is

```
~Jan/sqlllib/db2nodes.cfg.
```

The file *db2nodes.cfg* contains information used to translate DB2 node numbers to host names. The information in this file must agree with the host names specified in the WebSphere DataStage configuration file.

2. Optionally, set the environment variable DB2DBDFT to the name of the DB2 database that you want to use.
3. You must source the file *db2profile* and update your LIBPATH environment variable to include the location of the DB2 library containing your instance of DB2.

In the following example, the owner of the DB2 instance is the user Mary, the location of db2profile is *~Mary/sqlllib/db2profile*, and the location of the DB2 library is *~Mary/sqlllib/lib*.

For Korn and Bourne shells, include the following statements in your .profile or .kshrc:

```
. ~Mary/sqlllib/db2profile
export LIBPATH=~Mary/sqlllib/lib:${LIBPATH:-/usr/lib}
```

For C shell, put these lines in your .cshrc:

```
source ~Mary/sqlllib/db2profile
if (! $?LIBPATH) setenv LIBPATH /usr/lib
setenv LIBPATH ~Mary/sqlllib/lib:$LIBPATH
```

Setting up access to Oracle databases

After you install the suite, you need to configure access to your data sources. Configuring access to Oracle databases includes granting the appropriate access level and running jobs.

Granting access to the Oracle parallel server:

To access the Oracle parallel server, users must have SELECT access to the sys.gv_\$instance and sys.v_\$cache tables. This task is optional.

Procedure

To grant access to the Oracle parallel server:

Run the following SQL statements.

```
GRANT select ON sys.gv_$instance TO public;
GRANT select ON sys.v_$cache TO public;
```

If you do not have Oracle OPS on these platforms, you should set the APT_ORACLE_NO_OPS environment variable to disable OPS checking on the Oracle Enterprise stage.

Running WebSphere DataStage jobs using the Oracle Enterprise Stage:

Oracle 9i, Enterprise Edition, or later is required to run WebSphere DataStage jobs using the Oracle Enterprise Stage. Your Oracle installation should include the Oracle Database Utilities and Oracle Network software.

Procedure

To run WebSphere DataStage jobs using the Oracle Enterprise Stage:

1. Create the user defined environment variable ORACLE_HOME and set it to the \$ORACLE_HOME path. For example, /disk3/oracle10.
2. Create the user defined environment variable ORACLE_SID and set it to the correct service name. For example, ODBC SOL.
3. Add ORACLE_HOME/bin to your PATH and ORACLE_HOME/lib to your LIBPATH, LD_LIBRARY_PATH, or SHLIB_PATH.
4. Obtain login privileges to Oracle using a valid Oracle user name and corresponding password. The user name and password must be recognized by Oracle before you attempt to access it.
5. Obtain the SELECT privilege on:
 - DBA_EXTENTS
 - DBA_DATA_FILES
 - DBA_TAB_PARTITIONS
 - DBA_TAB_SUBPARTITIONS
 - DBA_OBJECTS
 - ALL_PART_INDEXES
 - ALL_PART_TABLES
 - ALL_INDEXES
 - SYS.GV_\$INSTANCE (Only if Oracle Parallel Server is used)APT_ORCHHOME/bin must appear before ORACLE_HOME/bin in your PATH.
6. Create a role that has the appropriate SELECT privileges. For example,

```
CREATE ROLE DSXE;
GRANT SELECT on sys.dba_extents to DSXE;
GRANT SELECT on sys.dba_data_files to DSXE;
GRANT SELECT on sys.dba_tab_partitions to DSXE;
GRANT SELECT on sys.dba_tab_subpartitions to DSXE;
GRANT SELECT on sys.dba_objects to DSXE;
GRANT SELECT on sys.all_part_indexes to DSXE;
GRANT SELECT on sys.all_part_tables to DSXE;
GRANT SELECT on sys.all_indexes to DSXE;
```
7. After the role is created, run the following command to grant it to users who run WebSphere DataStage jobs.

```
GRANT DSXE to oracle userid;
```

Configuring access to informix XPS databases

Before you begin

You must have the correct privileges and settings in order to use the Informix® Enterprise Stage. You must have a valid account and appropriate privileges on the databases to which you connect.

You require read and write privileges on any table to which you connect, and Resource privileges for using the Partition Table property on an output link or using create and replace modes on an input link.

Procedure

To configure access to Informix XPS databases:

1. Make sure that Informix XPS is running.
2. Make sure the INFORMIXSERVER is set in your environment. This corresponds to a server name in sqlhosts and is set to the coserver name of coserver 1. The coserver must be accessible from the node on which you invoke your DataStage job.
3. Make sure that INFORMIXDIR points to the installation directory of your INFORMIX server.
4. Make sure that INFORMIXSQLHOSTS points to the sql hosts path (e.g., */disk6/informix/informix_runtime/etc/sqlhosts*).

Using Informix XPS stages (AIX):

To run jobs containing Informix XMS stages on AIX systems, you need to have the Informix client sdk 2.81 version installed along with the Informix XPS server. The LIBPATH order should be set as follows:

```
LIBPATH=$APT_ORCHHOME/lib:$INFORMIXDIR/lib:'dirname $DSHOME'/branded_odbc/  
lib:$DSHOME/lib:$DSHOME/uvd11s:$DSHOME/java/jre/bin/classic:$DSHOME/java/jre/  
bin:$INFORMIXDIR/lib:$INFORMIXDIR/lib/cli:$INFORMIXDIR/lib/esql
```

Configuring access to ODBC data sources for WebSphere DataStage and WebSphere QualityStage (Windows)

If you installed ODBC drivers, you need to configure them to access data sources.

Procedure

To configure access to ODBC data sources: See the appropriate documentation:

- For drivers, see the vendor documentation.
- For the federated ODBC wrapper, see the topic, Adding ODBC data sources to a federated server, at publib.boulder.ibm.com/infocenter/db2luw/v9/topic/com.ibm.websphere.ii.foundation.conn.fw.odbc.doc/configuring/tlsodb01.html

Configuring access to ODBC data sources (Linux, UNIX)

After you install IBM Information Server, you need to configure the WebSphere DataStage ODBC connections.

About this task

You need to edit three files to set up the required ODBC connections. These are:

- dsenv
- .odbc.ini
- uvodbc.config

All three are located in the \$DSHOME directory. Copies of uvodbc.config are also placed in the project directories.

Non-wire drivers require different setup information than wire drivers. Non-wire drivers require information about the location of the database client software. Wire drivers require information about the database itself.

For information about configuring the ODBC environment for your specific database, see the Data Direct Drivers Reference manual `odbceref.pdf` file located in the `$DSHOME/Server/branded_odbc/books/odbc` directory. You should also check the `ODBCREAD.ME` file in the `branded_odbc` directory. There is also an html file located in `branded_odbc/odbchelp`.

dsenv file

The WebSphere DataStage server has a centralized file for storing environment variables called `dsenv` in `$DSHOME`. `$DSHOME` identifies the WebSphere DataStage installation directory. The default directory is `/opt/IBM/InformationServer/Server/DSEngine`.

You might need to add new environment variables as you configure WebSphere DataStage to connect to different databases using plug-ins or ODBC drivers. Any environment variables you need for interactive use of ODBC drivers to make a connection to an ODBC data source must be added to the `dsenv` file. This lets the WebSphere DataStage server inherit the proper environment for ODBC connections.

For a connection using a wire protocol driver, there are no changes required to `dsenv`.

For a connection using a non-wire protocol driver, you generally need to specify the following information in the `dsenv` file:

- environment variables required by the database client software
- database home location
- database library directory
- the `PATH` environment variable

The following table shows the library path environment variables for each operating system.

Table 29. Library path environment variables for WebSphere DataStage

Operating system	Environment variables
AIX	<code>LIBPATH</code>
Linux	<code>LD_LIBRARY_PATH</code>

The following example shows some typical entries for commonly used databases.

```
#Sybase 11 setup
LANG=
Export LANG
SYBASE=/export/home/sybase/sybase;export SYBASE
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$SYBASE/lib:/usr/lib:/lib;export LD_LIBRARY_PATH
# Oracle 9i
ORACLE_HOME=/space/oracle9i
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib:$ORACLE_HOME/rdbms/lib;export
LD_LIBRARY_PATH
ORACLE_SID=WSMK5
export ORACLE_HOME ORACLE_SID
#Informix XPS 9.3
INFORMIXDIR=/space/informix; export INFORMIXDIR
INFORMIXSERVER=solnew2.1;export INFORMIXSERVER
ONCONFIG=onconfig.solnew2; export ONCONFIG
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$INFORMIXDIR/lib:$INFORMIXDIR/lib/
esql:$INFORMIXDIR/lib/cli;export LD_LIBRARY_PATH
LANG=C
```

```
export LANG
#DB2 Version 9
DB2DIR=/opt/IBM/InformationServer/DB2;export DB2DIR
DB2INSTANCE=DB2inst1; export DB2INSTANCE
INSTHOME=/export/home/DB2inst1;export INSTHOME
PATH=$PATH:$INSTHOME/sql1lib/bin:$INSTHOME/sql1lib/adm:$INSTHOME/sql1lib/misc
export PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$INSTHOME/sql1lib/lib;export LD_LIBRARY_PATH
THREADS_FLAG=native;export THREADS_FLAG
```

Stop and restart the WebSphere DataStage server whenever you edit dsenv.

.odbc.ini file

The .odbc.ini files gives information about connecting to the database (wire protocol drivers) or the database client (non-wire protocol drivers). If your system uses a mix of drivers, your .odbc.ini file will contain a mix of entry types.

Configuration examples for the various platforms are provided in \$DSHome/branded_odbc/IBM_Tools in the odbc.ini file.

- aix_4.2.odbc.ini for AIX systems.
- linux_4.2.odbc.ini for Linux systems.

These files give the minimum parameter settings required to create a working odbc.ini file. They each give configurations for the following databases:

- DB2 on Linux, UNIX, Windows, and AS400
- Oracle on Linux and UNIX
- Informix on Linux, UNIX, and Windows
- Microsoft SQL Server 2000 connection
- Sybase on Windows

uvodbc.config file

Use the uvodbc.config file to specify the DSNs for the databases that you are connecting to through ODBC.

Entries in uvodbc.config have the following form:

```
<name>
```

```
DBMSTYPE = ODBC
```

name is the ODBC data source name (DSN), this should be the same name as specified in the .odbc.ini file. The spaces on either side of the equal sign are required.

The following example shows a typical uvodbc.config file:

```
[ODBC DATA SOURCES]

<localuv>
DBMSTYPE = UNIVERSE
network = TCP/IP
service = uvserver
host = 127.0.0.1
<Sybase1>
DBMSTYPE = ODBC
<Sybase2>
DBMSTYPE = ODBC
<Oracle8>
DBMSTYPE = ODBC
```

```
<Informix>  
DBMSTYPE = ODBC  
<DB2>  
DBMSTYPE = ODBC
```

A copy of the `uvodbc.config` file is also placed in each WebSphere DataStage project directory. For example, the default path for projects on Linux is `/opt/IBM/WDIS/Server/Projects/`. This is useful where you configure a data source that is known to some projects but not others. By default, WebSphere DataStage searches the current project directory for a `uvodbc.config` file and, if it finds one, uses this in preference to the file in `$DSHOME`. If you alter `uvodbc.config` after creating projects you should copy the edited file to the project directories.

Creating and binding a DB2 package to a DSN

Before you connect to a DB2 DSN you need to bind DB2 packages to the DSN. You must bind a package to every DSN you are going to use to connect to a DB2 database.

To create and bind a package to a DSN:

1. Open a terminal window and change to the `$DSHOME/../../branded_odbc/lib` directory.
2. Enter the command `bind20dsn`.
dsn is the ODBC data source name.
You are prompted for a user ID and password if one is required.

Configuring third party ODBC drivers

You can use ODBC drivers from third parties for connecting WebSphere DataStage to different data sources. You might require to connect to a database not supported by the Data Direct driver pack, for example, Redbrick. You can use only one ODBC driver manager at a time. As a result, you cannot use third party drivers at the same time that you use the WebSphere DataStage drivers, and you must perform some reconfiguration to switch between the two types.

Trial versions of the ODBC drivers are not supported.

Contact IBM support for assistance if you want to use an ODBC driver on an AIX platform for which you do not have a DataDirect driver.

To configure a third party ODBC driver:

1. Install the ODBC manager and drivers on your Linux or UNIX platform following the instructions supplied by the driver vendor.
2. Configure WebSphere DataStage to connect to the driver.
 - a. Add the required environment variables to the `dsenv` file.
 - b. Set up the connection to the database client or database in the `.odbc.ini` file. You must specify DSN, driver name, and log in information as a minimum.
 - c. Specify the DSN for connecting to the database in the `uvodbc.config` file.
3. Include the shared library `libodbc.xx` for the third party ODBC Manager in the WebSphere DataStage load library path. For the suffix *xx*, use `sl` for all platforms. Do this by editing `dsenv` and ensuring that the path to the third party Manager precedes that to `$DSHOME/../../branded_odbc/lib` in the shared library path environment variable.
4. After you have installed and configured your driver, test the connectivity.

To switch back to using Data Direct drivers, you must reverse the procedure in Step 3 and ensure that the WebSphere DataStage ODBC driver library precedes any third party driver libraries.

Testing ODBC driver connectivity:

After you configure ODBC connectivity, you can test your configuration.

Procedure

To test the connectivity of the ODBC connections that you have defined for WebSphere DataStage:

1. Change directories to \$DSHOME and set up the WebSphere DataStage environment by running the following command:
`. ./dsenv`
2. Start the WebSphere DataStage server shell:
`./bin/dssh`
The server shell starts.
3. Log on to the project:
`LOGTO project_name`
Where *project_name* is case sensitive.
4. Get a list of available DSNs by typing:
`DS_CONNECT`
5. Test the required connection by typing:
`DS_CONNECT DSN`
Where *DSN* specifies the connection you want to test. Enter a user name and password as requested, you should then be connected to the specified database.
6. Enter `.Q` to quit the connection.

Configuring user directories for IBM Information Server

User directories contain lists of valid users and user credentials. Any user that logs in to IBM Information Server must use a user name and password from the user registry.

IBM Information Server leverages the WebSphere Application Server user registry support for directory access. IBM Information Server uses the user registry that is configured in WebSphere Application Server.

During installation, you can select one of the following user directories:

IBM Information Server user directory

An IBM Information Server internal user registry. Suite users and suite groups are stored in the repository. You can use the console or the Web console to create, update, or delete suite users and suite groups. In the console and Web console, you can view the available users and groups and assign suite and suite component security roles to these users.

Local operating system directory

Suite users and suite groups are taken from the operating system on which WebSphere Application Server is installed. You can only create and delete users and groups from the operating system tools. In the console and Web console, you can view the available users and groups and assign suite and suite component security roles to these users.

You can change to a different user registry after you install. You can also configure any other LDAP-compliant user registry.

To configure the user directories, you must perform two configuration tasks:

1. Configure the user registry in WebSphere Application Server.
2. Configure the IBM Information Server directory to match the WebSphere Application Server user registry.

You should only change user registries as a post-installation task and not on a system that has been in production. For production systems, you might want to migrate to a new installation to avoid any security issues and risks. Otherwise, there might be a mismatch between the users of the old and new directories.

Related tasks



Using specific directory servers as the LDAP server

Configuring the local operating system user registry

You can configure the local operating system user registry after installation.

Configuring WebSphere Application Server to use the local operating system user registry

You configure the application server to use the local operating system user registry by selecting the local operating system registry option during the installation or by configuring it after the installation.

When you install the suite, you specify whether to use the IBM Information Server user registry (internal user registry) or the local operating system registry. If you choose the internal registry during the installation, you can configure the suite to use the local operating system user registry after the installation is complete. You might want to change user directories if you want user credentials verified against the local operating system or if you test your installation with the internal registry and want to use a local operating system in production. Refer to the WebSphere Application Server documentation topic Local operating system user registries for more information.

1. Start the WebSphere Application Server administration console.
2. In the administration console, click **Security** → **Global Security**.
3. In the list of user registries, click **Local OS**.
4. Type the WebSphere Application Server administrator credentials in the **Server ID** and **Server user password** fields. The credentials that you specify must match an existing user account on the local operating system of the computer where the application server is running. If you want to use an account in a domain where your computer is registered, see the WebSphere Application Server documentation for information.
5. Click **OK**.
6. Click the **Save** link at the top of the page, and click the **Save** button.
7. On the Global security page, ensure that **LTPA** is selected for the **Active authentication mechanism** setting.
8. Select **Local OS in the Active user registry**, and click **OK**. If an error occurs, the application server is unable to authenticate with the local operating system using the credentials that were provided on the Local OS page.
9. Click the **Save** link, and click the **Save** button.

10. Before the application server is restarted, configure the IBM Information Server directory service.

Configuring the IBM Information Server directory service to use the application server user registry

After you configure the WebSphere Application Server user registry to use the local operating system registry, you must configure the IBM Information Server directory service to use the application server user registry.

Even if the IBM Information Server directory service is already configured to use the application server user registry, you must complete this task to ensure that the default IBM Information Server administrator is properly configured.

Procedure

1. Log in to the Web console with suite administrator credentials.
 2. On the Administration tab, click **Domain Management** → **Directory Configuration**. A list of supported directory providers is shown. The directory provider that is being used is selected.
 3. Select **WebSphere User Registry** from the list.
 4. Click **Set as Active Provider**.
 5. Type a user name using the following format:
 - Windows: *DOMAIN\userid*. *DOMAIN* is the name of the computer in uppercase letters and *userid* is a user ID that exists in the local operating system user registry on the computer where the domain is installed. For a computer that is registered on a domain, ensure that the user ID does not exist in the domain and the local operating system user registry.
 - UNIX: *userid*
- The user name that you enter will be the default IBM Information Server administrator user name after the directory is configured. The user name is not verified when you type it, so ensure that it is typed accurately. You can change the suite administrator or add suite administrators after the configuration is complete.
6. Ensure that the clean the security registry option is selected before the WebSphere User Registry is set as active. Cleaning the security registry ensures that none of the security settings from the previous registry can accidentally be used with the new registry and cause unexpected security threats.
 7. Click **OK** to save your changes and continue.
 8. Click **OK** again.
 9. Restart the application server to complete the configuration. When you restart the application server, verify that the application server log files have no errors. Depending on your previous directory configuration, it is possible that the application server cannot be stopped using the standard stop commands. In this case, a manual stop of the WebSphere Application Server Java process is required.

Configuring the IBM Information Server user directory

You can configure the IBM Information Server user directory during installation or after installation. The IBM Information Server user directory stores users and groups in the metadata repository.

Configuring the IBM Information Server directory service to use the internal user registry

To use the internal user registry, you must configure the IBM Information Server directory service before you configure the application server to use the internal user registry.

Procedure

1. Start the IBM Information Server Web console.

The Web console is installed with the domain. If the domain is on Windows, click **Start** → **IBM Information Server** → **IBM Information Server Web Console**. If the domain is on Linux or UNIX, open Microsoft Internet Explorer on a Windows computer on the same network and go to URL of the Web console. The URL for the Web console is `http://host_name:port`. *host_name* is the URL or IP address of the computer where the domain is installed and *port* is the port number assigned to the Web console. The default port number is 9080.

2. Log in to the Web console with suite administrator credentials.
3. On the Administration tab, click **Domain Management** → **Directory Configuration**. A list of supported directory providers is shown. The directory provider that is being used is selected.
4. Select **Information Server Directory** from the list.
5. Click **Set as Active Provider**.
6. Type a user name and password. The credentials that you enter are the default IBM Information Server administrator credentials that should be used after the directory configuration change is completed. You can change the suite administrator or add suite administrators after the configuration is complete.
7. Ensure that the clean the security registry option is selected before the WebSphere User Registry is set as active. Cleaning the security registry ensures that none of the security settings from the previous registry can accidentally be used with the new registry and cause unexpected security threats.
8. Click **OK** to save your changes and continue.
9. Click **OK** again.
10. Before you restart the application server, configure the WebSphere Application Server user registry to use the IBM Information Server user directory.

Configuring the WebSphere Application Server to use the internal user registry

You configure the application server to use the internal user registry.

Complete the task Configuring the IBM Information Server directory service to use the internal user registry before you complete this task.

About this task

A WebSphere Application Server custom user registry is installed when you install the suite. The custom registry must be configured as the active user registry for the application server for it to work with the IBM Information Server user registry.

Procedure

To configure the application server to use the internal user registry:

1. Start the WebSphere Application Server administration console.

2. In the administration console, click **Security** → **Global Security**.
3. In the list of user registries, click **Custom User Registry**.
4. Enter the WebSphere Application Server administrator credentials in the **Server ID** and **Server user password** fields. The credentials that you specify must match an existing user account in the custom registry. The only user ID and password that is available in the custom registry is the one that you created when you completed the task “Configuring the IBM Information Server directory service to use the internal user registry” on page 90.
5. Ensure that the custom registry class name is `com.ascential.acs.security.directory.custom.websphere.ASBUserRegistry`.
6. Click **OK**.
7. Click the **Save** link at the top of the page, and click the **Save** button.
8. On the Global Security page, ensure that LTPA is selected for the **Active authentication mechanism** setting.
9. Select **Custom user registry in the Active user registry**, and click **OK**. If an error occurs, the application server is unable to authenticate with the local operating system using the credentials that were provided under Custom user registry on the Global Security page. Ensure that you specified the same credentials under **Custom User Registry** as you specified in the task Configuring the IBM Information Server directory service to use the internal user registry.
10. Click the **Save** link, and click the **Save** button.
11. Restart the application server.
12. When you restart the application server, verify that the application server log files have no errors.

Configuring LDAP for user registries

You can configure any LDAP-compliant user registry for IBM Information Server.

To configure LDAP for user registries:

1. Configure WebSphere Application Server for LDAP.
2. Start the IBM Information Server Web console.

The Web console is installed with the domain. If the domain is on Windows, click **Start** → **IBM Information Server** → **IBM Information Server Web Console**. If the domain is on Linux or UNIX, open Microsoft Internet Explorer on a Windows computer on the same network and go to URL of the Web console. The URL for the Web console is `http://host_name:port`. *host_name* is the URL or IP address of the computer where the domain is installed and *port* is the port number assigned to the Web console. The default port number is 9080.

3. Log in to the Web console with suite administrator credentials.
4. On the Administration tab, click **Domain Management** → **Directory Configuration**. A list of supported directory providers is shown. The directory provider that is being used is selected.
5. Select **WebSphere User Registry** from the list.
6. Click **Set as Active Provider**.
7. Type a user name. The format of the user name must be a LDAP full distinguished name, for example `cn=adminUsername, cn=users, dc=ibm, dc=com`, and the casing must match the casing used by the application server.
8. Ensure that the clean the security registry option is selected before the WebSphere User Registry is set as active. Cleaning the security registry

ensures that none of the security settings from the previous registry can accidentally be used with the new registry and cause unexpected security threats.

9. Click **OK** to save your changes and continue.
10. Click **OK** again.
11. Restart the application server to complete the configuration. When you restart the application server, verify that the application server log files have no errors. Depending on your previous directory configuration, it is possible that the application server cannot be stopped using the standard stop commands. In this case, a manual stop of the WebSphere Application Server Java process is required.

Configuring WebSphere Federation Server

If you installed WebSphere Federation Server, you need to configure the federated server, federated database, and access to the data sources in your enterprise.

About this task

This task describes the steps for configuring the federated server and database. For information about configuring access to the federated data sources, see the documentation for WebSphere Federation Server.

Procedure

1. Check the setup of the federated server
2. Create a federated database

Checking the setup of the federated server

You can avoid potential configuration problems by checking the key settings on the federated server.

When you install WebSphere Federation Server, the software attempts to configure the federated server for you. If problems occur during the installation, the federated server is not setup properly.

About this task

To check the setup of the federated server:

1. Confirm the link-edit of the wrapper library files to the data source client software (UNIX).
2. Check that the **FEDERATED** parameter is set to **YES**.

After you check the setup of the federated server, you must create a federated database.

Confirming the link-edit of the wrapper library files (UNIX)

On federated servers that run UNIX, some wrappers must be link-edited with the client software for the data source.

About this task

The link-edit step is attempted when you install WebSphere Federation Server. The link-edit step creates a wrapper library for each data source that the federated server will communicate with.

This task applies to only the following data sources:

- Informix
- Microsoft SQL Server
- Oracle
- Sybase
- Teradata

Before you configure the federated server and database to access data sources, you should confirm that the link-edit of the wrapper library files was successful.

Procedure

To confirm the link-edit of the wrapper library files:

Determine the status of the link-edit:

- If the link-edit was successful, the wrapper library file appears in the directory where WebSphere Federation Server is installed.
- If the link-edit failed, check the error message file in the directory where WebSphere Federation Server is installed.
- If the link-edit was not performed, neither the library file nor the message file appears in the directory where WebSphere Federation Server is installed. You will have to manually run the link script.

Checking the wrapper library files (UNIX):

You must confirm that a wrapper library file exists on the federated server for each data source that you want to access.

About this task

This task applies to only the following data sources:

- Informix
- Microsoft SQL Server
- Oracle
- Sybase
- Teradata

Procedure

To check for the wrapper library files:

Check for the library files in the default directory path where WebSphere Federation Server is installed. If the library files are not in that directory, you must manually link the wrapper libraries to the data source client software.

The directory path for the wrapper library depends on the data source.

Wrapper library files:

The wrapper library files are required so that you can access the data sources.

For some data sources, the wrapper library files are added to the federated server when you install WebSphere Federation Server. For other data sources a link-edit script must be run to create the library files. There are three wrapper library files for each data source.

The wrapper library files are required when you register the wrapper for the data source.

The wrapper library files are added to the default directory path. There should be a set of wrapper library files for each of the data sources that you want to access.

Checking the link-edit message files (UNIX):

If the link-edit fails, errors are listed in the message file in the library directory.

About this task

The existence of a message file does not mean that the link-edit failed. There is a message file in the library directory even if the link-edit is successful.

You must open the message file to determine if the link-edit failed.

This task is required for only the following data sources:

- Informix
- Microsoft SQL Server
- Oracle
- Sybase
- Teradata

Procedure

To determine if the link-edit failed:

1. Open the link-edit message files. The link-edit message files are in the directory where WebSphere Federation Server is installed, in the lib32 or lib64 subdirectory.

The names of the link-edit message files are listed in the following table.

Table 30. Link-edit message file names by data source

Data source	Message file names
Informix	djxlinkInformix.out
Microsoft SQL Server	djxlinkMssql.out
Oracle	djxlinkOracle.out
Sybase	djxlinkSybase.out
Teradata	djxlinkTeradata.out

2. Resolve the link-edit failure.

There are several reasons why the link might fail when you setup the federated server:

- If the data source client software is not installed before the link-edit is attempted, then the link-edit will fail. For example, if you do not install the Informix client software before you install the DB2 server software, the link-edit will fail. Likewise, if you do not install the Sybase Open Client

software before you install WebSphere Federation Server, the link-edit will fail. In these situations, you will have to perform the link manually.

- Verify that the version of the data source client software is supported. If the version of the data source client software that you have installed is not supported, the link-edit will fail. You will have to install a client version that is supported and then perform the link manually.

Manually linking the wrapper libraries to the data source client software:

If the wrapper library files are not in the directory path, you must manually link the wrapper libraries.

Before you begin

You need root authorization to run the link scripts.

The client software for the data sources that you want to access must be installed and configured on the federated server.

For the `djxlinkxxx` scripts to issue their messages in your language, at least one database instance must exist on the federated server. If an instance does not exist on the federated server, the scripts will still work. However, the scripts will issue all messages in English.

About this task

This task applies only to the following data sources:

- Informix
- Microsoft SQL Server
- Oracle
- Sybase
- Teradata

Procedure

To link the wrapper libraries to the data source client software:

1. Decide which method you want to use to perform the link:

Method	Step
Use the product CDs.	• For Informix data sources, run the DB2 for Linux, Unix, and Windows installation again and specify the Custom installation option. • For Microsoft SQL Server, Oracle, Sybase, and Teradata data sources run the WebSphere Federation Server installation again. From the launchpad, click Install Products and follow the instructions in the wizard.

Method	Step
Run the link-edit scripts.	1. Open a UNIX command prompt and run the link-edit script for each data source that you want to access. The names of the link-edit script are: • djxlinkInformix • djxlinkMssql • djxlinkOracle • djxlinkSybase • djxlinkTeradata 2. Issue the db2iupdt command on each federated database instance to enable federated access to the data sources. There is another script, the djxlink script, that attempts to create a wrapper library for every data source that is supported by WebSphere Federation Server. If you run djxlink script and have the client software for only some of the data sources installed, you will receive an error message for each of the data sources that you do not have installed.

2. After the link is performed, check the permissions on the wrapper libraries. Make sure that the libraries can be read and run by the database instance owners.

Checking the FEDERATED parameter

The FEDERATED parameter must be set to YES to enable the federated server to access to the data sources.

Before you begin

- WebSphere Federation Server must be installed on a server that will act as the federated server

About this task

Before you add data sources to the federated server and database, you should check the FEDERATED parameter setting.

Procedure

To check the FEDERATED parameter setting:

1. Issue the following DB2 command to display all of the parameters and their current settings:
GET DATABASE MANAGER CONFIGURATION
2. Check the CONCENTRATOR parameter setting. The CONCENTRATOR parameter and the FEDERATED parameter cannot be configured to YES at the same time. If the CONCENTRATOR parameter is set to YES, change the setting to NO. Issue the following DB2 command to change the setting:
UPDATE DATABASE MANAGER CONFIGURATION USING CONCENTRATOR NO

3. Check the FEDERATED parameter setting. If the FEDERATED parameter is set to NO, change the setting to YES. Issue the following DB2 command to change the setting:

```
UPDATE DATABASE MANAGER CONFIGURATION USING FEDERATED YES
```

Creating a federated database

After you set up the federated server, the instance owner must create a database on the federated server instance that will act as the federated database.

Before you begin

- You must have SYSADM or SYSCTRL authority to create a database.
- WebSphere Federation Server must be installed on a server that will act as the federated server

About this task

You must create a federated database before you can configure the federated server to access your data sources.

You can create the federated database by using the DB2 Control Center or from the DB2 command line.

Procedure

To create the federated database:

1. Determine the code set and collating sequence that you want to specify when you create the federated database.
2. Determine the method that you want to use to create the federated database:

Method	Step
Using the DB2 Control Center	Right-click on the Databases folder and click Create ->Database Using Wizard . The Create Database Wizard opens. Complete the steps in the wizard.
From the DB2 command line	Issue the CREATE DATABASE command. For example: <pre>CREATE DATABASE <i>federated</i> USING CODESET <i>ISO8859-15</i> TERRITORY <i>BR</i></pre>

Creating the database:

- Initializes a new database
- Creates the three initial table spaces
- Creates the system tables
- Allocates the recovery log

If your instance uses a multiple partition configuration, all of the partitions that are listed in the db2nodes.cfg file are affected when you create the database. The database partition from which this command is issued becomes the catalog partition for the new database.

Federated database code sets and collating sequences

When you create a federated database, you must decide the language that you want your data stored in and the order in which character data is sorted in a database.

The database is created using the settings of the application that creates the database unless you explicitly specify a territory, code set, and collating sequence. For example, if you create the federated database from a client that uses the ISO8859-15 code set, the database is with the same code set.

Recommendation: If the remote data sources that you need to connect to are using different or incompatible code sets, define the federated database as a Unicode database. To define the federated database as a Unicode database, specify `USING CODESET UTF-8` in the `CREATE DATABASE` command.

Code sets, code pages, and territories

You can specify the language for the federated database by including the `CODESET` and `TERRITORY` options on the `CREATE DATABASE` command.

A *code set* is an ISO term for unique bit patterns that are mapped to the characters contained in a specific natural language. After you create the database, you cannot change the specified code set. When you choose a code set, make sure that it can encode all the characters in the language that you will be using.

Code sets are mapped to IBM-defined *code pages*.

A *territory code* is used to provide region-specific information for the code set that you specify. After you create the database, you cannot change the specified territory.

Collating sequences

You specify the collating sequence for the federated database by including the `COLLATE USING` option on the `CREATE DATABASE` command.

For example:

```
CREATE DATABASE federated USING CODESET UTF-8 COLLATE USING UCA400-NO
```

After you create the database, you cannot change the collating sequence.

For relational data sources, if the federated database and the data source use the same collating sequence, you should set the `COLLATING_SEQUENCE` server option to 'Y'. Setting the `COLLATING_SEQUENCE` server option to 'Y' informs the federated server that the collating sequences of the federated database and the data source match. You set the `COLLATING_SEQUENCE` server option when you create the server definitions for the relational data sources.

The relational data sources that support the `COLLATING_SEQUENCE` server option are:

- DB2 family
- Informix
- Microsoft SQL Server
- ODBC

- OLE DB
- Oracle
- Sybase
- Teradata

Federated database national language considerations

When you create the federated database, you must decide which code set and territory to use. You can improve performance by using the same collating sequence that your data sources use.

For many data sources, the wrapper performs the following national language support (NLS) tasks when the wrapper connects to the data source:

1. Determines the code page and territory of the federated database.
2. Maps the code page and territory to a data source client locale name, if supported by the data source.
3. Depending on the data source, sets an environment variable. calls a data source API to tell the data source what the client locale is, or prepares to perform code set conversion.
4. The data is converted:
 - For data sources that perform code page conversion, the data source then converts character data between the code page of the remote database and the code page of the federated database.
 - For data sources that do not perform code page conversion, some of the wrappers perform the conversion.

For example, if the federated database uses code page 819, territory US, the equivalent Oracle client locale is American_America.WE8ISO8859P1. The wrapper sets the NLS_LANG variable to the Oracle client locale value. When data is sent from the Oracle database to the wrapper, the Oracle database converts the data from code set American_America.WE8ISO8859P1 to code page 819. When data is sent from the Oracle database to the wrapper, the Oracle server or client converts the data from the code page of the Oracle database to code page 819. When data is sent from the wrapper to the Oracle database, the Oracle server or client converts the data from code page 819 to the code page of the Oracle database.

Tip: Defining the federated database to use the same code set, territory, and collating sequence as your data source can improve performance. If you define the federated database to use the same code set, territory, and collating sequence as your data source, then the code page conversion is not necessary. Using the same national language settings can improve performance when you transfer large amounts of character data.

To specify the code set and territory on the federated database, you use the USING CODESET and TERRITORY parameters on the CREATE DATABASE command.

Collating sequences in a federated system

The collating sequence that you specify for the federated database impacts how queries with character sorts or comparisons are processed.

When the federated server receives a query, the SQL compiler consults information in the global catalog and the data source wrapper to help federated server process the query. As part of the SQL compiler process, the *query optimizer* analyzes the query. The SQL compiler develops alternative strategies, called *access plans*, for processing the query. The access plans might call for the query to be:

- Processed by the data sources
- Processed by the federated server
- Processed partly by the data sources and partly by the federated server

If the query requires character sorting or comparisons, the SQL compiler uses collating sequence information to determine which access plan to use. Performing character sorts and comparisons at the data source usually improves performance.

By default, the federated database collating sequence is case-sensitive. However some of the federated data sources use collating sequences that are case-insensitive. If the collating sequences of the federated database and the data source are different, the query results might differ depending on where the character sorts and comparisons are performed. Generally, if the operation is a character sort the same data is returned but the order of the results will be different. If the operation is a character comparison, the results returned might be different.

Where the sorting or comparison is processed depends on several factors:

- If the federated database collating sequence is the same as the data source collating sequence, the character sort or comparison can take place at the data source. The query optimizer can decide which is the most efficient way to complete the query, a local operation or a remote operation. It is assumed that all types of character comparisons and sorts by the data source would yield the same results as if those actions were performed by the federated database.
- If the federated database collating sequence is different than the data source collating sequence, but the data source collating sequence is case-sensitive, the character sort or comparison will take place at the federated database. It is assumed that the data source will yield the same results on character data for WHERE=, DISTINCT, and GROUP BY operations. But other operations, such as ORDER BY and WHERE with a greater than or less than predicate, will yield different results on character data.
- If the federated database collating sequence is different than the data source collating sequence, but the data source collating sequence is case-insensitive, the character sort or comparison will take place at the federated database. It is assumed that the data source will count uppercase and lowercase letters as equivalent and will include them both in a result, regardless of whether uppercase or lowercase was specified in the requested operation. WHERE=, WHERE with a greater than or less than predicate, ORDER BY, DISTINCT, and GROUP BY operations will not be pushed down to the data source.

For example, a case-insensitive data source assigns the same weights to the characters "S" and "s". A case-insensitive data source with an English code page considers the words STEWART, StEWArT, and stewart to be equal. However when a case-sensitive collating sequence is used, different weights are assigned to the characters. Depending on the sensitivity of the collating sequence, the result set of a character sort or comparison will be different.

When the collating sequences of the federated database and the data source differ, the federated server retrieves the data. The character sorts and comparisons are performed locally. The reason for performing these tasks locally is that the federated server users expect to see the query results ordered according to the collating sequence defined for the federated server. By ordering the data locally, federated server users are guaranteed that the result sets will be consistent. Retrieving data for local sorts and comparisons usually decreases performance.

If you need to see the character data ordered in the data source collating sequence, you can submit your query in a pass-through session.

To determine if a data source and the federated database have the same collating sequence, consider the following factors:

Code page

The code page scheme, such as ASCII and EBCDIC, that is used by the federated server and the data source impacts the results.

National language support (NLS)

The collating sequence is related to the language supported on a server. Compare the federated server NLS information for your operating system to the data source NLS information.

Data source characteristics

Some data sources are created using case-insensitive collating sequences, which can yield different results from the federated database in order-dependent operations.

Customization

Some data sources provide multiple options for collating sequences or allow the collating sequence to be customized.

There are several options that you can use to set the collating sequence:

- Set the collating sequence when you create the federated database
- Set the `COLLATING_SEQUENCE` option when you create the server definition for a data source. This option is available only for relational data sources.

Confirming and registering the Microsoft SQL Server Client Version 2000 driver (Windows)

The Microsoft SQL Server Client Version 2000 driver is typically installed when you install Windows. To access Microsoft SQL Server data sources, the driver must be installed, registered, and configured.

Procedure

To confirm that the driver is installed and to test the connection:

1. Open the Windows Control Panel.
2. Open the Administrative Tools folder.
3. Double-click **Data Sources (ODBC)** to open the ODBC Data Source Administrator window.
4. Click the **Drivers** tab and locate the SQL Server entry in the list of ODBC drivers that are installed on your system. If the SQL Server entry is in the list, the driver is installed.
5. Click the **System DSN** tab.
6. Click **Configure** to register the driver and test the connection to the Microsoft SQL Server data source. If you are using Microsoft SQL Server 2000 Personal Edition, you must use the SQL Server Client Network Utility to add a new SQL Server ODBC data source to your ODBC System DSN list.
7. See the installation procedures in the documentation that comes with the ODBC driver for specific details on how to install and configure the driver.

Configuring security for federated systems

When you configure your federated system to access your data sources, consider how you want to configure security. If you are connecting to Oracle databases with label security enabled, you must understand how label security works with the federated system. In addition, you can use an external repository for user mappings between the federated server and the data source.

Oracle Label Security and federated systems

When you create nicknames against data source objects with Oracle Label Security, the nicknames are not cached and materialized query tables cannot be created on them. The secure data remains on the Oracle system and only users with the appropriate authority can see it.

Administrators can use Oracle Label Security to apply security policies to each row in a table. The security policies determine a user's level of access to the data in the table that is based on the authorities that are granted to their user ID or session ID. See the documentation about Oracle Label Security from the Oracle Corporation for detailed information about configuring and managing security policies and labels.

Federated access to tables with Oracle Label Security

When you create a nickname on an Oracle data source object, the federated server automatically detects whether the data source uses Oracle Label Security. If Oracle Label Security is being used, the nickname is not cached. You can use the ALTER NICKNAME statement to allow or disallow caching. For example, if you created a nickname on a data source object with Oracle Label Security before federated support for the feature was available, you can alter the nickname to disallow caching. If you created a nickname on a data source object with Oracle Label Security and Oracle Label Security is removed, you can alter the nickname to allow caching.

A database administrator can choose to hide a label to prevent some users from knowing that it exists. In this case, the Label column is hidden in the table. Nicknames with hidden label columns are not cached.

Overview of the user mapping plugin for external repositories

You can develop a plugin that retrieves user mappings from an external repository instead of storing the user mappings on the federated server.

The plugin that you develop must retrieve the user mappings from the repository and pass the authentication information to the federated server. Federation includes a sample plugin that is designed for a Lightweight Directory Access Protocol (LDAP) server repository. To work with other types of external repositories, you can develop your own plugin or extend the sample LDAP plugin. After you create the plugin, you must configure the federated server to use the plugin by specifying an option on the wrapper or server definition.

Advantages of using an external repository to store user mappings

Storing user mappings in an external repository can provide improved security and reduced maintenance. An external repository can be shared by many federated servers.

By default, user mappings are stored in the catalog table on the federated server. You can query the catalog view SYSCAT.USEROPTIONS for details. The remote passwords that are stored in the database catalog use the encryption algorithm that is provided by the federated server. Because the remote passwords expire on a regular basis, and you store the user mappings on each federated server, the routine maintenance of updating passwords can become time consuming.

With an external repository, you can store user mappings in a central repository that many federated servers can use to retrieve user mappings.

An external repository such as a Lightweight Directory Access Protocol (LDAP) server, offers increased security. On an LDAP server, you can encrypt the remote passwords with an encryption algorithm and a secret key. If you choose to enable secure communication using secure sockets layer (SSL), you can protect the information that is passed between the federated server and the external repository. The external repository must support SSL communications to use this functionality.

You must develop your own plugin that matches the security and encryption settings of your external repository.

Relationship between the federated server and the user mapping plugin

The user mapping plugin allows the federated server to retrieve user mappings from an external repository. When DB2_UM_PLUGIN option is set and the federated server tries to establish a connection to the data source, the server uses the plugin to retrieve the user mapping for that data source.

The plugin acts as a gateway between the federated server and the external repository. When the federated server needs to access the repository for user mappings, the federated server provides the following information to the plugin:

- Local instance name
- Local database name
- Remote server name
- Local user ID

The plugin uses this information to locate the appropriate user mapping in the external repository. For example, the plugin finds the appropriate user entry in a Lightweight Directory Access Protocol (LDAP) server and requests that the LDAP server return the remote user ID and remote password.

Other types of external repositories will store and return the user mappings using other methods. You must develop a plugin that contains the code to allow the federated server to interface with the external repository.




Figure 4. Relationship between the federated server, the user mapping plugin, the external repository and the data source.. When a DB2 client requests a connection through the federated server to data that resides on the remote data source, the local user ID of the client is sent to the federated server (1). If the DB2_UM_PLUGIN option is set for a wrapper, the federated server loads the user mapping plugin and provides the plugin with the local instance name, local database name, remote server name, and local user ID. The plugin uses this information to locate (2) the user mapping in the repository. The plugin retrieves the remote user mapping and performs any necessary decryption (3). The plugin returns the remote user ID and remote password to the federated server. The federated server tries to connect to the data source through the wrapper (4). The remote user ID and remote password that was retrieved from the external repository is used to connect to the data source.

User mapping plugin architecture

You can develop your own plugin for retrieving user mappings from an external repository. The `UserMappingRepository`, `UserMappingCrypto`, `UserMappingEntry`, `UserMappingOption`, and `UserMappingException` classes provide the interface and utilities to develop a plugin to return the user mappings from the external repository to the federated server.

The following table lists the interface classes and utility classes that comprise the user mapping plugin architecture. You must extend the interface classes to work with your external repository. The utility classes can be used without modification.

Table 31. Interface classes and utility classes

Interface classes	Utility classes
UserMappingCrypto class	UserMappingEntry class
UserMappingRepository class	UserMappingOption class
	UserMappingException class

Some functions or methods that are in the interface classes act as utility functions. For example, you can use the `getChars()` function and `getBytes()` function in the `UserMappingCrypto` class without modification.




Figure 5. Diagram of the classes for the architecture of the user mapping plugin. The diagram shows the relationship of the classes to each other. The *0..n* term, means there can be multiple zero or more of those objects. For example, a `UserMappingEntry` object can have multiple `UserMappingOption` objects but there can be only be a single `UserMappingRepository` object. The `UserMappingCryptoLDAP` class and `UserMappingRepositoryLDAP` class are darker to show that they are extended from their parent classes.

UserMappingRepository class

The `UserMappingRepository` class is an abstract class that does not have a constructor. You must create a subclass of the `UserMappingRepository` class or modify the subclass in the LDAP sample plugin.

The `UserMappingRepository` class contains the following public methods: `getVersionNumber()`, `getCrypto()`, `connect()`, `disconnect()`, `fetchUM()`, and `lookupUM()`. You must create your own subclass of the `UserMappingRepository`

that contains these functions. In these functions, write the code that allows the plugin to interact with your external repository.

Public methods

int getVersionNumber()

Returns the version number of the plugin development kit that is used by the plugin.

UserMappingCrypto getCrypto()

Returns the UserMappingCrypto object that is associated with this UserMappingRepository object.

abstract void connect()

You must implement your own method for connecting to your repository within this function.

abstract void disconnect()

You must implement your own method for disconnecting from your repository within this function.

abstract void fetchUM(UserMappingEntry um)

You must implement your own method for retrieving the user mapping from your repository within this function. The **um** parameter contains the detailed query information that is used to determine which user mapping to retrieve.

UserMappingEntry lookupUM(UserMappingRepository repository, String iiInstanceName, String iiDatabaseName, String iiRemoteServerName, String iiAuthid)

This function is used primarily for testing the plugin. The function uses the **iiInstanceName**, **iiDatabaseName**, **iiRemoteServerName**, and **iiAuthid** parameters as input to create and initialize the UserMappingEntry class. The function calls the connect, fetchUM, and disconnect methods.

Sample plugin files

You can view the implementation of these functions in a sample plugin that retrieves user mappings from an LDAP server. The files are in the `sqlib/samples/federated/umplugin/ldap/` directory. The functions from this class are used in the `UserMappingRepositoryLDAP.java` and `UserMappingLookupLDAP.java` files.

UserMappingCrypto class

If your external repository encrypts or encodes your remote passwords, you must create your own subclass of the UserMappingCrypto class. The constructor of the subclass that you create is used to construct the cryptography object. The methods of the cryptography class are called by other classes when the user mapping passwords need to be encrypted, decrypted, encoded, or decoded.

The UserMappingCrypto class contains the following public methods: `encrypt()`, `decrypt()`, `encode()`, and `decode()`. In these functions, you must write your code for encrypting, decrypting, encoding, and decoding the remote password. The `getBytes()` and `getChars()` functions are utility functions that are inherited and can be used without modification.

Your encryption, decryption, encoding, and decoding methods must match the encryption and encoding methods that are used by your external repository for protecting the stored passwords.

Public methods

abstract byte[] encrypt(byte[] plainValue)

Implement the encryption algorithm that matches the encryption algorithm that is used by your external repository.

abstract byte[] decrypt(byte[] encryptedValue)

Implement the decryption algorithm that reverses the encryption algorithm that is used by your external repository and returns the password.

abstract string encode(byte[] bytes)

Write or implement a function that encodes the **bytes** parameter into a string. This function encodes the encrypted value, which is in bytes, into a string.

abstract byte[] decode(String[] string)

Write or implement a function that decodes the **string** parameter into bytes. This function is used to decode the retrieved password, which is a string, into bytes so that the value can be decrypted.

byte[] getBytes(char[] chars)

This function is inherited and can be used without modification. The function transforms each character of a string into a byte.

char[] getChars(byte[] bytes)

This function is inherited and can be used without modification. The function transforms each byte into a character.

Protected attributes

SecretKey key

The secret key that is used to encrypt and decrypt the remote passwords.

Cipher cipher

The algorithm that is used to encrypt the password by using the secret key.

Sample plugin files

You can view the implementation of these functions in a sample plugin that retrieves user mappings from an LDAP server. The files are located in the `sqlib/samples/federated/umplugin/ldap/` directory. The functions from this class are used in the `UserMappingRepositoryLDAP.java` and `UserMappingSetupLDAP.java` sample files.

UserMappingEntry class

The `UserMappingEntry` class is a utility class that creates and holds the user mapping options. The methods in the `UserMappingEntry` class are called by the `fetchUM()` and `lookupUM()` functions from the `UserMappingRepository` class.

The `UserMappingEntry` class contains the following public methods: `UserMappingEntry()`, `getRepository()`, `getIIInstanceName()`, `getIIDatabaseName()`, `getIIRemoteServerName()`, `getIIAuthID()`, `getFirstOption()`, and `addOption()`.

Public methods

UserMappingEntry(UserMappingRepository repository, String iiInstanceName, String iiDatabaseName, String iiRemoteServerName, String iiAuthID)

This constructor is used to instantiate the `UserMappingEntry` object with the following input parameters:

- **iiInstance** - Instance name of the federated server
- **iiDatabase** - Database name on the federated server

- **iiRemoteServerName** - Remote server name for the data source
- **iiAuthid** - Local user ID that is associated with the user mapping

UserMappingRepository getRepository()

Returns the UserMappingRepository object that is associated with this UserMappingEntry.

string getIIInstanceName()

Returns the name of the instance.

string getIIDatabaseName()

Returns the name of the database.

string getIIRemoteServerName()

Returns the name of the remote server.

string getIIAuthID()

Returns the name of the local user ID that is associated with the user mapping.

UserMappingOption getFirstOption()

Returns the first UserMappingOption object that belongs to this UserMappingEntry object.

void addOption(UserMappingOption newOption)

Adds a new UserMappingOption object to this UserMappingEntry object.

Sample plugin files

You can view the use of these functions in a sample plugin that retrieves user mappings from an LDAP server. The files are in the sqllib/samples/federated/umplugin/ldap/ directory. The functions from this class are used in the UserMappingRepositoryLDAP.java and UserMappingLookupLDAP.java files.

UserMappingOption class

The UserMappingOption class is a utility class that contains the functions for providing the federated server with the two components of the user mapping: the remote user ID and the remote password.

The UserMappingOption class contains the following public methods: getEntry(), getName(), setName(), getNextOption(), setNextOption(), and getValue(). The sqllib/samples/federated/umplugin/ldap/UserMappingRepositoryLDAP.java sample file, contains a sample of using the UserMappingOption object in an LDAP sample plugin.

Public methods

UserMappingEntry getEntry()

Returns the UserMappingEntry object that is associated with this UserMappingOption object.

string getName()

Returns the name of the option.

void setName()

Sets the name of the option.

UserMappingOption getNextOption()

Returns the next option.

void setNextOption(UserMappingOption nextOption)

Sets the next option.

abstract object getValue()

Returns the value of the option. This function must be implemented to return either a string value or binary value. See the StringOption and BinaryOption methods that are listed below.

abstract object setValue()

Sets the value of the option. This function must be implemented to set either a string value or binary value. See the StringOption and BinaryOption methods that are listed below.

StringOption class that is extended from the UserMappingOption class

The StringOption class contains the following public methods: getValue() and setValue().

Public methods**object getValue()**

Returns the value of the string option when the option is a string.

void setValue(string value)

Sets the value of the string option.

BinaryOption class that is extended from the UserMappingOption class

The BinaryOption class contains the following public methods: getValue() and setValue().

Public methods**object getValue()**

Returns the value of the binary option.

void setValue(byte[] value)

Sets the value of the binary option.

Sample plugin files

You can view the use of these functions in a sample plugin that retrieves user mappings from an LDAP server. The files are in the `sqllib/samples/federated/umplugin/ldap/` directory. The functions from this class are used in the `UserMappingRepositoryLDAP.java` and `UserMappingLookupLDAP.java` files.

UserMappingException class

The UserMappingException class is a subclass of the `java.lang.Exception` class and is used by the plugin to report errors.

The UserMappingRepository class contains the following public methods: `UserMappingException()`, `getErrorNumber()`, and `getErrorMessage()`.

Public methods**UserMappingException(int errorNumber)**

The constructor that is used for instantiating the UserMappingException object that is used for reporting errors. The **errorNumber** parameter that is sent to the UserMappingException object defines the type of error that is reported.

getErrorNumber()

Returns the error number of the exception.

The UserMappingRepositoryLDAP.java file, contains this function for catching and reporting errors for testing an LDAP plugin.

getErrorMessage()

Returns the error message of the exception.

The UserMappingRepositoryLDAP.java file, contains a method for catching and reporting errors for testing an LDAP plugin.

Error messages

Possible error numbers, constant names, and error messages are listed in the following table.

Table 32. Error numbers and messages

Error number	Constant name	Error message
1	INITIALIZE_ERROR	The plugin failed to initialize.
2	CONNECTION_ERROR	Unable to connect to the repository.
3	AUTHENTICATION_ERROR	Unable to authenticate with the repository.
4	LOOKUP_ERROR	Lookup on the repository failed.
5	DECRYPTION_ERROR	Decryption failed.
6	DISCONNECT_ERROR	Unable to disconnect from the repository.
7	INVALID_PARAMETER_ERROR	Invalid parameter.
8	UNAUTHORIZED_CALLER	The caller is not authorized to call the plugin.

Sample plugin files

You can view the use of these functions in a sample plugin that retrieves user mappings from an LDAP server. The files are in the sqllib/samples/federated/umplugin/ldap/ directory. The functions from this class are used in each of the sample Java files.

LDAP sample plugin

A sample plugin for retrieving user mappings from a Lightweight Directory Access Protocol (LDAP) server is provided in the sqllib/samples/federated/umplugin/ldap/ directory. Before you can use the sample plugin, you must modify the plugin to match the settings of your LDAP server.

By default, user mappings are stored locally on each federated server for a data source. An LDAP server stores objects (for example, user entries) in a directory tree. These objects can have attributes (for example, passwords). The LDAP server is often used store information about users and their related information such as the e-mail address of a user.

The LDAP plugin is used to retrieve user mapping information from an LDAP server. Many user mappings can be stored on one LDAP server and be accessed by many federated servers. The LDAP system administrators can choose the level of security that they want to protect the user attributes, such as passwords or other sensitive information that should not be transmitted as plain text.

Because the directory structure and security settings can vary on each server, you must modify the sample plugin to match the settings of your LDAP server.

Description of files for the LDAP sample plugin

The sample plugin contains Java source code that allows the plugin to connect to a Lightweight Directory Access Protocol (LDAP) server, retrieve the user mappings, and decrypt the remote password. Because server settings can vary, you must develop your own LDAP plugin to match your server's encryption settings, directory structure, and other settings.

By using the files that are described below, you can develop a plugin for retrieving user mappings from an LDAP server or you can extend the sample to use with any external repository. The sample code is in the `sqllib/samples/federated/umplugin/ldap/` directory. Before you begin to develop your plugin, copy the sample files to an empty working directory.

The sample plugin

You can develop your own LDAP plugin by using the sample code as a starting point. You can also find instructions in the sample folder for the LDAP plugin for testing the sample code with sample LDAP directory entries.

The sample plugin consists of four Java source files and two LDAP Lightweight Directory Interchange Format (LDIF) files. The following table contains a description of each file.

Table 33. Description of the files in the sample plugin

File name	Description
README.txt	This file contains a condensed version of the instructions and documentation for using the sample plugin.
UserMappingCryptoLDAP.java	This Java class contains the code that implements the security measures for encrypting, decrypting, encoding, and decoding the user mappings that are retrieved from the LDAP server. This file must be modified to work with your LDAP server.
UserMappingSetupLDAP.java	This Java class creates the configuration file that stores the LDAP connection information and other configuration parameters, including: IP address or host name, SSL or non-SSL, user ID, and password.
UserMappingRepositoryLDAP.java	This Java class contains the code for connecting, disconnecting, and fetching user mappings from the LDAP server. The code for this file uses the schema that is defined in the <code>schema.ldif</code> file. If you want to change the schema, you must also change a section in this file.
UserMappingLookupLDAP.java	This Java class contains the code to perform a LDAP lookup test. You can test your plugin using this file before testing the plugin on the federated server.
schema.ldif	This file is loaded into your LDAP server to define the schema. The Lightweight Directory Interchange Format (LDIF) file contains objects and attributes that are added to the LDAP server.

Table 33. Description of the files in the sample plugin (continued)

File name	Description
entry.ldif	The entry.ldif file adds user entries to the LDAP server. The user entries use the objects and attributes from the schema.ldif file to store the user mappings.

You can develop the plugin to use the settings that are appropriate for your LDAP server. For example, the encryption and encoding settings in the plugin must match the encryption and encoding settings on your LDAP server.

Developing a plugin for retrieving user mappings from an external repository

By using the sample code from the Lightweight Directory Access Protocol (LDAP) plugin as a starting point, you can develop a plugin for retrieving user mappings from an external repository. The sample files are located in the `/sqlib/samples/federated/umplugin/ldap/` directory.

You should restrict access to the code for the encryption and decryption algorithms to protect the security of your repository and your remote passwords.

Before you begin

To develop a plugin you must have the following installed:

- The Java Development Kit (JDK) version 1.4 or later
- The `db2umplugin.jar` file. This Java Archive (JAR) file is installed with the DB2 server installation or the DB2 Application Development client installation.
- The LDAP sample files, which are installed in the `/sqlib/samples/federated/umplugin/ldap/` directory. The sample files are installed with the DB2 Application Development Client installation.

About this task

The plugin that you develop must be able to connect to your repository, retrieve the user mappings, and decrypt the remote passwords. The repository that you are using might change the functions that you include in your plugin, for example, if you repository stores passwords as plain text then you would not need a decryption function. A sample plugin that you can reference is in the `/sqlib/samples/federated/umplugin/ldap/` directory. The following procedure is a general guide for developing your plugin.

Procedure

To develop your plugin, the following steps are required:

Extending the sample LDAP plugin files to other external repositories:

A sample plugin is provided for retrieving user mappings from a Lightweight Directory Access Protocol (LDAP) server. You can modify the plugin to retrieve user mappings from other repositories. The sample plugin provides a starting point to help you reduce your development time.

Each of the sample files accomplishes a different task in the process of retrieving user mappings. Many of the functions and classes in the sample code will need to

be modified to work with a different repository. Refer to the user mapping plugin architecture when you develop your plugin. The important functions are listed in the following table.

Table 34. Functions and classes that must be modified in the sample plugin to extend the plugin to other repositories.

File name	Function or class name	Class to reference
UserMappingCryptoLDAP.java	UserMappingCryptoLDAP() encrypt() decrypt() getKey() decode() encode()	UserMappingCrypto class UserMappingException class
UserMappingRepositoryLDAP.java	class UserMappingRepositoryLDAP(String configFilePath) connect() disconnect() fetchUM()	UserMappingRepository class UserMappingEntry class UserMappingOption class UserMappingException class
UserMappingSetupLDAP.java	You will need to modify the file to create a configuration file that stores the values that your plugin requires for connecting to and retrieving user mappings from the external repository. If you choose to manually create a configuration file, ensure that any passwords that you store in the file are encrypted. The configuration file name should match the name of the repository class (for example, UserMappingRepositoryXXXX class and UserMappingRepositoryXXXX.cfg file).	

Security considerations for the user mapping plugin:

When developing and using your plugin, you are sending sensitive user IDs and passwords between multiple sources. You can protect your information by restricting access to code, auditing plugin usage, and encrypting communication.

You must develop your plugin to use security settings that match the security settings that are used by your external repository. You should choose an external repository that allows you to encrypt sensitive information that is stored. The user mapping cryptography class must contain the encryption schema and secret key that allows for decrypting and decoding the passwords. You should restrict access to the source code of the plugin so that this information stays secure.

If auditing is turned on, any attempt to access a user mapping through plugin by the federated server has a VALIDATE audit record. You can configure the db2audit tool to capture VALIDATE records: db2audit configure scope VALIDATE.

You can also protect the communication between the federated server and the external repository by using a secure socket layer (SSL). When you create your configuration file for the user mapping plugin, you can specify that the plugin uses SSL to protect communications.

Modifying the UserMappingCryptoLDAP sample file:

You need to modify the code that is in the UserMappingCryptoLDAP.java file to match the encryption or encoding that is used on your Lightweight Directory Access Protocol (LDAP) server. Because the encryption methods should be secret and unique, this task only directs you to the sections or functions that you must modify.

You should restrict access to the code for the encryption and decryption algorithms to protect the security of your repository and your remote passwords.

Before you begin

To use the user mapping plugin for LDAP servers, you need to know how the passwords are protected in the LDAP server.

About this task

You will need to write the code for implementing your security measures that match your LDAP configuration. The `UserMappingCryptoLDAP.java` file provides the functions for encrypting, decrypting, encoding, and decoding the remote passwords for the data sources.

Procedure

To modify the security functions in the `UserMappingCryptoLDAP` file:

1. Open the `UserMappingCryptoLDAP.java` file with a text editor.
2. Below the IBM copyright and legal disclaimer, import the packages that your code will reference. The sample plugin uses the `javax.crypto` and `javax.crypto.spec` Java packages. These packages provide the classes for the cipher (encrypting and decrypting) and the key and algorithm parameters. If you choose to use your own packages, replace these packages with your packages.
3. Update the following functions:

public UserMappingCryptoLDAP()

Replace the code for the cipher with the code for the cipher that matches the password encryption that is used by your LDAP server.

public byte[] encrypt(byte[] plainValue)

This function provides the code that encrypts the passwords so that they can be stored on the LDAP server. This function also encrypts the LDAP connection password that is stored in the configuration file.

Replace this function's code with your own code that encrypts the **plainValue** parameter.

public byte[] decrypt(byte[] encryptedValue)

Replace this function's code with your own code that decrypts the **encryptedValue** parameter.

private SecretKey getKey()

Replace this function's code with the code to provide the plugin with the key that is used to encrypt and decrypt your passwords.

public byte[] decode(String string)

The passwords are first encrypted and then encoded. This function provides the code for decoding the passwords before the passwords are decrypted. The encrypted passwords are encoded to transform the binary output of the encrypted password into ASCII characters.

Replace this function's code with your own code that decodes the **string** parameter.

public String encode(byte[] bytes)

The passwords are first encrypted and then encoded. This function provides the code for encoding the binary output of the encrypted passwords. The encrypted passwords are encoded to transform the binary output of the encrypted password into ASCII characters.

Replace this function's code with your own code that encodes the **bytes** parameter.

Modifying the UserMappingRepositoryLDAP sample file:

The UserMappingRepositoryLDAP.java sample file contains the functions for connecting to the Lightweight Directory Access Protocol (LDAP) server, retrieving the user mappings, and disconnecting. You need to modify the sample code to match the object classes that are defined in the directory structure of your LDAP server.

About this task

To retrieve the user mappings from the LDAP server, the plugin must search the directory for the user entries that have the attributes that define the user mapping entry. The following information is stored as attributes of a user:

- Remote server name
- Instance name
- Database name
- Remote user name
- Remote user password

The sample code assumes that the user entry will be identified by the *inetOrgPerson* object class and that the user mapping entry will be identified by the *IIUserMapping* object class. The Lightweight Directory Interchange Format (LDIF) sample files (schema.ldif and entry.ldif) are used to load the sample schema and sample entries into the LDAP server.

The code in the UserMappingRepositoryLDAP.java file assumes that the schema.ldif file contains the LDIF code for objects and attributes with the following names:

- *IIUserMapping* object
 - *IIRemoteServerName* attribute
 - *IIInstanceName* attribute
 - *IIDatabaseName* attribute
 - *IIRemotePassword* attribute
 - *uid* attribute

You must modify the UserMappingRepositoryLDAP.java file to match the schema in your LDAP server. The UserMappingRepositoryLDAP.java file must search for and retrieve the user mapping entries from the LDAP server. The LDIF files are provided as a sample schema and method of storing user mapping entries.

Procedure

To modify the schema that is used by the UserMappingRepositoryLDAP.java file:

1. Locate the code: `private String UserObjectClassName = "inetOrgPerson"` and replace the *inetOrgPerson* value with the name of the object class for user entries that is used by your LDAP server.
2. Optional: Change the attribute names that are used by the plugin. Replace the values for the *IIRemoteServerAttrName*, *IIInstanceAttrName*, *IIDatabaseAttrName*, and *IIRemotePasswordAttrName* variables with the attribute names that you choose.

3. If you choose to use the LDIF files, ensure that the schema in the LDIF files matches the structure that is searched by the UserMappingRepositoryLDAP.java file.

Compiling the user mapping plugin files:

When you finish modifying your source files for your user mapping plugin, you must compile the Java source files.

About this task

In the commands below, if your full path contains spaces, then the full path must be enclosed in quotation marks, for example, "C:\program files\sqllib\java\db2umplugin.jar". %DB2PATH% is the path to your Windows DB2 sqllib folder. \$DB2PATH is the path to your the sqllib folder on UNIX.

The commands below assume that you are using the file naming convention that is based on the names of the classes. Replace the XXXX with the name that you chose.

Procedure

To compile the Java source files:

1. Issue the compile command:

Windows:

```
javac -classpath %DB2PATH%\sqllib\java\db2umplugin.jar; ^
"%CLASSPATH%" -d . ^
.\UserMappingRepositoryXXXX.java ^
.\UserMappingCryptoXXXX.java ^
.\UserMappingSetupXXXX.java ^
.\UserMappingLookupXXXX.java
```

UNIX:

```
javac -classpath $DB2PATH/java/db2umplugin.jar: \
$CLASSPATH -d . \
./UserMappingRepositoryXXXX.java \
./UserMappingCryptoXXXX.java \
./UserMappingSetupXXXX.java \
./UserMappingLookupXXXX.java
```

2. Archive the Java class files into a single Java Archive (JAR) file: The period symbol after the output file name, directs the command to find and place the files in the same directory. If you change the directory, use the appropriate file path for your operating system (for example, /home/user/folder or C:\test\folder).

```
jar -cfM0 UserMappingRepositoryXXXX.jar .
```

Creating the configuration file for the user mapping plugin:

The configuration file for the sample LDAP plugin stores the connection information for your LDAP server. After compiling the Java source files, you can run the configuration program that prompts you to enter your LDAP connection information.

About this task

The application uses your input to create the UserMappingRepositoryLDAP.cfg file to store the configuration information. If a password is required to connect to your LDAP server, the password is stored in the configuration file. The password is encrypted with the algorithm that you provided in the UserMappingCryptoLDAP.java file.

In the commands below, if your full path contains spaces, then the full path must be enclosed in quotation marks, for example, "C:\program files\sqllib\java\db2umplugin.jar". %DB2PATH% is the path to your Windows DB2 sqllib folder. \$DB2PATH is the path to your the sqllib folder on UNIX.

The application prompts you to enter the following configuration information:

- Host name or IP address of the LDAP server
- Port number (default is 389) of the LDAP server
- Distinguished Name of the LDAP subtree (for example, ou=ii,o=ibm,c=us)
- User ID for connecting to the LDAP server
- Password for connecting to the LDAP server
- SSL configuration

Procedure

To create the configuration file for the sample LDAP plugin:

Windows:

```
java -classpath %DB2PATH%\sqllib\java\db2umplugin.jar; ^
.\UserMappingRepositoryLDAP.jar;%CLASSPATH% UserMappingSetupLDAP
```

UNIX:

```
java -classpath $DB2PATH/sqllib/java/db2umplugin.jar: \
./UserMappingRepositoryLDAP.jar:$CLASSPATH UserMappingSetupLDAP
```

Testing the user mapping plugin:

You can test the functionality of your user mapping plugin outside of the federated server. To test outside of the federated server, you can develop an application that attempts to connect to and retrieve user mappings from your external repository.

About this task

You can develop a simple program that calls the lookupUM() method that your UserMappingRepositoryXXXX class inherited from the UserMappingRepository class to connect and retrieve user mappings from your external repository. You can view the UserMappingLookupLDAP.java file that is located in the sqllib/samples/federated/umplugin/ldap/ directory.

In the commands below, if your full path contains spaces, then the full path must be enclosed in quotation marks, for example, "C:\program files\sqllib\java\db2umplugin.jar". %DB2PATH% is the path to your Windows DB2 sqllib folder. \$DB2PATH is the path to your the sqllib folder on UNIX.

The test program must take the parameters that are required for identifying the user mapping:

- *remoteServerName* - Remote server name for the data source
- *iiAuthid* - Local user ID that is associated with the user mapping

- *iiInstance* - Instance name of the federated server
- *iiDatabase* - Database name on the federated server

Procedure

To test the user mapping plugin:

Windows:

```
java -classpath %DB2PATH%\sqllib\java\db2umplugin.jar; ^
.\UserMappingRepositoryXXXX.jar;%CLASSPATH%" ^
UserMappingLookupXXXX -server remoteServerName ^
-authid iiAuthid ^
-instance iiInstance -database iiDatabase
```

UNIX:

```
java -classpath $DB2PATH/sqllib/java/db2umplugin.jar: ^
./UserMappingRepositoryXXXX.jar:$CLASSPATH \
UserMappingLookupXXXX -server remoteServerName \
-authid iiAuthid \
-instance iiInstance -database iiDatabase
```

Installing the user mapping plugin files

After compiling and testing the Java files for your user mapping plugin, you must install the files on the federated server.

About this task

In the commands below, *sqllib* is the full path to your DB2 installation.

The commands below assume that you are using the file naming convention that is based on the names of the classes. Replace the XXXX with the name that you chose.

Procedure

To install the compiled files:

Copy your UserMappingRepositoryXXXX.jar and UserMappingRepositoryXXXX.cfg files into the *sqllib/function/directory*.

Placing the files in this directory allows the federated server to load and call the classes contained in those files.

You can now configure the wrapper or server option on your federated server to use the user mapping plugin to retrieve the user mappings from your repository.

Configuring access to the user mapping plugin

You must set the DB2_UM_PLUGIN option to configure the federated server to access the user mappings in the external repository.

Before you begin

Before you configure the federated server to access the user mappings in an external repository, you must:

- Develop a user mapping plugin
- Install the plugin on the federated server

- Update your database manager configuration:

```
db2 update dbm cfg using JDK_PATH <your_jdk_path>
db2 terminate
db2stop
db2start
```

About this task

The DB2_UM_PLUGIN option must contain the full path of the class, including the package name. If you develop a package, you must include the package name before the class name for example, 'package.classname'. If you use the LDAP sample plugin, specify the value 'UserMappingRepositoryLDAP' in the DB2_UM_PLUGIN option. The sample plugin is not developed as a package.

Where you specify the DB2_UM_PLUGIN option impacts how the federated server uses the user mapping plugin:

- If you specify the DB2_UM_PLUGIN option on a wrapper, all of the servers that use the wrapper will use the value in that DB2_UM_PLUGIN wrapper option.
- If you specify the DB2_UM_PLUGIN option on a server definition, that server will use the value in that DB2_UM_PLUGIN server option.
- If you specify the DB2_UM_PLUGIN option on both a wrapper and server definition, the value that is specified on the server definition takes precedence over the value that is specified on the wrapper. A warning message, SQL20351W, is returned that indicates that the server option is not overwritten by the wrapper option.
- If the DB2_UM_PLUGIN is specified for a wrapper or server option, you can still use DDL to create a user mapping. A warning message, SQL20352W, is returned that indicates that the user mapping is stored in the federated database catalog but will not be used to access the remote server.

Procedure

To configure the federated server to access the user mapping external repository:

Choose how you want to implement the user mapping plugin:

Method	SQL statement
Specify the user mapping plugin when you create a wrapper	<pre>CREATE WRAPPER net8 OPTIONS (DB2_UM_PLUGIN 'UserMappingRepositoryLDAP');</pre>
Alter an existing wrapper to specify the user mapping plugin	<pre>ALTER WRAPPER drda (ADD DB2_UM_PLUGIN 'UserMappingRepositoryLDAP');</pre>
Specify the user mapping plugin when you create a server definition	<pre>CREATE SERVER server_name TYPE date_source_type VERSION version_number WRAPPER net8 OPTIONS (DB2_UM_PLUGIN 'UserMappingRepositoryLDAP');</pre>
Alter an existing server definition to specify the user mapping plugin	<pre>ALTER SERVER server_name OPTIONS (ADD DB2_UM_PLUGIN 'UserMappingRepositoryLDAP');</pre>

After you set the DB2_UM_PLUGIN option, the federated server uses the plugin to retrieve user mappings from the external repository server by using the connection information that you specify in the UserMappingRepositoryXXXX.cfg file. XXXX is the name of the plugin.

To alter the wrapper to use a different plugin

```
ALTER WRAPPER drda (  
    SET DB2_UM_PLUGIN 'com.package_name.um.UserMappingRepositoryXXXX'  
);
```

To alter a server definition to use a different plugin

```
ALTER SERVER server_name  
OPTIONS (  
    SET DB2_UM_PLUGIN 'UserMappingRepositoryXXXX'  
);
```

Configuring a connection to the analysis database for WebSphere Information Analyzer

When you analyze data, the results of your data analysis are stored in the analysis database. The connection to the analysis database must be configured after you install the suite.

If you installed DB2, the analysis database name is IADB by default. If you did not install DB2, you must use the name that you provided when you manually created this database.

The user ID that you use to log in to the IBM Information Server console must have the Information Analyzer Data Administrator and Information Analyzer Project Administrator suite component roles.

To configure a connection to the analysis database:

1. On the Windows computer where you installed the clients, click **Start → All Programs → IBM Information Server → IBM Information Server Console** and log in to the console with credentials that have WebSphere Information Analyzer Data Administrator and WebSphere Information Analyzer Project Administrator privileges. The console opens to the My Home workspace.
2. Click **Home → Configuration → Analysis Settings**. In the Analysis Database tab, provide the required information about the analysis database and the connection. The credentials that you supply on the Analysis Database tab are used by the WebSphere Information Analyzer client, the domain component, and the engine component to access the analysis database.
3. Type the ODBC DSN identifier. The engine components use this identifier to access the analysis database. The ODBC DSN must exist on the computer where the engine is installed. If the analysis database is DB2, the ODBC DSN must use the IBM DB2 Wire Protocol driver.
4. Ensure that the user name and password represent a valid database user for the specified database. The user must have create, drop, delete, select, insert, and update privileges in the database.
5. Click **Validate Settings** to test the client connection to the analysis database. This test verifies the ODBC connection from the client to the analysis database. It does not test the ODBC Connection from the engine to the database. Verify the connection from the ODBC Administrator on the computer where the engine is installed using the user name and password that was entered here.

6. In the Analysis Engine tab, verify the WebSphere DataStage user name and password. The WebSphere DataStage user must be the name of the WebSphere DataStage OS user on the computer that is running the WebSphere DataStage engine.
7. Click **Validate Settings** to test the WebSphere DataStage credentials.
8. Click **Save All**.
9. Click the **My Home** tab and expand the **Analyzing information** section to begin analyzing your information.

Validating your installation

After you install and configure IBM Information Server and the components that you have a license for, you can validate that the installation was successful.

Complete all the installation and configuration tasks for the suite and for the components that you have a license for.

Product or component	Administrator role	User role
WebSphere DataStage and WebSphere QualityStage	To validate the configuration of WebSphere DataStage and WebSphere QualityStage: 1. On the Windows system where you installed the clients, click Start → All Programs → IBM Information Server → WebSphere DataStage and QualityStage Administrator. 2. Log in to the Administrator client by using suite credentials. You will be prompted for the server where the domain is installed. Type the server name and the port number in the Domain field. For example, if the domain is installed on the xyz server and you used the default port number, you specify xyz:9080 as the server when you log in to the Administrator client. Type the WebSphere DataStage server name in the DataStage server field.	To start the Designer client: 1. On the Windows system where you installed the clients, click Start → All Programs → IBM Information Server → WebSphere DataStage and QualityStage Designer. 2. Log in to the Designer client by using suite credentials. You will be prompted for the server where the domain is installed. Specify the server name and the port number in the Domain field. For example, if the domain is installed on the xyz server and you used the default port number, you specify xyz:9080 as the server when you log in to the Designer client. You will also be prompted for a WebSphere DataStage server and project to attach to. Type the name of the project and the server name in the Project field. For example, if the WebSphere DataStage server is installed on the xyz server and you want to attach to the project named my-project, you specify xyz/my-project.

Product or component	Administrator role	User role
WebSphere Business Glossary	To validate the configuration of WebSphere Business Glossary: 1. Start Microsoft Internet Explorer and type the URL for the Web console. The URL takes the form, <i>host_server:port</i> where <i>host_server</i> is the URL or IP address of the computer where WebSphere Application Server is installed and <i>port</i> is the port number that is assigned to the Web console. The default port number is 9080. 2. Enable popups for this site in your Web browser. 3. Type the user name and password for the IBM Information Server Web console, and click Enter. 4. Click the Glossary tab.	To start WebSphere Business Glossary: 1. Start Microsoft Internet Explorer and type the URL for the Web console. The URL takes the form, <i>host_server:port</i> where <i>host_server</i> is the URL or IP address of the computer where WebSphere Application Server is installed and <i>port</i> is the port number that is assigned to the Web console. The default port number is 9080. 2. Enable popups for this site in your Web browser. 3. Type the user name and password for the IBM Information Server Web console, and click Enter. 4. Click the Glossary tab.

Product or component	Administrator role	User role
WebSphere Information Analyzer	To validate the configuration of WebSphere Information Analyzer: 1. On the Windows system where you installed the clients, click Start → All Programs → IBM Information Server → IBM Information Server Console and log in to the console. 2. In the Home navigator menu, select Configuration → Analysis Settings. 3. Click the Analysis Database tab. 4. In the Analysis Database tab, click Validate Settings. If you have configured the connection to the analysis database correctly, a window confirms that the settings are valid. 5. Click the Analysis Engine tab. 6. In the Analysis Engine tab, click Validate Settings. If you correctly configured the WebSphere DataStage user name and password that are used by the application server to access the WebSphere DataStage engine, a window confirms that the setting is valid. For the WebSphere DataStage credentials, you must use an operating-system level user name and password. Do not the suite user name and password. 7. On the computer where the WebSphere DataStage engine is installed, verify that the WebSphere Information Analyzer database DSN is configured and points to the WebSphere Information Analyzer database which might be local or remote.	To start WebSphere Information Analyzer: 1. On the Windows system where you installed the clients, click Start → All Programs → IBM Information Server → IBM Information Server console and log in to the console. You can log in to the console with credentials that have Suite Administrator authority or that have user authority. The console opens to the My Home workspace. 2. Expand the Analyzing information section in the Getting Started pane to begin creating a project, configuring resources, and analyzing information.

Product or component	Administrator role	User role
WebSphere Information Services Director	To validate the configuration of WebSphere Information Services Director 1. On the Windows system where you installed the clients, click Start → All Programs → IBM Information Server → IBM Information Server Console and log in to the console. You can log in to the console with credentials that have Suite Administrator authority or that have user authority. The console opens to the My Home workspace. 2. Expand the Enabling Information Services section in the Getting Started pane to begin deploying integration logic as services, cataloging and registering services, and setting up credentials.	To start WebSphere Information Services Director: 1. On the Windows system where you installed the clients, click Start → All Programs → IBM Information Server → IBM Information Server Console and log in to the console. You can log in to the console with credentials that have Suite Administrator authority or that have user authority. The console opens to the My Home workspace. 2. Expand the Enabling Information Services section in the Getting Started pane to begin deploying integration logic as services, cataloging and registering services, and setting up credentials.

Tutorials

Tutorials and sample data are provided to help you understand how to use features of the IBM Information Server suite.

The *IBM WebSphere QualityStage Tutorial* demonstrates the basic functions of IBM WebSphere QualityStage. You can learn WebSphere QualityStage process flow by designing and running jobs.

The exercises, sample data, and sample import files for the QualityStage tutorial are in the TutorialData\QualityStage\Data folder on the suite CD, DVD, or downloaded installation image.

The *IBM WebSphere DataStage Parallel Job Tutorial* demonstrates how to design and run IBM WebSphere DataStage parallel jobs.

The exercises, sample data, and sample import files for the parallel job tutorial are in the TutorialData\DataStage\parallel_tutorial folder on the suite CD, DVD, or downloaded installation image.

Chapter 6. Installing additional components

After you install IBM Information Server, you can install additional components.

Updating the license file

The suite installation wizard registers your license when you install IBM Information Server. If you purchase additional components after you install the suite, you must update your license file.

Before you begin

- You must have suite administrator authority to complete this task.
- You must use Microsoft Internet Explorer to access the Web console.

Procedure

To install license for suite components:

1. Ensure that your new license is on a computer that can run the Web console.
2. Log in to the Web console as a suite administrator.
 - Windows: Click **Start** → **IBM Information Server** → **IBM Information Server Web Console**.
 - Linux, UNIX: Access the Web console from a Windows computer using Microsoft Internet Explorer. Open the browser and go to the URL for the Web console. The URL takes the form, *domain_server:port*. *domain_server* is the URL or IP address of the server where the domain is installed. *port* is the port that is assigned to the Web console. The default port is 9080.
3. On the Administration tab, click **Domain management** → **Licensing**.
4. Click **Upload** and upload your license file.

Enabling the SAS connector (Linux, UNIX)

If you add the SAS connector to an existing suite installation, you must run the `install.liborchsasop` script to enable it.

Before you begin

You must update your license in the IBM Information Server Web console before you complete this task.

Procedure

To enable the SAS connector:

1. On the computer where the parallel engine is installed, change to the `/isHome/Server/PXEngine/install` directory. *isHome* is the directory where you installed the suite.
2. Run the `install.liborchsasop` script.
`./install.liborchsasop`

WebSphere MetaBrokers and bridges

After you install the IBM Information Server clients, you can install WebSphere MetaBrokers and bridges.

Before you begin

- Ensure that the system and software requirements are met before you install WebSphere MetaBrokers and bridges.
- You must have administrator access to the installation system to complete this task.

Procedure

To install WebSphere MetaBrokers and bridges:

1. Log on to the system where you installed the clients as an administrator.
2. In the root of the IBM Information Server CD, DVD, or downloaded installation image, open the MetaBrokersAndBridges directory and double-click **setup.exe**.
3. Follow the instructions in the wizard to install WebSphere MetaBrokers and bridges.

The Multi-Client Manager

You can install the Multi-Client Manager alone or as part of the upgrade process during an IBM Information Server suite installation. This topic describes how to install it alone.

Before you begin

- You must have administrator access to the installation system to complete this task.
- At least one WebSphere DataStage and WebSphere QualityStage client must be installed on the computer where you install the Multi-Client Manager.

Procedure

To install the Multi-Client Manager:

1. Log on as an administrator to the system where you installed the clients.
2. In the top-level directory of the IBM Information Server DVD, CD, or downloaded installation image, open the MultiClientManager directory and double-click **install.exe**.
3. Follow the instructions in the installation wizard to install the Multi-Client Manager.

MVS components for WebSphere DataStage

If you want to compile and run WebSphere DataStage jobs in MVS™, you must first install COBOL copybook files and the Runtime Library (RTL) in MVS.

The installation process for MVS components includes the following steps:

1. “Installing COBOL copybook files” on page 127
2. “The runtime library for MVS jobs” on page 128

3. Running the Installation Verification Process (IVP). Instructions are in the x:\Utilities\Supported\IVP\IVP_installation_instructions.doc file on the installation DVD.

Installing COBOL copybook files

Upload the COBOL copybook files to MVS to add them as members of a copybook library (copylib).

To install the COBOL copybook files:

1. Using FTP, upload the COBOL copybook files to copylib members. On the installation DVD, the files are in the \Mainframe Components folder.

Option	Description
Copy these files	To these copylib members
ArdtIMsg.cob	ARDTIMSG
ArdtMsg1.cob	ARDTMSG1
ArdtMsg2.cob	ARDTMSG2
ArdtUBgn.cob	ARDTUBGN
ArdtUCod.cob	ARDTUCOD
ArdtUDat.cob	ARDTUDAT
ArdtUEnd.cob	ARDTUEND
ArdtUFct.cob	ARDTUFCT
ArdtUFdf.cob	ARDTUFDF
RTLMsgs.cob	RTLMSGs

2. In the JCL templates for compiling the Cobol program, specify the name of the copylib that contains the uploaded copybook files.
 - a. Start the Designer client.
 - b. Under the Tools menu option, edit the CompileLink, DB2CompileLinkBind, and TDCompileLink JCL templates.
 - c. In the SYSLIB DD statement of each COBCOMP step, specify the copylib.

Copylib members

The copylib members for WebSphere DataStage contain message text or code templates that you can customize.

Message text

The following table describes copylib members that contain runtime message text.

Table 35. Copylib members with message text that you can customize

copylib member	Customizes
ARDTIMSG	Message text for IMS™ DB errors.
ARDTMSG1	Message text that displays after a job runs.
ARDTMSG2	Message text about file processing errors.
RTLMSGs	Message text about errors and warnings that the runtime library generates.

Templates

The following table describes copylib members that can control the initialization or termination of the generated COBOL program for a WebSphere DataStage job.

Table 36. Copylib members that can control the operation of generated COBOL programs

copylib member	Adds
ARDTUBGN	Program initialization code for the generated COBOL program.
ARDTUCOD	Paragraphs to the end of the generated COBOL program.
ARDTUDAT	Code before the PROCEDURE DIVISION statement of the generated COBOL program. For example, you can add WORKING-STORAGE variables and a LINKAGE SECTION.
ARDTUEND	Program termination code before each GOBACK statement.
ARDTUFCT	Statements, such as SELECT, to the end of FILE-CONTROL paragraph in the ENVIRONMENT division.
ARDTUFDF	FD statements at the end of the FILE SECTION in the DATA DIVISION.

The runtime library for MVS jobs

The runtime library for mainframe jobs contains routines that are used when a mainframe job is run.

To install the runtime library for MVS jobs:

1. Log on to MVS.
2. Allocate a runtime library data set.
 - a. In ISPF, access the function 3.2 (Data Set Utility).
 - b. In the Option line, type A.
 - c. Tab to the Data Set Name prompt, type a data set name, and press Enter.
 - d. In the Allocate New Data Set panel, in the Space units field, type TRKS.
 - e. In the Allocate New Data Set panel, type the following values:

Field	Value
Primary quantity	10
Secondary quantity	10
Directory blocks	0
Record format	FB
Record length	80
Block size	3120

3. Transfer the runtime library data file to the allocated runtime library dataset.
 - a. On your workstation, start FTP.
 - b. Connect to MVS and log on.
 - c. Set the transfer mode to BIN.
 - d. Use the PUT command to send the file ArdtRTL.data to the allocated runtime library dataset. This file contains the WebSphere DataStage runtime library in binary format.
4. In MVS, receive the runtime library dataset.
 - a. Run the TSO RECEIVE command `RECEIVE INDA('input_dataset_name')`

input_dataset_name is the dataset that you specified in Step 2.

- b. At the prompt, type DA('output_dataset_name')
output_dataset_name is the name of the dataset that the TSO RECEIVE command creates. This output dataset is a sequential file that is used by the job that you create and run in steps 5, 6, and 7.
5. Transfer the file ArdtRTL.jcl to MVS. This file contains sample JCL to create the runtime library.
 - a. On your workstation, start FTP.
 - b. Connect to MVS and log on .
 - c. Set the transfer mode to ASCII.
 - d. Use the PUT command to send the file ArdtRTL.jcl to a member of a JCL library.
6. In MVS, edit the JCL library member specified in step 5d.
 - a. Modify the JOB statement to meet your site standards.
 - b. In the SEQ1 statement, change the DSN value to the name of the output dataset that the TSO RECEIVE command creates in Step 4b. The output data set is the source dataset for an IEBCOPY operation.
 - c. In the PDS2 statement, change the DSN value to the name of the link library to contain the contents of the DataStage runtime library. The link library is the target for the IEBCOPY operation.
7. Submit the JCL file from the previous step to copy the dataset to the link library.
8. Start the Designer client.
9. Edit the JCL templates for CompileLink and DB2CompileLinkBind. In the LKED step, change the ARDTLIB DD statement to use the dataset name that you specified in the PDS2 statement in Step 6c.

Installation verification process

The installation verification Process (IVP) validates that an MVS environment is configured for compiling and executing MVS jobs.

The installation verification process uses three WebSphere DataStage jobs to validate an MVS environment.

DataStage job	Purpose
IVPFLAT	Tests the runtime library.
IVPDB2	Tests connectivity to DB2 Universal Database™ for z/OS® and the runtime library.
IVPIMS	Tests connectivity to IMS DB and the runtime library.

For information about installing and running the installation verifications process, see the x:\Utilities\Supported\IVP\IVP Installation Instructions.doc file on the installation DVD.

Installing additional components and updates for WebSphere Federation Server

You can add wrappers, install fix packs, and add tools to your federated system after the initial installation. You can also update the documentation in the DB2 Information Center.

Updating WebSphere Information Integration documentation

WebSphere Information Integration product documentation is available online in an information center or in PDF books.

The most up-to-date information is available in the online information center at <http://publib.boulder.ibm.com/infocenter/db2help/>.

Updates might also be available for topics in your locally-installed information center. To determine if there is an update available for a specific topic, compare the 'Last updated' date at the bottom of the locally installed topic to the same topic on the online information center. The 'Last updated' value can be found at the bottom of most topics.

For information on how to perform an update, see Updating the DB2 Information Center installed on your computer or intranet server.

Installing fix packs

You can install fix packs on an IBM WebSphere Information Integration system to get the latest updates and fixes for problems.

About this task

Fix packs contain updates, fixes for problems, and updates that are based on Authorized Program Analysis Reports (APARs). APARs are problems that are found during internal testing at IBM. The APARLIST.TXT file accompanies every fix pack. The APARLIST.TXT file describes the problem fixes that are contained in the fix pack.

Procedure

To install fix packs:

1. Obtain the latest WebSphere Information Integration fix pack from the support Web site at www.ibm.com/software/data/integration/db2ii/support.html.
2. Follow the instructions on the Web site to install the fix pack.

Adding the wrapper development kit to a system where DB2 is installed (Windows)

If DB2 Version 9 is installed on your system, you can add the software development kit (SDK) for developing wrappers in C++ and Java. If an earlier version is installed, migrate or upgrade to DB2 Version 9 to install the wrapper development kit.

Procedure

To add the wrapper development kit to a system where DB2 Version 9 is installed:

1. Stop all of the DB2 services that are running.
2. Click **Start** → **Settings** → **Control Panel** → **Add/Remove Programs**.
3. In the Currently installed programs window, click **DB2 Enterprise Server Edition**, and click **Change**. The DB2 Setup wizard opens.
4. Click **Next**. The Program Maintenance page opens.
5. Click **Modify**.
6. Click **Next**. The "Select the features you want to install" page opens.
7. Expand the **Application Development tools** tree.
8. Click **Base Application Development Tools**, and select an installation option from the drop-down list.
9. Optional: To install the sample wrapper, click **Sample Applications**, and select an installation option from the drop-down list.
10. Follow the instructions in the wizard to complete the installation.

Adding the wrapper development kit to a system with DB2 installed (UNIX)

If DB2 Version 9 is installed on your system, you can install the software development kit (SDK) for developing wrappers in C++ and Java without reinstalling DB2. If an earlier version is installed, migrate or upgrade to DB2 Version 9 to install the wrapper development kit.

Before you begin

Before you install a DB2 client on UNIX:

- Ensure that your system meets all of the memory, hardware, and software requirements to install your DB2 product.
- If you are installing a DB2 client on Solaris, update your kernel configuration parameters and restart your system.

Procedure

To add the wrapper development kit to a system with DB2 installed:

1. Log in as a user with root authority.
2. Insert and mount the appropriate CD.
3. Change to the directory where the CD is mounted by entering the `cd /cdrom` command where `/cdrom` is the CD mount point.
4. Enter the following command at a command prompt:
`./db2setup`

The DB2 Setup wizard opens.

5. Click **Install Products**. The "Select the products you want to install" page opens.
6. Expand the **Application Development tools** tree.
7. Click **Base Application Development Tools**, and select an installation option from the drop-down list.
8. Follow the instructions in the wizard to complete the installation.

Chapter 7. Troubleshooting installations

In addition to the topics in this chapter, you can find troubleshooting information in the information centers and technotes for some product modules and components.

Installation log files

The installation wizard creates log files during the installation process. You can use the log files to troubleshoot installation errors. By default, log files are stored in the InformationServer\logs directory.

If the installation fails, log files are saved to `x:\Documents and Settings \ user\ Local Settings \Temp\ibm_is_logs\`. *x* is the drive where your Documents and Settings folder is located. *user* is the user that the installation was running under.

Other log files

Log files provide a useful source of information when you are investigating an issue.

Client log files

WebSphere DataStage creates log files that provide information about recent sessions. The log files are in `x:\Documents and Settings\user name\ds_logs\datastage_wrapper_trace_nn.log`. *x* is the drive on which IBM Information Server is installed. *user name* is the Windows logon name. *nn* represents the number of DataStage sessions (up to 20) that logs are available for.

Log files for the application server

The application server records information in the following log files:

- `x:\IBM\WebSphere\AppServer\profiles\default\heapdump.date.phd`
- `x:\IBM\WebSphere\AppServer\profiles\default\javacore.date.txt`
- `x:\IBM\WebSphere\AppServer\profiles\default\*.log`
- `x:\IBM\WebSphere\AppServer\profiles\default\server1\*.log`

x is the drive on which the IBM Information Server domain layer is installed. *date* is the date of the session.

Troubleshooting the installation process when it appears to be stopped

The installation might appear to be stopped if the installation requires user input and the window that requires input is hidden.

If the installation wizard opens a window that requires user input, the main installation task is stopped until you respond to the window. If you are running other programs during the installation and return to the installation wizard by clicking the main installation window or by clicking the task bar, the window might be hidden by the wizard. If the installation wizard appears to be stopped, verify that a window that requires user input is not hidden behind another window.

Troubleshooting command-line scripts that fail

Certain characters in user names or passwords might cause command-line scripts to fail.

Avoid the following characters in user names and passwords: % ^ &) " < > \ | ! \$.

Each operating system has its own requirements. Review your operating system documentation for rules about user names and passwords. You can find naming conventions for DB2 database systems and WebSphere products in their respective information center or online.

DB2 database systems

publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp

WebSphere products

publib.boulder.ibm.com/infocenter/wasinfo/v6r0/index.jsp

Runtime error during removal of the ASBAgent component

You might receive a runtime error when the uninstall process is removing the ASBAgent.exe component.

This problem with the Visual C++ Runtime dll (msvcrt.dll) is documented in the Microsoft Knowledge Base at support.microsoft.com/kb/884538. The uninstallation process completes successfully with no errors encountered in the process.

Troubleshooting the user mapping plugin for federated systems

After you configure the federated server to use the user mapping plugin, there are several ways that you can troubleshoot problems.

Before you test the plugin in the federated server, you should develop a Java application to test if the plugin can connect to and retrieve user mappings from the external repository. If the test is successful, then you can try to integrate the plugin with the federated server.

JAVA_HEAP_SZ parameter

The JAVA_HEAP_SZ parameter determines the maximum size of the heap that is used by the Java interpreter to service Java stored procedures and user defined functions.

If you encounter memory problems, try setting the heap size to 1024 or higher.

Error messages

If an error message is returned when using the plugin, the information in the error message can help you to determine the problem. There are reason codes and user responses in the error messages that should help you determine the problem.

SQL20349N

The federated server cannot access the user mappings that are in the external repository. One reason why you might receive this error is that the user mapping plugin is not in the correct location. If you use the sample LDAP plugin on a UNIX federated server, ensure that

UserMappingRepositoryLDAP.jar file and the UserMappingRepositoryLDAP.cfg file are in sqllib/function directory.

SQL20350N

The authentication at the user mapping repository failed. The proper credentials, such as passwords or certificates, to retrieve the user mappings were not communicated to the user mapping repository.

SQL20351W

You specified the DB2_UM_PLUGIN option on a wrapper and the DB2_UM_PLUGIN option is already specified on a server definition that uses the wrapper. The value that is set on the server definition takes precedence over the value that you set on the wrapper.

SQL20352W

You specified DB2_UM_PLUGIN option for the server. User mappings for this server are read from the external user mapping repository. The CREATE USER MAPPING, ALTER USER MAPPING, and DROP USER MAPPING statements only affect user mappings in the federated catalog table.

db2diag.log file

The diagnostic log is a file that contains text information that is logged by the federated database. This information is used for problem determination and is intended for IBM customer support.

You can view the file with a text editor. You can also use the db2diag tool can help you to analyze the information that is contained in the db2diag.log file to troubleshoot problems with your user mapping plugin.

Example 1

In the db2diag.log file, search for the term JNI_Env::check_exception. This term marks the first error record for the user mapping plugin. The record indicates the line in the plugin source code where an error has occurred. In the following example, the line in the code that is in error is line 119 of the UserMappingRepositoryLDAP.java file.

```
2006-02-18-19.05.43.966176-480 I6780A253          LEVEL: Severe
PID      : 3932214          TID : 2828          PROC : db2fmp (Java)
INSTANCE: einstein        NODE : 000
FUNCTION: DB2 UDB, Query Gateway, Sqlqg_JNI_Env::check_exception, probe:150
2006-02-18-19.05.43.966953-480 I7034A388          LEVEL: Error
PID      : 3932214          TID : 2828          PROC : db2fmp (Java)
INSTANCE: einstein        NODE : 000
FUNCTION: DB2 UDB, Query Gateway, Sqlqg_JNI_Env::check_exception, probe:190
DATA #1 : String, 108 bytes
com.ibm.ii.um.UserMappingException
        at UserMappingRepositoryLDAP.<init>(UserMappingRepositoryLDAP.java:119)
```

Example 2

In the db2diag.log file, search for the term Error fetching user mappings. The record for this error includes the plugin name, the instance name, the database name, the remote server name, and the local authentication ID. These values are passed from the federated server to the plugin as input. Ensure that the values , including the case sensitivity, are correct.

```
2006-02-18-19.05.43.967308-480 I7423A371          LEVEL: Severe
PID      : 3932214          TID : 2828          PROC : db2fmp (Java)
INSTANCE: einstein        NODE : 000
MESSAGE : Error fetching user mappings through plugin:
```

```
DATA #1 : String, 110 bytes  
plugin=UserMappingRepositoryLDAP,instance=EINSTEIN,database=SAMPLE,  
remote_server=DRDASERVER1,local_authid=NEWTON
```

db2trc - trace command

If you are unable to troubleshoot the problem, you can use db2trc command to generate a trace. You will need to send the trace record to IBM for analysis.

Chapter 8. Removing IBM Information Server

It might be necessary to remove IBM Information Server from your computers for various reasons, such as upgrading to a newer computer or to clean up a test computer.

Removing IBM Information Server from your system (Windows)

You can use the Add or Remove Programs utility in the Windows Control Panel to remove the suite from your system.

Before you begin

You must have administrator access to the installation system to complete this task.

About this task

If you installed the suite with DB2, and if the suite installation wizard has created a metadata repository database, then the metadata repository database is dropped when IBM Information Server is removed. If you used Oracle or Microsoft SQL Server for the metadata repository, you must drop the database manually. The default database name is xmeta.

Procedure

To remove IBM Information Server from Windows:

1. Use the Add or Remove Programs utility in the Windows Control Panel to remove the following software.
 - a. Any MetaBrokers and bridges that are installed.
 - b. IBM Information Server. If multiple versions of WebSphere DataStage clients are installed on your computer, you must enable another client version in the Multi-Client Manager after you uninstall WebSphere DataStage.
2. Optional: If you intend to reinstall IBM Information Server, you must manually remove the following files from your system, if they still exist. (In the following steps, *x* represents the drive where the suite was installed.)
 - a. Delete the *x*:\\IBM\\InformationServer directory.
 - b. If WebSphere Application Server was installed by the suite installation wizard, remove the application server directory. The default directory is C:\\IBM\\WebSphere.
 - c. If you used an existing installation of WebSphere Application Server, remove the profile that you created for the suite.
 - d. If DB2 was installed by the suite installation wizard, delete the following directories:
 - \\DB2Home. DB2Home is the directory where DB2 is installed. The default directory is C:\\IBM\\sqlib.
 - *x*:\\DB2.
 - e. Delete all of the entries that contain references to the application server installation directory. The default directory is *x*:\\IBM\\WebSphere\\AppServer. If the file is empty, delete the file.
3. Restart your computer.

Removing profiles (Windows)

If you used an existing installation of WebSphere Application Server to host WebSphere Metadata Server, you must remove the profile that you created for the metadata server when you remove the suite from your system.

If the suite installation wizard installed the application server, the profile is removed when you remove the suite from your computer.

Procedure

To remove profiles from Windows:

1. Click **Start** → **All Programs** → **IBM WebSphere** → **Application Server** → **Profiles** → *profileName* → **Stop the server**. *profileName* is the name of the profile to be deleted.
2. Open a Windows command prompt.
3. Change to the *WASHome*\bin directory. *WASHome* is the directory where the application server is installed. The default directory is C:\Program Files\IBM\WebSphere\AppServer.
4. Run the following command.
`wasprofile.bat -delete -profileName profile`

profile is the name of the profile to be deleted. The profile that is created by the suite installation wizard is named default.

5. Delete the directory for the profile. The directory is in the *WASHome*\profiles directory.

Removing the Multi-Client Manager (Windows)

You can remove the Multi-Client Manager from your computer.

Before you begin

You must have administrator access to the installation system to complete this task.

Procedure

To remove the Multi-Client Manager from your computer:

1. Open the Add or Remove Programs utility.
2. Select **DataStage Multi-Client Manager** from the list of programs, and click **Remove**.

Removing IBM Information Server from your system (Linux, UNIX)

You can remove the suite from your system with the product uninstallation utility.

Before you begin

You must have root authority on the system to complete this task.

About this task

If the suite installation wizard installed the database server, the metadata repository database is dropped when the application server is removed. The default name of the metadata repository database is xmeta.

Procedure

To remove the suite from Linux and UNIX:

1. Run the following commands to start the wizard that removes IBM Information Server.

```
cd /opt/IBM/InformationServer/_uninst/suite
./uninstall
```

2. Follow the instructions in the wizard to remove the suite.
3. Delete the `/opt/IBM/InformationServer` directory.
4. If you installed WebSphere Application Server using the wizard, delete the application server directory. The default directory is `/opt/IBM/WebSphere/AppServer`.
5. If you installed DB2 using the wizard, remove the following directories.
 - The default directory for AIX, is `/opt/IBM/db2/V9.1`. The default directory for Linux is `/opt/ibm/db2/V9.1`.
 - `/DB2Inst/sqllib` where *DB2Inst* is home directory of the instance owner. The default directory is `/db2home/db2inst1`.

If the application server was not installed by the suite installation wizard, you must remove the profile for the metadata server.

Removing profiles (Linux, UNIX)

If you used an existing installation of WebSphere Application Server to host WebSphere Metadata Server, you must remove the profile you created for the metadata server when you remove the suite from your system.

If the suite installation wizard installed the application server, the profile is removed when you remove the suite from your system.

Procedure

To remove profiles from Linux and UNIX:

1. Open a Linux or UNIX shell prompt.
2. From a Linux or UNIX shell prompt, change to the *WASHome*/bin.
WASHome is the directory where the application server is installed. The default directory is `/opt/IBM/WebSphere/AppServer`
3. Run the following command to delete the profile.

```
./wasprofile.sh -delete -profileName profile
```

profile is the name of the profile to be deleted.
4. Delete the directory for the profile. The directory is in the *WASHome*\profiles directory.

Removing the WebSphere QualityStage migration script from your system

You can remove the WebSphere QualityStage migration script from your system.

Procedure

To remove the WebSphere QualityStage migration script from your system:

Delete the directory where the script is installed.

Accessing information about the product

IBM has several methods for you to learn about products and services.

You can find the latest information on the Web:
www.ibm.com/software/data/integration/info_server/

To access product documentation, go to publib.boulder.ibm.com/infocenter/iisinfsv/v8r0/index.jsp.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at www.ibm.com/shop/publications/order.
- To order publications by telephone in the United States, call 1-800-879-2755.

To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at www.ibm.com/planetwide.

Providing comments on the documentation

Please send any comments that you have about this information or other documentation.

Your feedback helps IBM to provide quality information. You can use any of the following methods to provide comments:

- Send your comments using the online readers' comment form at www.ibm.com/software/awdtools/rcf/.
- Send your comments by e-mail to comments@us.ibm.com. Include the name of the product, the version number of the product, and the name and part number of the information (if applicable). If you are commenting on specific text, please include the location of the text (for example, a title, a table number, or a page number).

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing 2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

(C) (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. (C) Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM trademarks and certain non-IBM trademarks are marked at their first occurrence in this document.

See www.ibm.com/legal/copytrade.shtml for information about IBM trademarks.

The following terms are trademarks or registered trademarks of other companies:

Adobe, Acrobat, Portable Document Format (PDF), and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Cell Broadband Engine and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Intel[®], Intel logo, Intel Inside[®] logo, Intel Centrino[™], Intel Centrino logo, Celeron[®], Intel Xeon[™], Intel SpeedStep[®], Itanium[®] and Pentium[®] are trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT[®] and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL[®] is a registered trademark and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library[®] is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Other company, product or service names may be trademarks or service marks of others.

Index

Special characters

/etc/magic
editing 72

A

access plans
 federated systems 99
accessibility 141
analysis databases
 configuring 120
application server
 creating server profiles 30, 40, 44
 creating users and groups for installation
 Windows 31, 43, 44
 removing profiles
 Linux, UNIX 139
 Windows 138
Application Server log files 133
application server profiles
 creating 30, 40, 44
architecture
 domain 8
 engine 9
architecture, IBM Information Server 1
AuditStage
 installation requirements 20
authentication
 configuring PAM 65

B

BCPLoad
 stages 66
BCPLoad stage 66
bridges
 installing 126

C

C++ compilers
 setting environment variables for
 Microsoft Visual Studio .NET 2003 C++ 61
 Windows 20
client log files 133
clients
 described 6
 licensing controls 15
COBOL copybook files
 installing 127
code pages
 conversion 99
 federated database 98
 federated systems 99
code sets
 federated database 98

collating sequences
 federated database 98
 federated systems 99
common connectivity 1
common services 1
component installer
 described 9
components
 enabling 125
CONCENTRATOR parameter 96
configuration files
 examples
 PAM 65
configuring 65
 plug-in stages 65
connectivity
 DB2 79
 ODBC drivers 14
 testing ODBC drivers 87
connectors
 enabling SAS 125
console installations
 syntax 48
console mode
 installing 48
 Linux, UNIX 51
 installing the suite 18
creating
 federated databases 97
credential mappings
 creating 53

D

data sources
 configuring ODBC access 83
 Linux, UNIX 83
 Windows 83
databases
 federated, creating 97
DataStage users, setting up 59
DB2
 creating users and groups for installation
 Windows 31, 43, 44
DB2 database systems
 configuring WebSphere DataStage 78
DB2 Setup program
 running 79
db2nodes.cfg file 97
disk space requirements
 Installing IBM Information Server on
 Linux, UNIX 20
distributed database management
 system 11
dijlink script 95
documentation
 accessible 141
 accessing 13
 IBM Information Server 12

domain
 order of installation 19
 preparing to install 29
 Windows system requirements 20
driver
 confirming and registering the
 Microsoft SQL Server client 101
dsenv file
 sample entries 84

E

engine
 order of installation 19
 preparing to install 30
 Windows system requirements 20
engine layer
 overview 9
Enterprise stage jobs
 granting user privileges 79
environment variables
 configuring for WebSphere DataStage
 server 64
 parallel engine 72, 75
 setting for Microsoft Visual Studio
 .NET 2003 C++ 61
 setting for the parallel engine 72
 setting for WebSphere DataStage
 server 84
 WebSphere DataStage libraries 64
external user registry
 creating credential mappings 53

F

federated databases
 code sets 98
 collating sequence 98, 99
 creating 97
FEDERATED parameter 96
federated server
 setup 92
federated systems
 configuring 53
 overview 11
fix packs
 installing 130
fsz
 settings
 AIX 42

G

graphical mode
 installing
 Linux, UNIX 50
 Windows 47
 installing the suite 18

H

- heapsize
 - settings
 - AIX 42
- help
 - installing 12

I

- IBM Information Server
 - accessing documentation 13
 - architecture and concepts 1
 - installing
 - Linux, UNIX 50
 - Windows 47
 - installing documentation 12
 - installing silently
 - Linux, UNIX 51
 - Windows 48
 - ODBC drivers 14
 - planning to install 19
 - removing
 - Linux, UNIX 138
 - Windows 137
 - user registries 18
- IBM Information Server directory service
 - configuring 89, 90
- IBM Information Server Web console
 - described 8
- information center
 - installing 12
- Informix
 - djxlinkInformix script 95
- installation
 - planning 19
 - preparing Windows computers 21
- installation methods
 - supported 18
- installing
 - console syntax 48
 - in console mode 48
 - Linux, UNIX 51
 - silent installation syntax 49, 52
- installing in graphical mode 50
 - Windows 47
- internal user registries
 - configuring for IBM Information Server 90
- internal user registry
 - configuring for IBM Information Server 90
 - creating credential mappings 53

J

- jobs
 - configuring the parallel environment 74

K

- kernel parameters
 - Linux, UNIX 41

- kernels
 - configuring
 - AIX 42

L

- layers
 - client 6
 - domain 8
 - engine 9
 - order of installation 19
 - product components contained in 4
- LDAP
 - configuring 91
 - configuring for IBM Information Server 88
- legal notices 143
- libcIntsh
 - creating 42
- libraries
 - environment variables for WebSphere
 - DataStage 64
 - wrappers 94
- library files
 - link-edit 92
 - wrappers
 - error message files 94
 - link manually 95
 - wrappers, check installation 93
- library path environment variables
 - WebSphere DataStage libraries 84
- license
 - upgrading 15
- licenses
 - installing 125
- link-edit
 - error message files 94
 - wrapper library files 92
- Linux, UNIX 50
- local operating system user registries
 - configuring LDAP 88
- local operating system user registry
 - configuring 89, 90
- Local operating system user registry
 - configuring the IBM Information Server directory service 89, 90
- log files
 - Application Server 133
 - client 133

M

- magic
 - editing the /etc/magic file 72
- mainframe components
 - installing 127
- Mainframe jobs
 - installing runtime libraries 128
- mainframes
 - running jobs 129
- massively parallel processing
 - configuring a node (Windows) 62, 76
- memory
 - setting for parallel jobs
 - AIX 70

- metabrokers and bridges
 - installing 126
- metadata repository
 - configuring Oracle 25, 35
 - creating users and groups for installation
 - Windows 31, 43, 44
 - Oracle database 25, 26, 27, 36, 37
 - order of installation 19
 - preparing to install
 - Linux, UNIX 34
 - Windows 23
 - repository 12
 - SQL Server database 24, 25
 - Windows system requirements 20
- Microsoft SQL Server
 - djxlinkMssql script 95
- Microsoft SQL Server client
 - confirming and registering 101
- Microsoft Visual Studio .NET 2003 C++
 - setting environment variables 61
- migration 15
- migration utility
 - removing for WebSphere
 - QualityStage 139
- Multi-Client Manager
 - installing 126
 - removing 138
- MVS
 - verifying the installation of MVS components 129
- MVS components
 - installing 127
- MVS jobs
 - installing runtime libraries 128

N

- NLS
 - configuring for WebSphere DataStage and QualityStage 66
- NLS (national language support)
 - federated systems 99

O

- ODBC data sources
 - configuring access
 - Linux, UNIX 83
 - Windows 83
- ODBC drivers
 - testing connectivity 87
 - WebSphere DataStage and WebSphere QualityStage 14
- operators
 - enabling 125
- Optimizer, query 99
- Oracle
 - configuring as the metadata repository 25, 35
 - djxlinkOracle script 95
- Oracle 10g
 - metadata repository 25, 26, 27, 36, 37
- Oracle Enterprise Stage
 - running WebSphere DataStage jobs 82

- Oracle gencntsh script
 - editing 42
- Oracle parallel server
 - granting access 81

P

- paging space
 - setting 71
 - setting on AIX 71
 - setting on Linux 71
- PAM
 - configuring for WebSphere DataStage and QualityStage 65
 - example configuration files 65
- parallel engine
 - configuring 74
 - configuring ssh 67
 - copying 69
 - editing the /etc/magic file 72
 - environment variables 72, 75
 - monitoring system resources 9
 - remote shells 67
 - running rsh without a password 68
 - setting environment variables 72
 - setting paging space 71
 - AIX 71
 - Linux 71
 - tutorials 124
- parallel jobs
 - setting the maximum memory
 - AIX 70
 - setting the maximum number of processes 69
 - user privileges 74, 76
- parallel processing
 - configuring a node (Windows) 62, 76
- parameters
 - kernel parameters
 - Linux, UNIX 41
- permissions and groups 54, 56
- Permissions page 59
- plug-in stages 65
- plug-ins
 - enabling 125
- processes
 - setting the maximum number for parallel jobs 69
- processing node
 - Windows 62, 76
- product licenses
 - installing product components 15
- products components
 - layers contained in 4
- profiles
 - creating application server
 - profiles 30, 40, 44
 - removing
 - Linux, UNIX 139
 - Windows 138

Q

- queries
 - access plans 99

- query optimizer
 - federated systems 99

R

- registries
 - user registries 18
- remote shell
 - running rsh without a password 68
- repository
 - metadata repository 12
 - preparing to install
 - Linux, UNIX 34
 - Windows 23
- Resource Tracker
 - monitoring system resources 9
- response files
 - installing IBM Information Server
 - Linux, UNIX 51
 - Windows 48
- rsh
 - parallel engine 67
- rsh command
 - running without a password 68
- running with Oracle Enterprise Stage 82
- runtime components
 - preparing to install 30
- runtime engine
 - starting the WebSphere DataStage server
 - Windows 49
 - starting the WebSphere DataStage server (Linux, UNIX) 52, 64
 - stopping the WebSphere DataStage server
 - Windows 49
 - stopping the WebSphere DataStage server (Linux, UNIX) 52, 64
- runtime libraries
 - installing for MVS jobs 128

S

- SAS connector
 - enabling 125
- screen readers 141
- Secure Shell
 - configuring for the parallel engine 67
- security
 - configuring for IBM Information Server 87
 - configuring PAM 65
 - configuring the IBM Information Server directory service 89, 90
 - example PAM configuration files 65
- server ID
 - application server authentication
 - Linux, UNIX 16
 - Windows 15
- shmttest utility
 - options 61
 - syntax 61
- tuning the WebSphere DataStage and WebSphere QualityStage server 60

- silent installations
 - installing IBM Information Server
 - Linux, UNIX 51
 - Windows 48
 - syntax 49, 52
- silent mode
 - installing the suite 18
- SQL Server 2005
 - metadata repository 24, 25
- ssh
 - configuring for the parallel engine 67
- stack size
 - settings
 - AIX 42
- suite
 - installing
 - Linux, UNIX 50
 - Windows 47
 - removing
 - Linux, UNIX 138
 - Windows 137
- suite users
 - mapping between the WebSphere DataStage and WebSphere QualityStage server and the suite 53
- Sybase
 - djxlinkSybase script 95
- sys.gv_\$instance table
 - granting access 81
- sys.v_\$cache table
 - granting access 81
- system requirements
 - Windows 20

T

- Teradata
 - djxlinkTeradata script 95
- trademarks 145
- troubleshooting
 - appears stopped 133
 - command-line script fails 134
- tutorials
 - parallel engine 124
 - WebSphere QualityStage 124

U

- unified metadata 1
- unified parallel processing engine 1
- unified user interfaces 1
- updates
 - installing 130
- user privileges
 - Enterprise stage jobs 79
 - parallel jobs 74, 76
- user registries
 - configuring 87
 - LDAP 16, 91
 - local operating system 15
- user roles 59
- user user registries
 - internal user 15

- users
 - mapping credentials 57
- users and groups
 - creating for installation
 - Windows 31, 43, 44
 - required
 - Linux, UNIX 16
 - Windows 15

V

- validating the installation 121
 - WebSphere Business Glossary 121
 - WebSphere Information Analyzer 121
 - WebSphere Information Services Director 121

W

- WebSphere Application Server
 - creating server profiles 30, 40, 44
 - creating users and groups for installation
 - Windows 31, 43, 44
 - removing profiles
 - Linux, UNIX 139
 - Windows 138
 - user registry 18
- WebSphere Business Glossary
 - described 8
 - preparing to install 29
- WebSphere DataStage
 - configuring 53
 - configuring NLS 66
 - configuring on DB2 servers 78
 - installing the Multi-Client Manager 126
 - installing the parallel engine on multiple computers 69
 - removing the Multi-Client Manager 138
 - tutorials 124
- WebSphere DataStage and WebSphere QualityStage 121
 - tuning the server 60
- WebSphere DataStage jobs 82
- WebSphere DataStage server
 - configuring the environment 64
 - mapping credentials 57
 - preparing to install 30
 - shmttest utility syntax and options 61
 - starting
 - Windows 49
 - starting (Linux, UNIX) 52, 64
 - stopping
 - Windows 49
 - stopping (Linux, UNIX) 52, 64
- WebSphere Federation Server
 - configuring 53
 - described 9
 - ODBC drivers 14
 - preparing to install 30
- WebSphere Information Analyzer
 - AuditStage 20
 - configuring 53
- WebSphere Information Analyzer
 - (continued)
 - configuring analysis databases 120
- WebSphere Information Integration
 - Oracle genclntsh script
 - libclntsh file 42
- WebSphere Metabrokers and bridges
 - installing 126
- WebSphere Metadata Server
 - mapping credentials 57
 - preparing to install 29
- WebSphere QualityStage
 - configuring 53
 - Configuring NLS 66
 - configuring the environment 64
 - removing the migration utility 139
 - tutorials 124
- Windows
 - preparing computers for installation 21
- wrapper development kit
 - installing on UNIX 131
 - installing on Windows 130
- wrapper library files
 - check installation 93
 - link-edit 92
- wrappers
 - library files 94
 - error message files 94
 - link manually 95
 - link-edit
 - error message files 94



Printed in USA

GC19-1048-01



Spine information:

IBM Information Server Version 8.0.1

IBM Information Server Planning, Installation, and Configuration Guide

