IBM InfoSphere Information Server
Version 11 Release 3

Planning, Installation, and Configuration Guide



GC19-4316-00

IBM InfoSphere Information Server
Version 11 Release 3

Planning, Installation, and Configuration Guide



GC19-4316-00

Note

Before using this information and the product that it supports, read the information in "Notices and trademarks" on page 453.

© Copyright IBM Corporation 2005, 2014.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Part 1. Planning to install IBM InfoSphere Information Server 1

Chapter 1. Tiers and components 3
Client tier 3
Services tier 4
Engine tier 6
Repository tier 9
Tier relationships 11

Chapter 2. Planning to install IBM InfoSphere Information Server 13
Understanding your installation options 13
- Options for installing the database software . . 13
- Options for installing the application server software 14
- Choosing a topology 17

Reviewing installation checklists 63
- Basic installation checklist 63
- Basic isolated engine tier and parallel engine installation checklist 69
- High availability (active-passive configuration) installation checklist 76
- High availability (clustered configuration) installation checklist 83

Setting up the tier computers 94
- Reviewing system requirements and release notes 94
- Running ISA Lite to check prerequisites 94
- Directories that must be present on all tier computers before installation 95
- Configuring your network 96
- Modifying kernel parameters and user limits (AIX Linux) 101
- Directory and user permissions 103
- Setting up administrator and owner accounts 107
- Pre-installing databases 113
- Setting up a preinstalled application server . . 148
- Setting the locale and character encoding for the engine tier target computer 163
- Creating a two-server active-passive high availability configuration 166

Part 2. Installing IBM InfoSphere Information Server software 191

Chapter 3. Preparing to install software 193
Selecting the installation order in a multi-computer topology 193
Obtaining IBM InfoSphere Information Server software 193
Configuring your browser 194
Preparing to install InfoSphere Metadata Asset Manager and bridges to import and export metadata 195
Preparing to run the installation program 197

Chapter 4. Installing IBM InfoSphere Information Server software 201
Installing IBM InfoSphere Information Server by using the wizard (graphical mode) 201
- Taking action if the installation prerequisites check fails 205
- Specifying tiers and product modules to install 206
- Running the installation step 211
- Taking action if the installation program encounters an error 213
- Completing the installation process 214

Installing IBM InfoSphere Information Server in silent mode by using the response file 215
- Creating a response file 215
- About the response files 217
- Starting the installation program with a response file 240

setup command 244
Monitoring log files to follow installation progress 246
- Monitoring the InfoSphere Information Server installation 247
- Monitoring the WebSphere Application Server Network Deployment installation 248
- Monitoring the WebSphere Application Server Network Deployment application deployment . 248
- Monitoring the WebSphere Application Server Liberty Profile application deployment 249
- Monitoring InfoSphere DataStage component installations 250

Verifying and testing the installation 251
- Running the ISALite tool to verify the installation and troubleshoot 251
- Testing the installation of IBM InfoSphere DataStage and IBM InfoSphere QualityStage . . 252
- Testing the IBM InfoSphere Information Services Director installation 254

Chapter 5. Installing components that are not in the suite installation program 257
Installing the IBM InfoSphere Glossary Anywhere client 257
- Preparing for enterprise-wide distribution . . . 257
- Configuration parameters in the config.ini file 258
- Installing the client 263
- Uninstalling IBM InfoSphere Glossary Anywhere 268
- Deploying the client to large groups 269
- Retaining settings for InfoSphere Glossary Anywhere 270

© Copyright IBM Corp. 2005, 2014

Installing IBM InfoSphere Information Governance Catalog for Eclipse 271

Part 3. Configuring InfoSphere Information Server software 277

Chapter 6. Configuring IBM InfoSphere Information Server 279

Configuring security for IBM InfoSphere Information Server 279
WebSphere Application Server cluster configuration 280
Configuring a single front-end web server . . . 280
Connecting the front-end web server with InfoSphere Information Server 281
Changing the front-end dispatcher location . . 283
Registering and creating data stores after installation 285
Changing configuration of databases and data stores 287
Changing the metadata repository database host name and port 290
Changing configuration of other data stores to high availability 297
Modifying the connection properties for data sources for the InfoSphere QualityStage Standardization Rules Designer database . . . 299
Changing application server host names and ports 299
Changing WebSphere Application Server Network Deployment HTTP server and transport ports 299
Changing other WebSphere Application Server Network Deployment ports 302
Changing WebSphere Application Server Liberty Profile ports 304
Setting up firewall exceptions 306

Chapter 7. Configuring a parallel processing environment. 309

Adding computers to share engine processing . . 309
Configuring computers to be compute nodes for in an MPP system (AIX Linux) 309
Configuring computers to be compute nodes for MPP (Windows) 313
Distributing transform libraries to compute nodes 321
Setting up a C++ compiler 321
Setting up a compiler on IBM AIX 322
Setting up a GCC compiler (Linux on 64-bit Intel systems) 323
Setting up a Microsoft Visual Studio .NET C++ compiler 323
Setting environment variables for the parallel engine 324
Environment variables for the parallel engine 324
Configuring the parallel engine (AIX Linux) . . . 326
Setting the maximum number of processes for parallel jobs (AIX) 326
Setting per-user kernel parameters for parallel jobs (AIX) 327
Setting the paging space for the parallel engine 328
Updating the /etc/magic file 329
Configuring the parallel engine (Windows Server) 330
Configuring the Windows registry: Enabling auto-tuning for paged pool size 330
Configuring the Windows registry: Enabling auto-tuning for system page size 331
Configuring the Windows registry: Setting the threshold for de-committing memory 331
Configuring the Windows registry: Increasing the noninteractive desktop heap size 332
Configuring the Windows registry: Specifying TCP/IP settings (Windows Server 2003) . . . 333
Configuring Windows: Editing the boot.ini file 334
Configuring database connectivity in a parallel environment. 335

Chapter 8. Configuring product modules in the suite 337

Configuring IBM InfoSphere DataStage and IBM InfoSphere QualityStage 337
Configuring the engine to use PAM (Linux, UNIX) 338
Ensuring that InfoSphere DataStage users have the correct localization settings (AIX Linux) . . 339
Configuring ODBC access (AIX Linux) 341
Configuring ODBC access (Microsoft Windows) 345
Tuning the InfoSphere Information Server engine for large numbers of users or jobs (Windows Server) 345
Configuring IBM InfoSphere QualityStage Match Designer 346
Setting up application servers to run investigation, SQA, or Match Statistics reports . 347
Configuring the Oozie Workflow Activity stage (Linux) 348
Configuring IBM InfoSphere Information Governance Catalog 349
Stopping and starting the IBM InfoSphere Information Server server engine 351

Chapter 9. Processes used by InfoSphere Information Server 353

Part 4. Applying updates or adding tiers or products 357

Chapter 10. Updating software 359

Chapter 11. Adding product modules to an installation 361

Chapter 12. Adding additional tiers to an installation 365

Part 5. Removing IBM InfoSphere Information Server software 367

Chapter 13. Preparing to remove software 369
Shutting down tasks and applications before removing software (service tier and engine tier) . . 370
Configuring your browser for software removal 371

Chapter 14. Removing IBM InfoSphere Information Server software by using the software removal program 373
Removing individual product modules 374
Removing IBM InfoSphere Information Server by using the software removal program (graphical mode) 376
Specify options to run the software removal program 379
Saving a response file and removing the product 379
Removing IBM InfoSphere Information Server by using a response file (silent mode) 380
Software removal program encounters an error . . 383

Chapter 15. Completing software removal 385
Completing software removal (Windows) 385
Removing leftover files from the services tier (Windows) 385
Removing leftover files from the engine tier (Windows) 386
Removing leftover files from the repository tier (Windows) 387
Completing software removal (AIX Linux) . . . 388
Removing leftover files from the services tier (AIX Linux) 388
Removing leftover files from the engine tier (AIX Linux) 389
Removing leftover files from the repository tier (AIX Linux) 390

Chapter 16. Removing IBM InfoSphere Information Server software manually . 393
Manually removing IBM InfoSphere Information Server (Windows) 393
Manually removing InfoSphere Information Server from the services tier (Windows) . . . 393
Manually removing InfoSphere Information Server from the engine tier (Windows) 395
Manually removing InfoSphere Information Server from the repository tier (Windows) . . . 396
Manually removing InfoSphere Information Server from the client tier (Windows) 398
Manually removing IBM InfoSphere Information Server (AIX Linux) 399
Manually removing InfoSphere Information Server from the services tier (AIX Linux) . . . 400
Manually removing InfoSphere Information Server from the engine tier (AIX Linux) . . . 401
Manually removing InfoSphere Information Server from the repository tier (AIX Linux) . . 402

Part 6. Troubleshooting InfoSphere Information Server software installation 405

Chapter 17. Log files 407
Installation log files 407
Application server system log files 407
InfoSphere Information Server console client log file 409

Chapter 18. General installation problems 411
ExecuteScriptAction failed: File Names in the following package are too long 411
Recovering from a UVREGEN installation exception 412
"User credentials are not valid" error appears unexpectedly during installation 413
IWAV0003E warning in installation log file . . . 413
Warnings occur during installation of DB2. . . . 414
Installation problems with VMware ESX 414
"The publisher could not be verified" message . . 415
"There is no script engine for file extension '.js'" message appears during installation 415
Installation process hangs when a web browser session is closed 416
Fix pack version errors during IBM InfoSphere Information Server installation 416

Chapter 19. Authentication problems 419
Confirming user privileges 419

Chapter 20. WebSphere Application Server Network Deployment server problems 421
Resolving IBM WebSphere Application Server installation problems 421
Installation fails when the IBM WebSphere Application Server Network Deployment port is changed after installation 422
Resolving connection problems 422
Unable to start IBM WebSphere Application Server after the user name changes or expires 423
Adding an IBM WebSphere Application Server administrator user to an internal user registry . . 425
Repairing the WebSphere Application Server Network Deployment registry after switching to LDAP 425
Services tier performance 426

Chapter 21. IBM InfoSphere Information Server reporting problems 429
Cannot view report result from InfoSphere Information Server console or InfoSphere Information Server Web console 429
Enabling and disabling trace logging of the reporting engine 429

Chapter 22. Client and console problems 433
Java Virtual Machine fails to load 433
Rich client login failure 433
Resolving problems logging in to the WebSphere Application Server administrative console 435
Connection problems after IBM WebSphere Application Server cluster members are removed . 437

Chapter 23. Removal problems 439
IBM InfoSphere Information Server installation fails repeatedly 439
Ignoring runtime error during uninstallation . . . 439
MKS Toolkit errors occur when removing clients 440

Part 7. Appendixes 443

Appendix A. Product accessibility . . 445

Appendix B. Reading command-line syntax 447

Appendix C. Contacting IBM 449

Appendix D. Accessing the product documentation 451

Notices and trademarks 453

Index 459

Part 1. Planning to install IBM InfoSphere Information Server

Before you install IBM® InfoSphere® Information Server, review the system requirements, learn about tiers and components, design your topology, and determine your installation steps.

© Copyright IBM Corp. 2005, 2014

Chapter 1. Tiers and components

You install IBM InfoSphere Information Server product modules in logical tiers. A *tier* is a logical group of components within InfoSphere Information Server and the computers on which those components are installed.

Each tier includes a subgroup of the components that make up InfoSphere Information Server. The tiers provide services, job execution, and storage of metadata and other data for your product modules.

Client tier

The client tier consists of the client programs and consoles that are used for development, administration, and other tasks and the computers where they are installed.

The following tools are installed as part of the client tier, based on the products and components that you select:

- IBM InfoSphere Information Server console
- IBM InfoSphere DataStage® and QualityStage® Administrator client
- IBM InfoSphere DataStage and QualityStage Designer client
- IBM InfoSphere DataStage and QualityStage Director client
- IBM InfoSphere FastTrack client
- Metadata interchange agent and InfoSphere Metadata Integration Bridges. The metadata interchange agent enables the use of bridges with InfoSphere Metadata Asset Manager.
- IBM InfoSphere Connector Migration Tool
- IBM InfoSphere Information Server istool command line. The istool framework is installed on the engine tier and client tier. Commands for IBM InfoSphere Information Analyzer, IBM InfoSphere Information Governance Catalog, and InfoSphere FastTrack are installed on the clients only when those products are installed.
- The Multi-Client Manager is installed when you install a product that includes InfoSphere DataStage and InfoSphere QualityStage client tier components. The Multi-Client Manager enables you to switch between multiple versions of InfoSphere DataStage clients. For example, you can switch between Version 8.5 and Version 7.5 clients.
- The MKS Toolkit is installed in the client tier. This toolset is used by the InfoSphere QualityStage migration utility.

The following diagram shows the client tier.

© Copyright IBM Corp. 2005, 2014

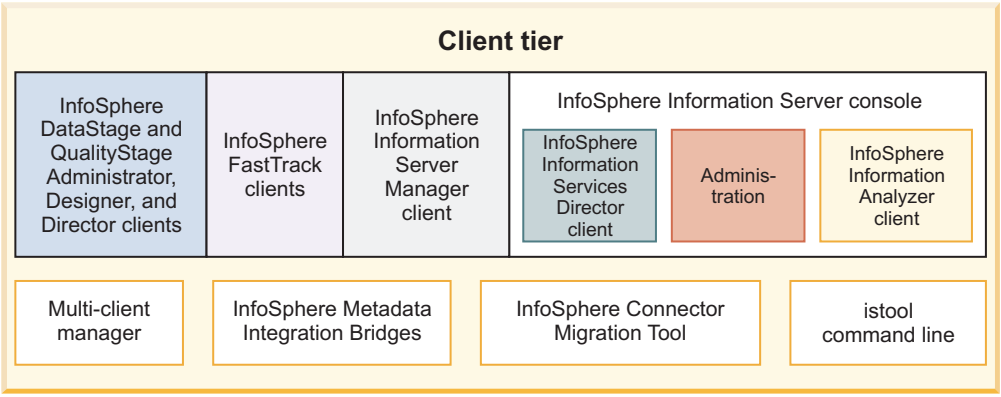


Figure 1. Client tier components

Services tier

The services tier consists of the application server, common services, and product services for the suite and product modules. The services tier provides common services (such as security) and services that are specific to certain product modules. On the services tier, IBM WebSphere® Application Server hosts the services. The services tier also hosts InfoSphere Information Server applications that are Web-based.

Some services are common to all product modules. Other services are specific to the product modules that you install. The services tier must have access to the metadata repository tier and the engine tier.

An application server hosts these services. IBM WebSphere Application Server is included with the suite for supported operating systems. You can choose to use WebSphere Application Server Liberty Profile or WebSphere Application Server Network Deployment. Alternatively, you can use an existing instance of WebSphere Application Server Network Deployment, if the version is supported by InfoSphere Information Server.

The following diagram shows the services that run on the application server on the services tier.

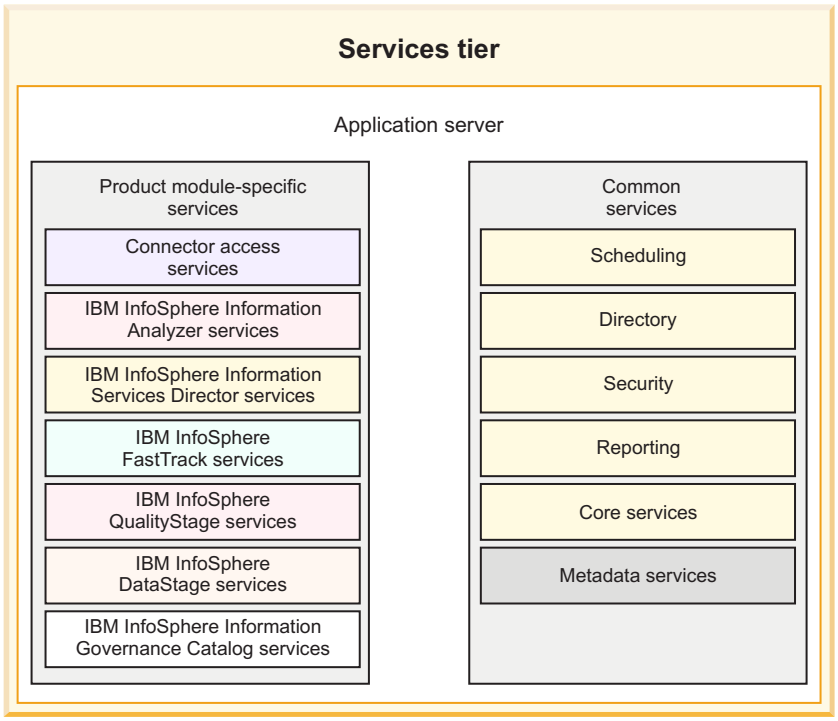


Figure 2. Services tier services

Product module-specific services for IBM InfoSphere Information Analyzer, IBM InfoSphere Information Services Director, IBM InfoSphere FastTrack, IBM InfoSphere DataStage, IBM InfoSphere QualityStage, and IBM InfoSphere Information Governance Catalog are included on the services tier. The services tier also includes connector access services that provide access to external data sources through the ODBC driver components and the connector access services agent on the engine tier.

The common services include:

Scheduling services
: These services plan and track activities such as logging, reporting, and suite component tasks such as data monitoring and trending. You can use the InfoSphere Information Server console and Web console to maintain the schedules. Within the consoles, you can define schedules, view their status, history, and forecast, and purge them from the system. For example, a report run and the analysis job within InfoSphere Information Analyzer are scheduled tasks.

Logging services
: These services enable you to manage logs across all the InfoSphere Information Server suite components. You can view the logs and resolve problems by using the InfoSphere Information Server console and Web console. Logs are stored in the metadata repository. Each InfoSphere Information Server suite component defines relevant logging categories.

Directory services
: These services act as a central authority that can authenticate resources and manage identities and relationships among identities. You can base directories on the InfoSphere Information Server internal user registry. Alternatively, you can use external user registries such as the local

operating system user registry, or Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory registries.

Security services
These services manage role-based authorization of users, access-control services, and encryption that complies with many privacy and security regulations. If the user registry internal to InfoSphere Information Server is used, administrators can use the InfoSphere Information Server console and Web console to add users, groups, and roles within InfoSphere Information Server.

Reporting services
These services manage runtime and administrative aspects of reporting for InfoSphere Information Server. You can create product module-specific reports for InfoSphere DataStage, InfoSphere QualityStage, and InfoSphere Information Analyzer. You can also create cross-product reports for logging, monitoring, scheduling, and security services. You can access, delete, and purge report results from an associated scheduled report execution. You can set up and run all reporting tasks from the InfoSphere Information Server Web console.

Core services
These services are low-level services such as service registration, life cycle management, binding services, and agent services.

Metadata services
These services implement the integrated metadata management within InfoSphere Information Server. Functions include repository management, persistence management, and model management.

The following InfoSphere Information Server Web-based applications are installed as part of the services tier.

- Information Governance Catalog
- The IBM InfoSphere Information Server Web console. A browser shortcut to the IBM InfoSphere Information Server Web console is created during the InfoSphere Information Server installation. The Web console consists of administration and reporting tools.
- IBM InfoSphere Information Server Manager client

Engine tier

The engine tier is the logical group of engine components, communication agents, and so on. The engine runs jobs and other tasks for product modules.

Several product modules require the engine tier for certain operations. You install the engine tier components as part of the installation process for these product modules. The following product modules require the engine tier:

- IBM InfoSphere DataStage
- IBM InfoSphere Information Analyzer
- IBM InfoSphere Information Services Director
- IBM InfoSphere QualityStage
- IBM InfoSphere Information Server istool command line. The istool framework is installed on the engine tier and client tier. Commands for InfoSphere Information Analyzer are installed on the engine tier only when the product is installed.

IBM InfoSphere FastTrack, Information Governance Catalog, and IBM InfoSphere Glossary Anywhere do not require an engine tier.

AIX® Linux: The following configurations are supported:

- Multiple engines, each on a different computer, all registered to the same InfoSphere Information Server services tier.
- Multiple engines on the same computer. In this configuration, each engine must be registered to a different services tier. This configuration is called an *ITAG* installation.

Microsoft Windows: Only one InfoSphere Information Server engine can be installed on a single computer.

The installation program installs the following engine components as part of an engine tier:

InfoSphere Information Server engine
Runs tasks or jobs such as discovery, analysis, cleansing, or transformation. The engine includes the server engine and parallel engine and other components that make up the runtime environment for InfoSphere Information Server and its product components.

ASB agents
Java™ processes that run in the background on each computer that hosts an InfoSphere Information Server engine tier. When a service that runs on the services tier receives a service request that requires processing by an engine tier component, the agents receive and convey the request.

AIX Linux: The agents run as daemons that are named ASBAgent.

Microsoft Windows: The agents run as services that are named ASBAgent.

ASB agents include:

Connector access services agent
Conveys service requests between the ODBC driver components on the engine tier and the connector access services component on the services tier.

InfoSphere Information Analyzer agent
Conveys service requests between the engine components on the engine tier and the InfoSphere Information Analyzer services component on the services tier.

InfoSphere Information Services Director agent
Conveys service requests between the engine components on the engine tier and the InfoSphere Information Services Director services component on the services tier.

ODBC drivers

The installation program installs a set of ODBC drivers on the engine tier that works with InfoSphere Information Server components. These drivers provide connectivity to source and target data.

Resource Tracker

The installation program installs the Resource Tracker for parallel jobs with the engine components for InfoSphere DataStage and InfoSphere QualityStage. The Resource Tracker logs the processor, memory, and I/O usage on each computer that runs parallel jobs.

dsrpcd (DSRPC Service)
Allows InfoSphere DataStage clients to connect to the server engine.

AIX Linux: This process runs as a daemon (dsrpcd).

Microsoft Windows: This process runs as the DSRPC Service.

Job monitor
A Java application (JobMonApp) that collects processing information from parallel engine jobs. The information is routed to the server controller process for the parallel engine job. The server controller process updates various files in the metadata repository with statistics such as the number of inputs and outputs, the external resources that are accessed, operator start time, and the number of rows processed.

Operational MetaData monitor
A Java application (OMDMonApp) that processes the operational metadata XML files that are generated by job runs if the collection of operational metadata is enabled. The information in the XML files is stored in the metadata repository, and the XML files are deleted.

DataStage engine resource service
Microsoft Windows: Establishes the shared memory structure that is used by server engine processes.

DataStage Telnet service
Microsoft Windows: Allows users to connect to the server engine by using Telnet. Useful for debugging issues with the server engine. Does not need to be started for normal InfoSphere DataStage processing.

MKS Toolkit
Microsoft Windows: Used by the InfoSphere Information Server parallel engine to run jobs.

The following diagram shows the components that make up the engine tier. Items marked with asterisks (*) are only present in Microsoft Windows installations.

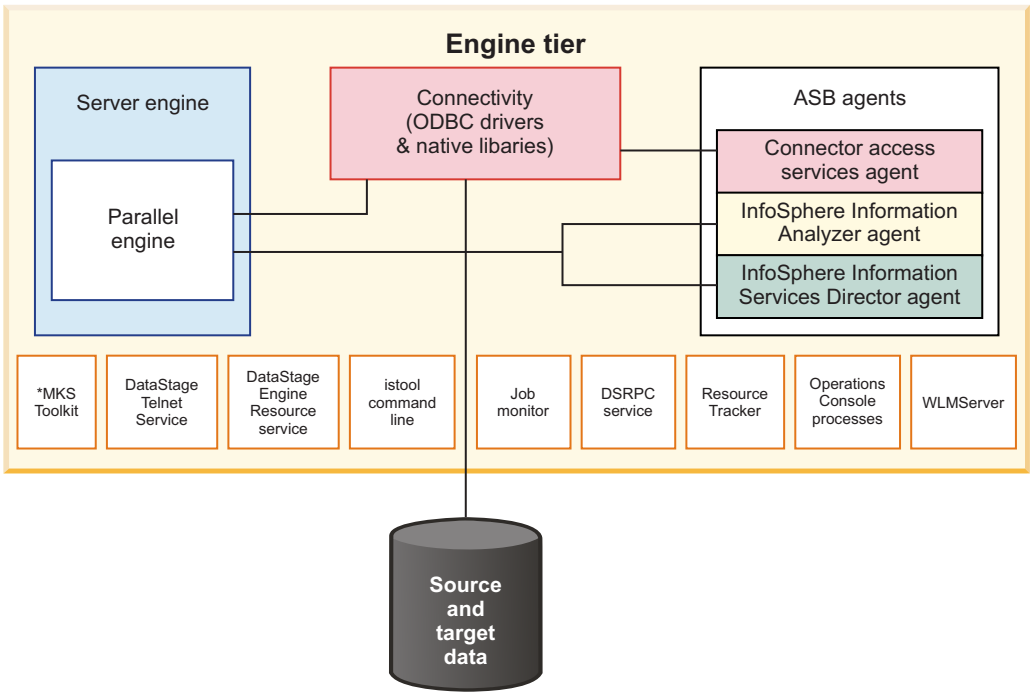


Figure 3. Engine tier components

Note: InfoSphere Metadata Integration Bridges are installed only on the client tier, not on the engine tier.

Repository tier

The repository tier consists of the metadata repository and, if installed, other data stores to support other product modules. The metadata repository contains the shared metadata, data, and configuration information for InfoSphere Information Server product modules. The other data stores store extended data for use by the product modules they support, such as the operations database, which is a data store that is used by the engine operations console.

The repository tier includes the metadata repository database for IBM InfoSphere Information Server. The metadata repository exists as its own schema in this database. The metadata repository is a shared component that stores design-time, runtime, glossary, and other metadata for product modules in the InfoSphere Information Server suite.

The repository tier also includes other data stores. Some of these data stores might be referred to as databases or repositories throughout the documentation, based on legacy naming conventions. However, they might exist as either separate database schemas in a shared database or as schemas in their own separate databases in the product suite. Some of these data stores can exist on other computers, and in that sense the repository tier can be thought of as a logical tier. However, when this documentation refers to the repository tier computer, it is the computer that hosts the database for the metadata repository. Location and connection information for the other data stores in the repository tier is stored in the metadata repository.

The repository tier can include these data stores:

- As part of IBM InfoSphere Metadata Asset Manager, a data store called the *staging area* is installed as a separate schema in the database for the metadata repository.
- If InfoSphere Information Analyzer is installed, the repository tier also includes one or more analysis databases, one per project, for example, which are installed as distinct databases outside of the metadata repository database. The analysis databases are used by InfoSphere Information Analyzer when it runs analysis jobs.
- An operations database can be installed with IBM InfoSphere DataStage and QualityStage as a separate schema in the metadata repository database or as a separate schema in another database. Additional operations databases can be created, one per InfoSphere Information Server engine, if desired.
- A data store called the Standardization Rules Designer database is installed with Standardization Rules Designer. By default, it is installed as a separate schema in the metadata repository database. However, you can choose to install it in a separate schema in another existing database.
- A data store called the exceptions database is installed with IBM InfoSphere Data Quality Console. By default, it is installed as a separate schema in the metadata repository database. However, you can choose to install it in a separate schema in another existing database.

The services tier must have access to the metadata repository. When product modules store or retrieve metadata, services on the services tier connect to the metadata repository and manage the access to the databases from the product modules.

The engine tier and the client tier must have direct access to the analysis databases and operations databases, which are part of the repository tier.

The following diagram shows some of the components that make up the repository tier.

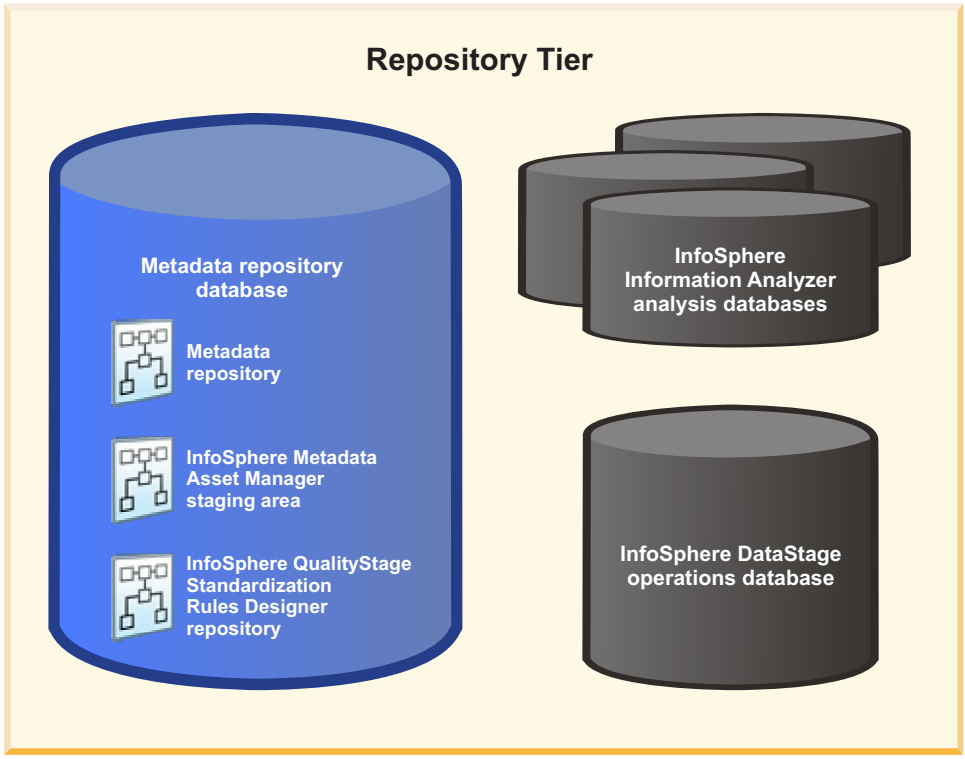


Figure 4. Repository tier components

Tier relationships

The tiers provide services, job execution, and storage of metadata and other data for the product modules that you install.

The following diagram illustrates the tier relationships.

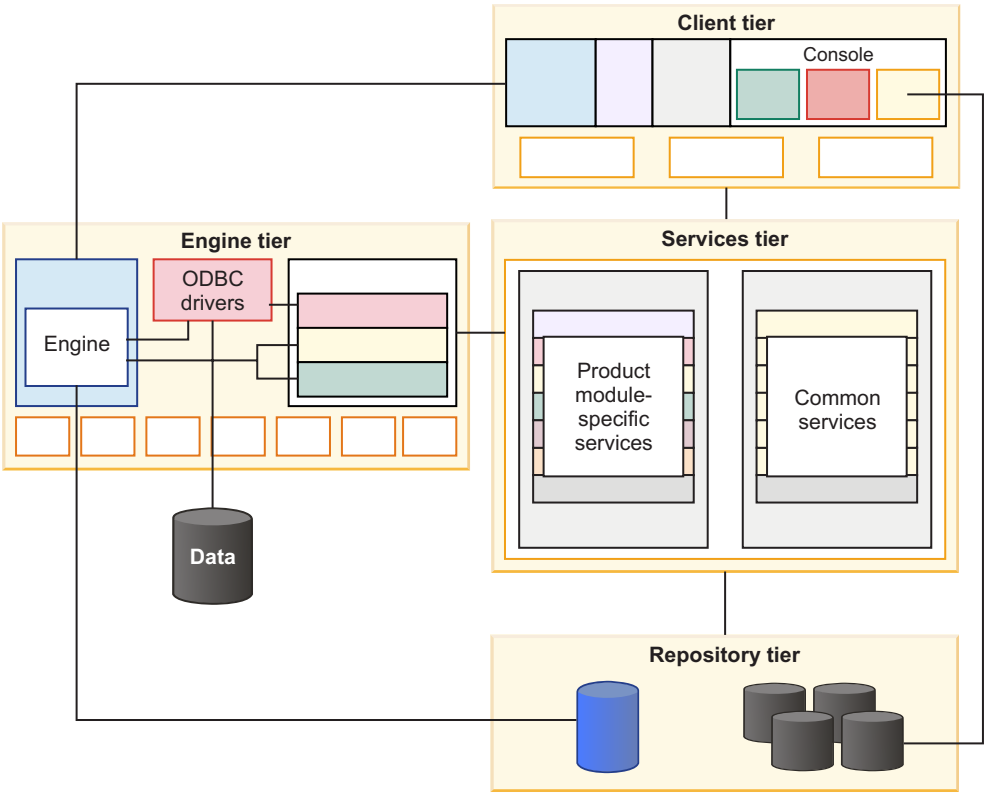


Figure 5. Tier relationships

The tiers relate to one another in the following ways:

- Relationships differ depending on which product modules you install.
- Client programs on the client tier communicate primarily with the services tier. The IBM InfoSphere DataStage and QualityStage clients also communicate with the engine tier.
- Various services within the services tier communicate with agents on the engine tier.
- Metadata services on the services tier communicate with the metadata repository.
- ODBC drivers on the engine tier communicate with external databases.
- InfoSphere Metadata Integration Bridges on the client tier can import data from external sources. Some bridges can also export data.
- With the IBM InfoSphere Information Analyzer product module, the engine tier communicates directly with the analysis databases on the repository tier. The InfoSphere Information Analyzer client also communicates directly with the analysis databases.

Chapter 2. Planning to install IBM InfoSphere Information Server

Before you install IBM InfoSphere Information Server, review the system requirements, learn about tiers and components, design your topology, and determine your installation steps.

Understanding your installation options

IBM InfoSphere Information Server can be installed in a variety of ways depending on your needs. It can be set up for different topologies, application server configurations, metadata repository database types, and high availability scenarios. Review this section to better understand the options that you have available.

Options for installing the database software

You must determine how you will install the database system for each database, and how you will create the databases within the database system.

The procedure for setting up IBM InfoSphere Information Server databases differs depending upon the database system that you use and the location of the database system.

To use a clustered database configuration for high availability, you must first install the database system and configure the cluster. Then, use the scripts on the InfoSphere Information Server installation media to create the databases within the instance. When you run the InfoSphere Information Server installation program, specify the database system for the installation. See "Pre-installing the metadata repository database" on page 114 for more information about installing databases.

IBM DB2

There are two major options to configure the DB2 10.5 Enterprise Server Edition database:

Installing the DB2® database system by using the InfoSphere Information Server installation program

Use the InfoSphere Information Server installation program to install DB2 10.5 Enterprise Server Edition on the target computer and create the database within it. DB2 is bundled with InfoSphere Information Server. Within the installation program, select the **Repository tier** installation option to install DB2 and create the databases.

Preinstalling the DB2 database system

Before you run the installation program, preinstall the DB2 database system (any supported version) on the target computer or another computer. Then, create the databases by using one of the following methods:

- If the DB2 instance is located on the target computer, use the InfoSphere Information Server installation program to create the databases within the instance. Within the installation program, select the **Repository tier** installation option to create the databases.
- If the DB2 instance is not located on the target computer, create the databases within the instance before running the InfoSphere Information

© Copyright IBM Corp. 2005, 2014

Server installation program. To create the databases, use the scripts that are provided with the InfoSphere Information Server software.

The following table summarizes the different installation options for IBM DB2 databases:

Table 1. Summary of database system installation and database creation options for IBM DB2 databases

Configuration	Database system installation options	Database creation options
DB2, non-clustered configuration	You can preinstall DB2 before you run the InfoSphere Information Server installation program. If you don't preinstall DB2, you can use the InfoSphere Information Server installation program to install it.	Use either of the following methods: • Use the InfoSphere Information Server installation program. • Use the scripts on the InfoSphere Information Server installation media.
DB2, clustered configuration	Preinstall, and then create the cluster.	Use the scripts on the InfoSphere Information Server installation media.
DB2, high availability disaster recovery (HADR) configuration (a supported version earlier than what is bundled with InfoSphere Information Server).	Preinstall, and then create the HADR configuration.	Use either of the following methods to create the database on the primary node: • Use the InfoSphere Information Server installation program. • Use the scripts on the InfoSphere Information Server installation media.
DB2, high availability disaster recovery (HADR) configuration (the same version that is bundled with InfoSphere Information Server).	Use either of the following methods: • Use InfoSphere Information Server installation program to install • Preinstall before you run the InfoSphere Information Server installation program	Use either of the following methods to create the database on the primary node: • Use the InfoSphere Information Server installation program. • Use the scripts on the InfoSphere Information Server installation media.

Options for installing the application server software

The central component within the services tier is IBM WebSphere Application Server. This application server hosts the common and product-specific services for IBM InfoSphere Information Server.

Options

You have four options available on how to set up WebSphere Application Server.

Install WebSphere Application Server Network Deployment with the InfoSphere Information Server program

As part of the services tier installation, the installation program creates a server profile that is named `InfoSphere` and an application server instance that is named `server1`. This application server is licensed for use only with InfoSphere Information Server.

The advantage of this option is that the installation program does much of the configuration for you and you get a fully functional application server with wider authentication and security options to handle more complex configurations.

Install WebSphere Application Server Liberty Profile with the InfoSphere Information Server program

As part of the services tier installation, the installation program installs WebSphere Application Server Liberty Profile under the InfoSphere Information Server installation directory in a folder named `wlp`. This application server is licensed for use only with InfoSphere Information Server.

WebSphere Application Server Liberty Profile offers the following benefits:

- The application server uses less memory and disk space.
- The installation is considerably faster.
- There is little post-installation configuration.
- Specialized administration expertise is not required.

Be mindful of the following limitations with using WebSphere Application Server Liberty Profile that are not present with WebSphere Application Server Network Deployment:

- You use the application server that is installed with InfoSphere Information Server. You cannot use your own preinstalled version of WebSphere Application Server Liberty Profile.
- Updates and patches to the application server are offered only through an InfoSphere Information Server update or patch when it is available. Patches to the application server cannot be applied separately.
- Application server clustering is not supported. (You can, however, use WebSphere Application Server Liberty Profile with a clustered database.)
- LDAP authentication is supported, but LDAP attribute mapping is not. (That is, user and group attributes other than the user and group ID, such as name or email address that might be in the LDAP registry, will not made available to populate various interface screens.)
- Local operating system user registry is not supported.
- PAM user registry is not supported.
- Windows desktop SSO for IBM InfoSphere Information Governance Catalog is not supported.

In general, WebSphere Application Server Liberty Profile is a good choice for smaller businesses and organizations that do not require the complex configuration options that come with WebSphere Application Server Network Deployment.

Preinstall or use an existing version of WebSphere Application Server Network Deployment in a stand-alone configuration

If you use an existing installation, you must create the profile manually. You must create an empty and unsecure profile that will be solely

dedicated to InfoSphere Information Server. The InfoSphere Information Server installation program will secure the profile and do the necessary configuration.

The application server that is bundled with InfoSphere Information Server is licensed for use only with InfoSphere Information Server. If you use your own preinstalled separately licensed installation of WebSphere Application Server Network Deployment, you can run other applications in separate profiles on the same computer as InfoSphere Information Server.

Preinstall or use an existing version of WebSphere Application Server Network Deployment in a clustered configuration

If you want to set up a clustered WebSphere Application Server configuration for high availability and scalability, you must preinstall or use an existing version of WebSphere Application Server Network Deployment. You configure the cluster, profiles, and optionally the front-end dispatcher before you install InfoSphere Information Server. You can configure the front-end web server after installation. InfoSphere Information Server requires a dedicated cluster (not shared with any other applications). The cell that contains the dedicated cluster can be shared, however.

For more information, see "Setting up a preinstalled application server" on page 148.

Version support

WebSphere Application Server Network Deployment (64-bit) 8.5.5.1 and future fix packs are the only supported versions of WebSphere Application Server Network Deployment, whether the application server is pre-installed or installed by the installation program.

WebSphere Application Server Liberty Profile 8.5.5.1 is the version that can be installed by the InfoSphere Information Server installation program, if this is the application server that you choose to use.

IBM InfoSphere Information Services Director binding support

Table 2. InfoSphere Information Services Director bindings. List of the InfoSphere Information Services Director bindings that are supported with each WebSphere Application Server type in this version of InfoSphere Information Server

Binding	Network Deployment	Liberty profile	Not supported in Version 11.3
REST			X
REST2	✓	✓	
SOAP over HTTP	✓	✓	
Text over HTTP	✓	✓	
EJB	✓		
Text over JMS	✓		
SOAP over JMS			X

Choosing a topology

A successful deployment of IBM InfoSphere Information Server must include a topology that meets the needs of your organization, such as performance, ease of maintenance, security, high availability, and scalability. Determining the optimal deployment requires a firm understanding of your requirements and the topology that you plan to implement.

This information provides explanations, illustrations, scenarios, and checklists to help you design an ideal topology based on a set of available resources and a set of functional requirements (such as high availability and scalability). Each of these variables represents different dimensions of a topology, and a change in any of these dimensions can greatly affect the resulting topology. Therefore, identifying and quantifying these dimensions is important to remain within your constraints while still meeting requirements.

Designing an installation topology

The overall topology of an installation can range widely in complexity. Optimize the design of the topology to accommodate the product modules that you install and any requirements for high availability, performance and throughput, and user concurrency.

IBM InfoSphere Information Server supports a variety of topolgies. In these topology descriptions, the term *computers* refers to separate physical servers, logical partitions (LPARs), or virtual machines.

Basic installation topologies:

If you do not need a high availability solution and do not anticipate scaling the installation for higher capacity in the future, choose a basic topology.

Single computer topology:

In this topology, the engine tier, repository tier, services tier, and client tier are all installed on a single computer. This topology is supported on Microsoft Windows only.

The computer must be a server-class computer that meets all system requirements, such as memory, disk space, and other requirements. To support the client tier, the computer must run Microsoft Windows and meet all system requirements for this platform.

This topology is suitable for demonstration systems and for small-scale development.

The following diagram illustrates this topology.

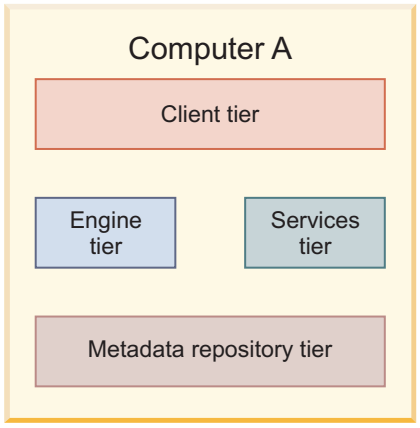


Figure 6. Single computer topology

Client/server topology:

In this topology, the engine tier, repository tier, and services tier are all installed on a single computer. The client tier is installed on separate computers.

The client tier computers must run Microsoft Windows. The computer that hosts the other tiers can run any operating system that IBM InfoSphere Information Server supports.

This topology centralizes administration and isolates it from client users. The following diagram illustrates this topology. Although the diagram shows only one client tier computer, you can include multiple client tier computers.

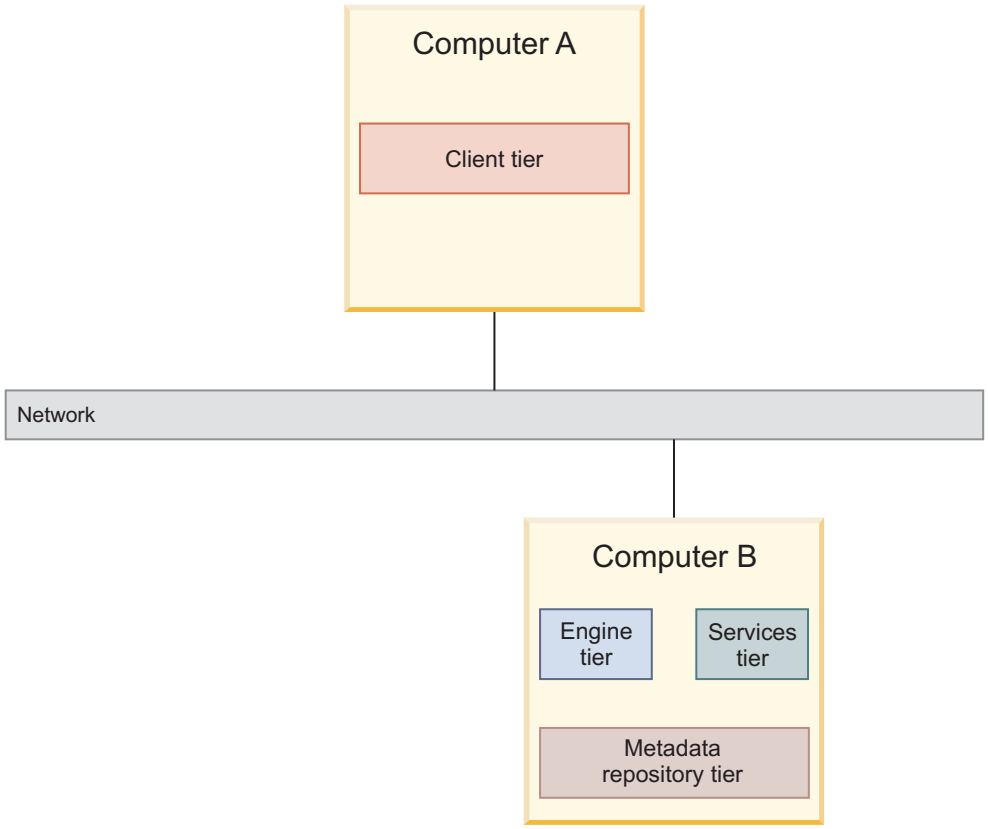


Figure 7. Client/server topology

Dedicated engine tier topology:

In this topology, the services tier and repository tier are installed on one computer. The engine tier is installed on another computer. The client tier is installed on separate computers.

The client tier computer must run Microsoft Windows. The other computers can run any operating system that IBM InfoSphere Information Server supports.

Installing the repository tier with the services tier provides optimal performance because there is no network latency between the tiers. Also, higher engine tier activity does not affect the operations of the services tier and repository tier. This topology is suitable for environments that have high IBM InfoSphere DataStage and QualityStage job requirements and for IBM InfoSphere Information Analyzer installations.

The following diagram illustrates this topology. Although the diagram shows only one client tier computer, you can include multiple client tier computers.

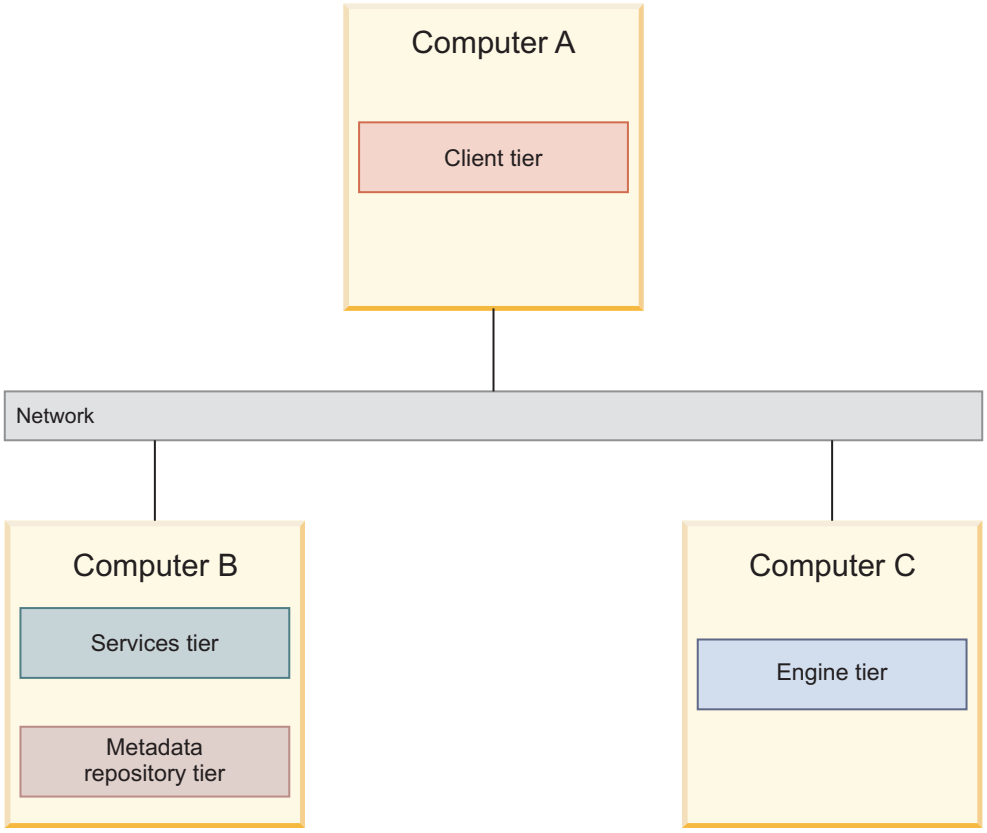


Figure 8. Dedicated engine tier topology

Dedicated computer for each tier topology:

You can host each tier on a separate computer.

This topology provides each tier with dedicated computing resources. If you choose this topology, minimize network latency between all tiers. In particular, you must have a high-bandwidth connection between the services tier and the metadata repository tier.

The services tier and the metadata repository tier must be on the same subnet.

The following diagram illustrates this topology. Although the diagram shows only one client tier computer, you can include multiple client tier computers.

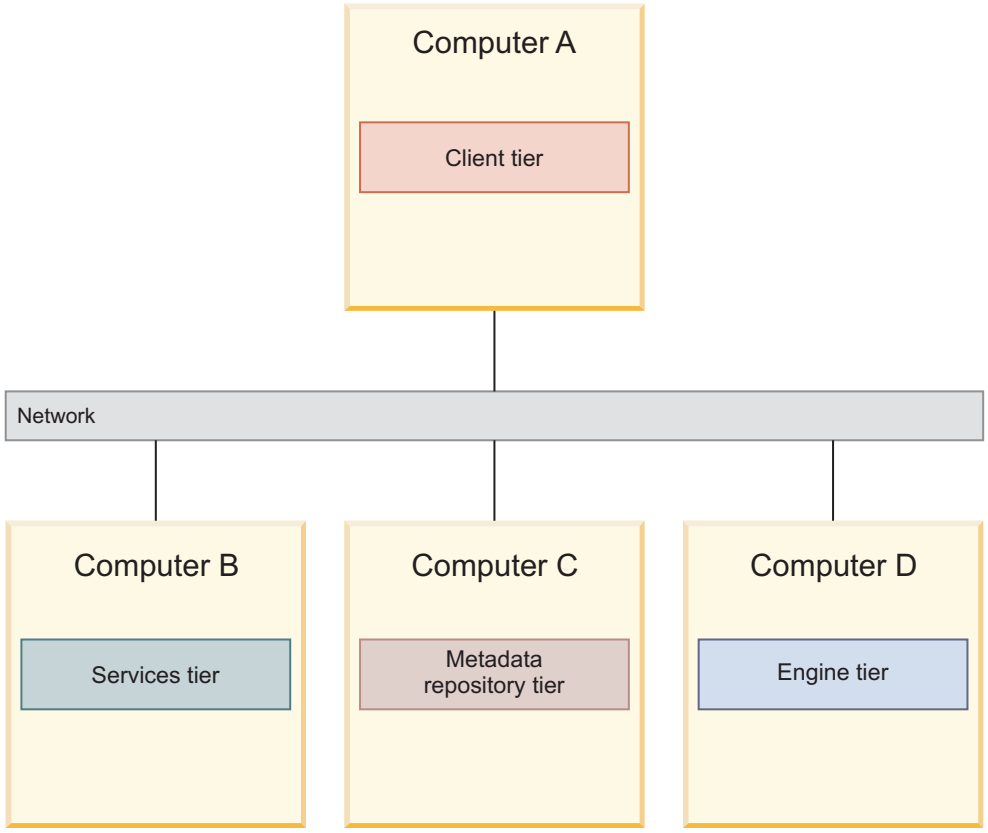


Figure 9. Topology with a dedicated computer for each tier

Topology with separate instances of IBM InfoSphere Information Server:

AIX Linux: You can install more than one instance of InfoSphere Information Server on the same set of computers. **Microsoft Windows**: You cannot install multiple versions of services and engine tiers, but you can install multiple clients and use the Multi-Client Manager to switch server connections.

The instances are isolated from one another, with separate engine tiers, services tiers, and repository tiers. These tiers cannot be shared between instances. The instances can exist on the same computer in different installation directories or on different computers. The instances can be the same software version or different versions.

If the instances are the same InfoSphere Information Server version, you can configure a single client tier instance to communicate with multiple InfoSphere Information Server instances. The following diagram shows a topology with two separate InfoSphere Information Server instances that are installed on two separate computers. A single client tier installation communicates with both instances.

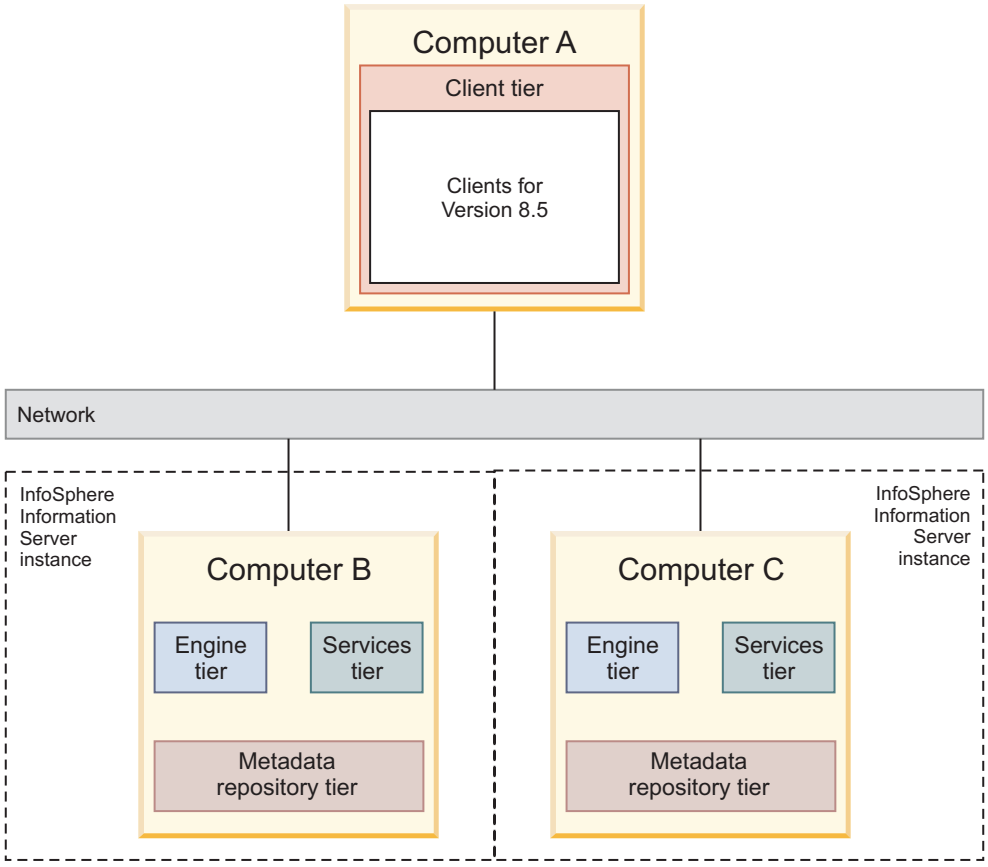


Figure 10. Topology with separate instances on separate computers

If the instances are different versions, the client tier must contain clients for each version. The following diagram shows a topology with two separate InfoSphere Information Server instances that are installed on two separate computers. The instances are different versions. The client tier contains clients for both versions. The Multi-Client Manager is used to switch between one version of IBM InfoSphere DataStage to another version that is also installed on the client tier.

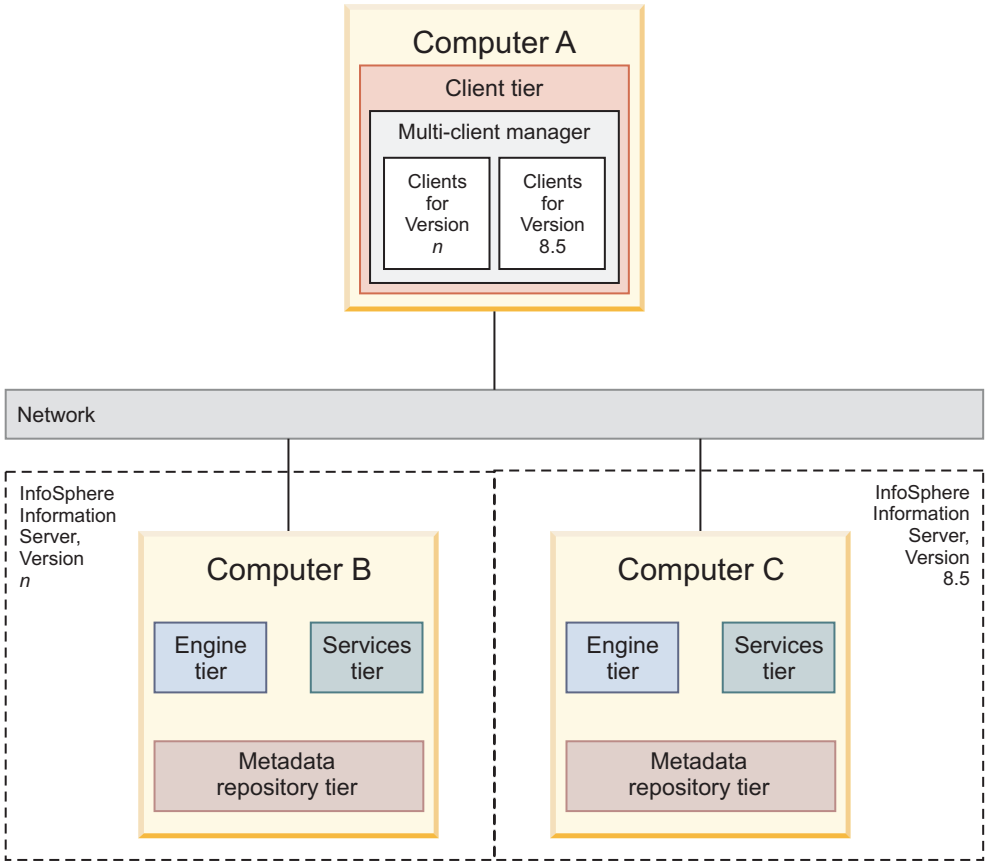


Figure 11. Topology with separate different version instances on separate computers

If the services tiers for more than one instance exist on the same computer, specify different ports for each services tier. The following diagram shows a topology where the services tiers and repository tiers for two different InfoSphere Information Server version instances are installed on the same computer. Each services tier communicates with a separate engine tier installed on a separate computer.

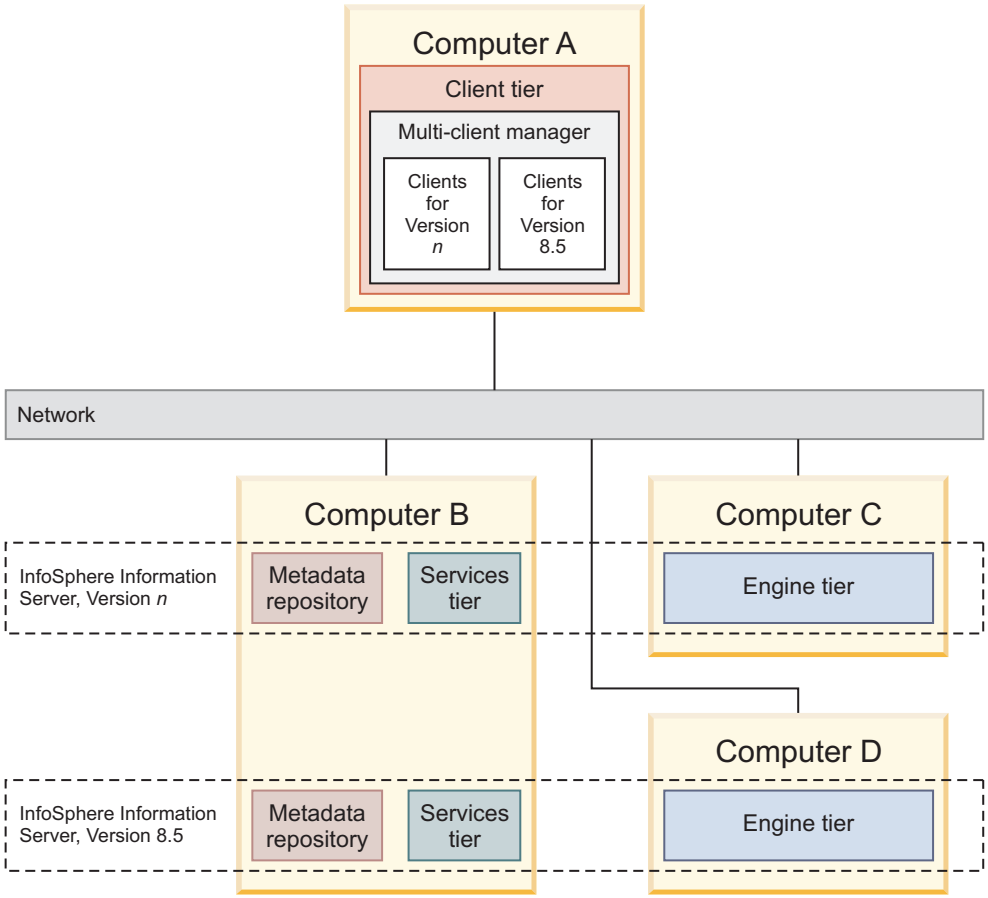


Figure 12. Topology where services tiers share a common computer

To install more than one instance of the InfoSphere Information Server engine on the same computer, create an *installation tag* (ITAG). An ITAG must be exactly three hexidecimal digits. The ITAG value is used when assigning shared memory segments so that each installation creates its shared memory in different locations. To perform this installation, run the installation program multiple times to create multiple engine tier instances. During the installation, assign each instance a unique identification tag and port number. All projects and jobs associated with a server engine instance are assigned the tag for the server engine instance.

Each engine tier instance must be registered to a different InfoSphere Information Server service instance.

The following diagram shows a topology where the engine tiers for two InfoSphere Information Server instances are installed on the same computer. Each engine tier has an associated ITAG. (ITAGs are three-character identifiers that are used to distinguish each engine instance on the computer.)

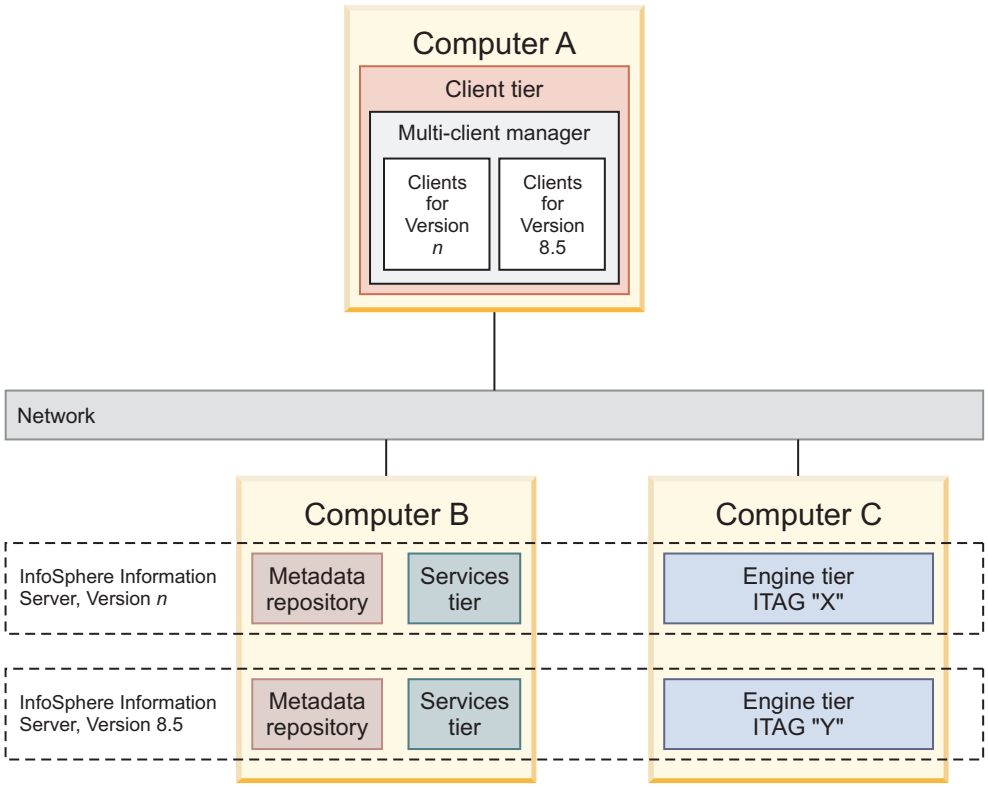


Figure 13. ITAG installation

Note: ITAG installations are not supported on Microsoft Windows computers.

Topology with multiple engine tiers:

You can create an installation topology that includes more than one engine tier. In this topology, each engine tier includes the IBM InfoSphere Information Server engine, which includes the server engine, the parallel engine, and other components that make up the runtime environment for InfoSphere Information Server and its product components.

You might choose this topology for project and resource management reasons. For example, you might dedicate one engine tier to IBM InfoSphere DataStage jobs and another engine tier to IBM InfoSphere Information Analyzer jobs. Or you might dedicate multiple engine tiers to separate sets of InfoSphere Information Analyzer jobs, each with its own analysis database.

Each engine tier instance must exist on a separate computer. You cannot configure the services tier to communicate with multiple engine tiers that are installed on the same computer.

The following diagram shows a topology with two engine tiers that are installed on separate computers. A single services tier instance communicates with both engine tiers.

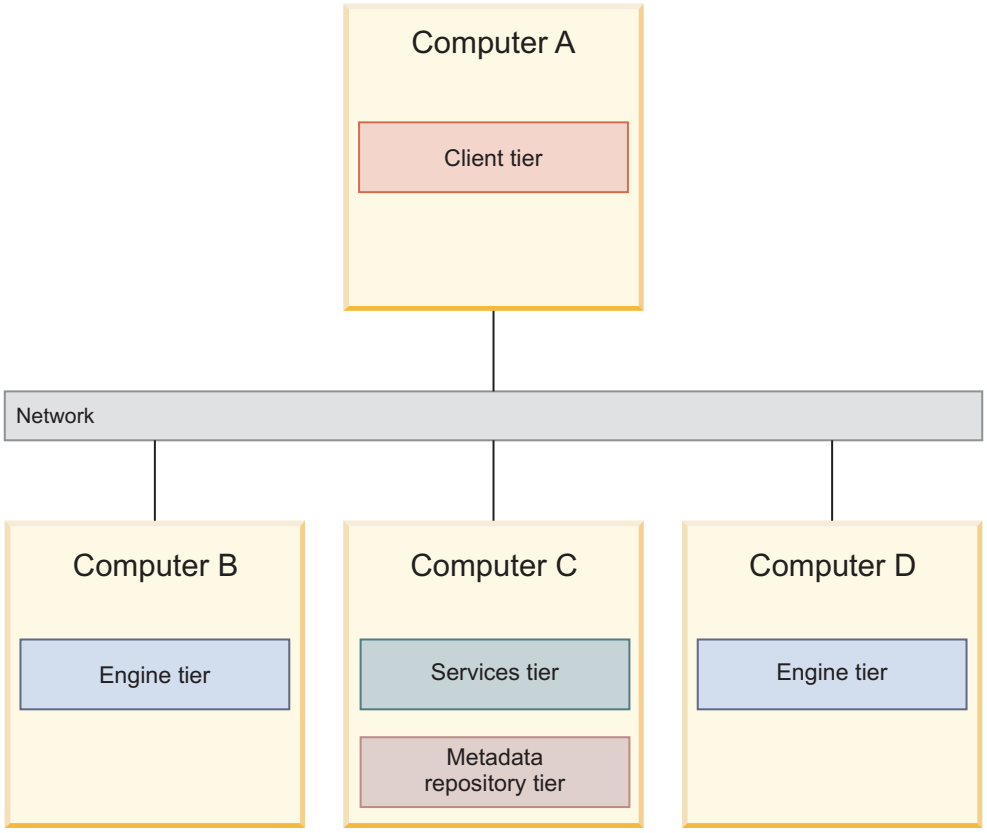


Figure 14. Topology with two engine tiers

Database topologies:

You use the metadata repository to store imported metadata, project configurations, reports, and results for all components of IBM InfoSphere Information Server. The metadata repository is included as part of the metadata repository tier, which can include additional product module data stores as separate databases or database schemas. The metadata repository is also used to keep a registry of these data stores, to include their locations and connection information. The staging area is created along with the metadata repository.

Depending on which product modules you install, the repository tier can include the following data stores:

- one or more InfoSphere Information Analyzer analysis databases
- the InfoSphere QualityStage Match Designer database
- the Standardization Rules Designer database
- an operations database to support the IBM InfoSphere DataStage and QualityStage Operations Console
- IBM InfoSphere Data Quality Console Exceptions database

For most of the databases, the installation program will register them in the metadata repository and create and configure them as necessary. You can also choose to manually register and set them up with the help of the **RepositoryAdmin** tool as a post-installation step. (See "Managing repositories" and "Changing configuration in repositories" in the *IBM InfoSphere Information Server Administration Guide*.)

Note: The term *data stores* is new for this release to replace *repositories* to prevent confusion between them and the *metadata repository*. However, the name of the tool to manage them is currently still called **RepositoryAdmin**, and the topics that describe the use of the tool still refers to data stores as repositories in certain contexts.

Staging area

Stores metadata that is imported from external data sources so that it can be examined before it is moved to the active metadata repository. Managed from InfoSphere Metadata Asset Manager.

Analysis databases

Analysis databases store high-volume, detailed analysis results, such as column analysis, primary key analysis, and domain analysis. InfoSphere Information Analyzer projects can share an analysis database, or you can associate each project with a specific analysis database.

If you install InfoSphere Information Analyzer, you must provide the location for one or more analysis databases. After the installation, you can add additional databases by using the InfoSphere Information Server console.

The analysis database might be used by a single InfoSphere Information Analyzer project, or it might be shared by multiple projects. For example, two InfoSphere Information Analyzer projects might use two different analysis databases, or they might share the same analysis database.

Match Designer database

The InfoSphere QualityStage Match Designer is a component of InfoSphere QualityStage that is used to design and test match specifications. Match specifications consist of match passes that identify duplicate entities within one or more files.

The InfoSphere Information Server installation program does not create the Match Designer results database. You can create the database before or after the installation, as long as the database is configured and accessible when you use the Match Designer. You can create the database on a computer where the client or engine tier is installed or on any computer that is accessible to both of these tiers. You must configure the database to receive the type of data that is processed in the Match Designer. For example, you must configure the database to receive double-byte data if the Match Designer processes Asian data.

Standardization Rules Designer database

The Standardization Rules Designer database is a component of InfoSphere QualityStage that is used to enhance standardization rule sets. After you enhance rule sets in the Standardization Rules Designer, you can apply the enhanced rule sets in a Standardize stage.

By default, the InfoSphere Information Server installation program creates the Standardization Rules Designer database as a separate schema in the metadata repository database.

Exceptions database

In InfoSphere Information Server products and components, entities that might require additional information or investigation are called exceptions. Information about each set of exceptions is provided by exception descriptors that are stored in the exceptions database. The exceptions database is managed from the IBM InfoSphere Data Quality Console.

These data stores can reside in the same database system installation as distinct databases or database schemas. Alternatively, other than the staging area, which must be located in the same database (but in a different schema) as the active metadata repository, you can locate the data stores and the metadata repository on different computers.

If you are creating a database system installation for a database, determine where to locate the instance and the database. The database must be accessible by the computers where the services tier and engine tiers are installed.

The following two diagrams illustrate a three-computer topology where the data stores are located on the same computer as the metadata repository database. They can be created in separate databases or in separate schemas in the same or different databases. Figure 16 on page 29 illustrates data stores as separate schemas within the metadata repository database. This is possible for all data stores except the analysis databases, which must be in their own database.

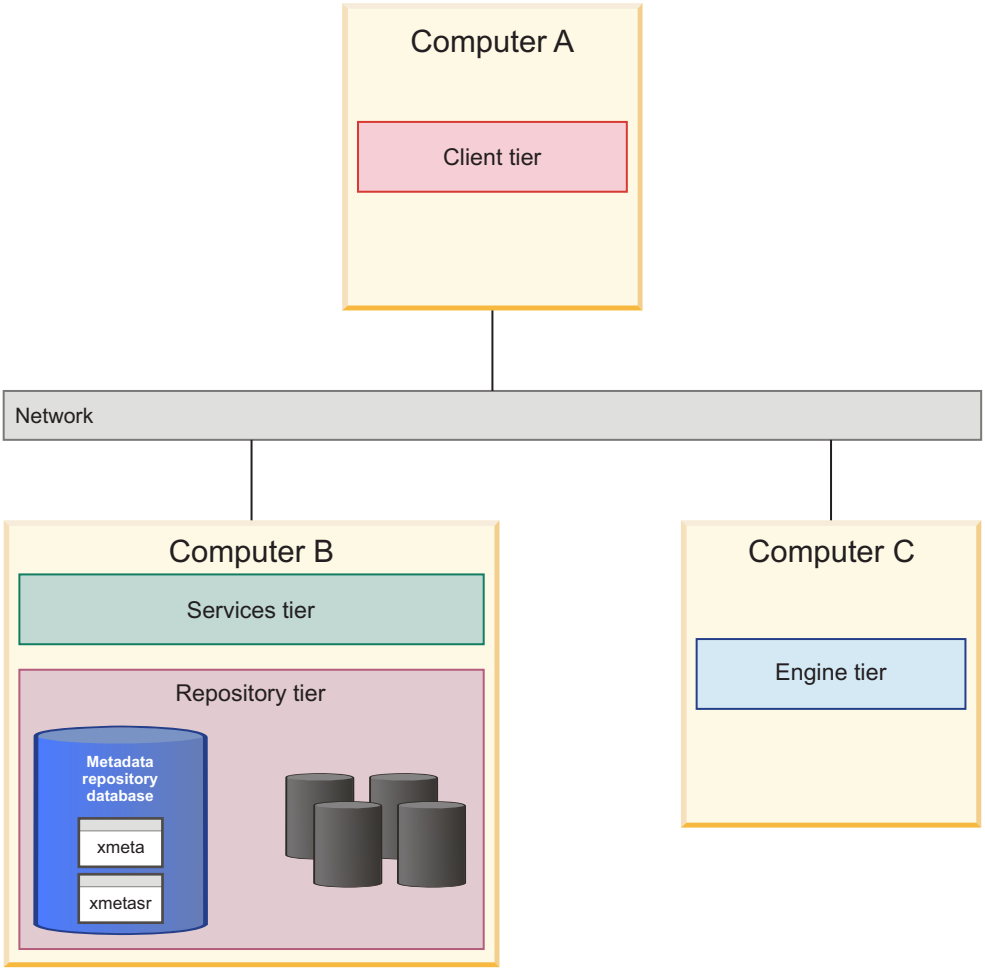


Figure 15. Topology with data stores and metadata repository database on the same database server installation

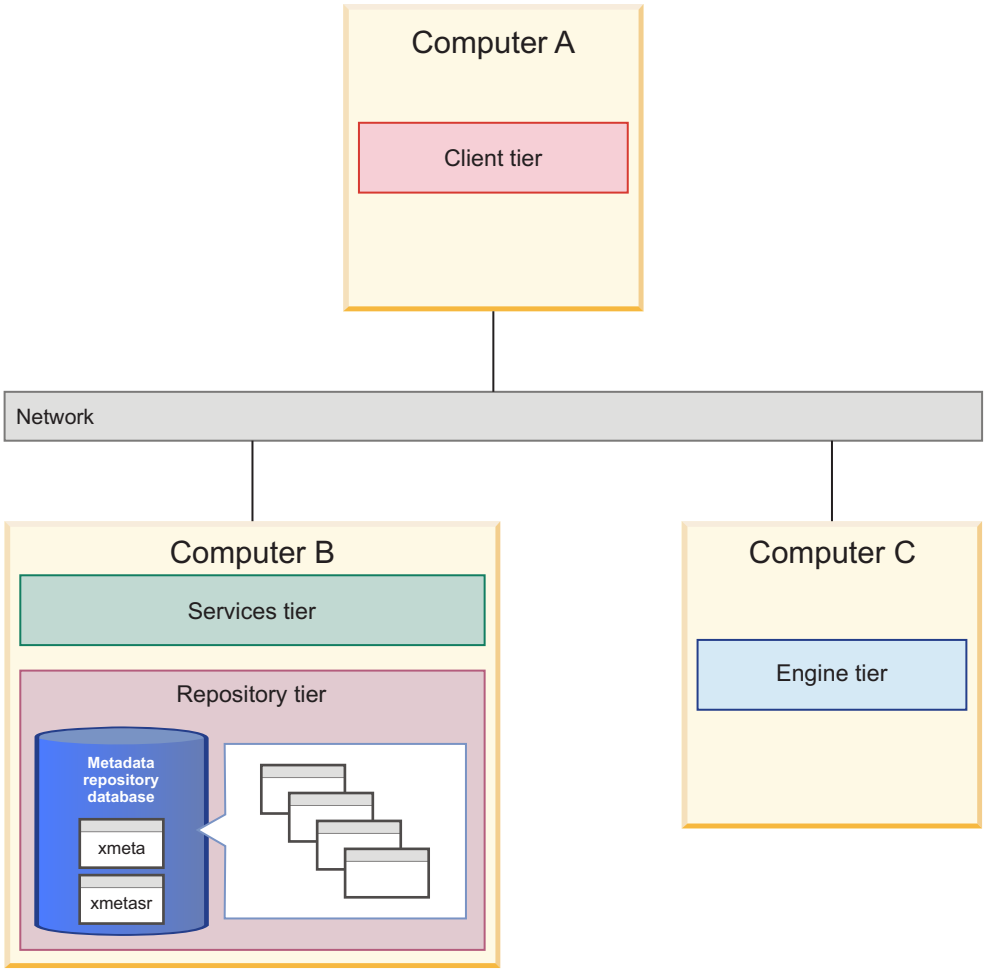


Figure 16. Topology with data stores as separate schemas within the metadata repository database

The following diagram illustrates a four-computer topology where the analysis databases and other data stores are on a separate computer from the metadata repository database.

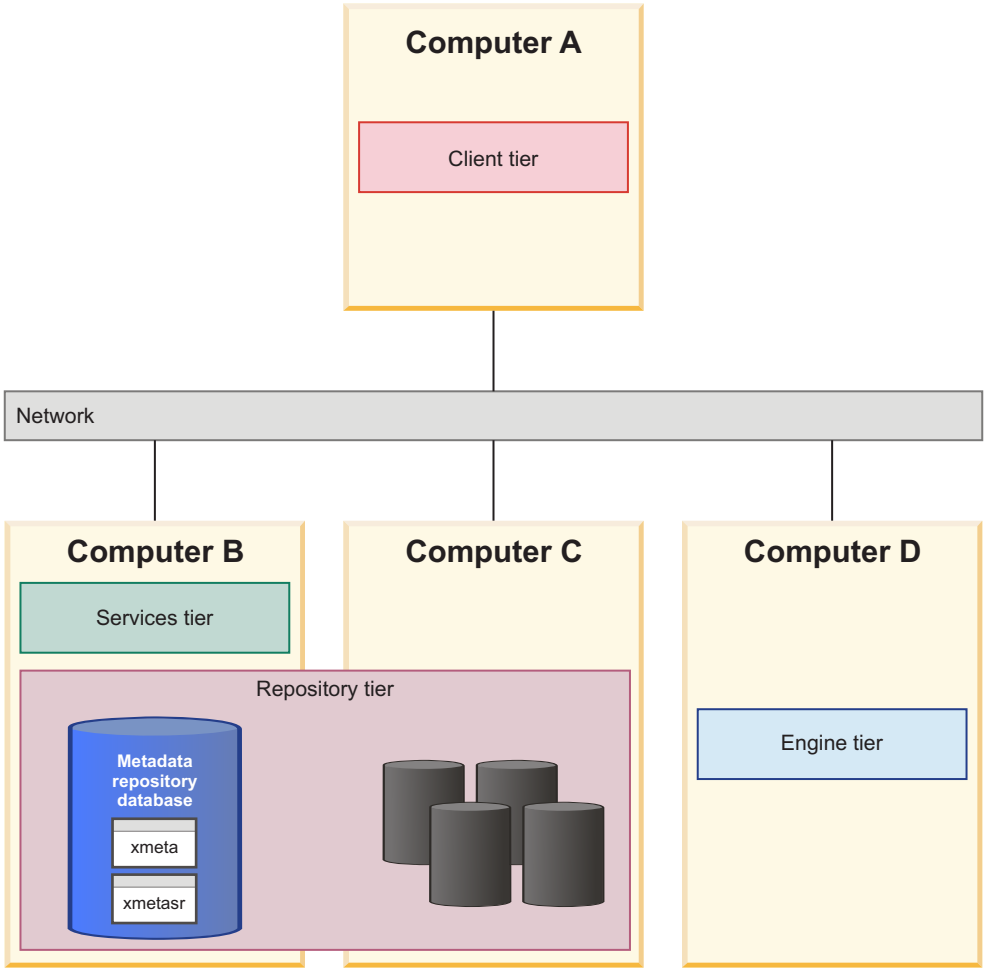


Figure 17. Topology with data stores and metadata repository database on different computers

High availability topologies:

IBM provides different high availability solutions for each IBM InfoSphere Information Server tier.

Increasing availability (implementing a *high availability* solution) refers to maximizing the percentage of time that the system services are operational. To increase availability, you implement topologies and technologies that introduce redundancy. The aim is to reduce or eliminate the number of *single points of failure*. Single points of failure are elements whose failure causes critical aspects of the system to stop operating.

IBM provides different high availability solutions for each InfoSphere Information Server tier. With each solution, you can design many different highly available configurations, from relatively simple setups to complex installations.

The following table lists high availability solutions for each tier:

Table 3. Tiers and high availability solutions

Tier	Solutions
Engine tier	Active-passive topology managed by high availability cluster-management software, such as IBM Tivoli® System Automation for Multiplatforms
Services tier	Either of the following solutions: • Active-passive topology managed by high availability cluster-management software such as Tivoli System Automation for Multiplatforms • Active-active topology with IBM WebSphere Application Server clustering
Metadata repository tier	Any of the following solutions: • Active-passive topology managed by high availability cluster-management software such as Tivoli System Automation for Multiplatforms • IBM DB2 clustering • DB2 high-availability data recovery (HADR)

In general, the higher the level of overall availability that you want to achieve, the more complex the system that you must design and maintain. High availability systems typically require more hardware. For these reasons, give careful consideration to the level of availability that you require within each software tier in the system. You might want a different level of availability within your development system than you have within your production system.

To install a highly available topology, you must have a solid understanding of network technologies such as protocols, tiers, and devices. Successfully deploying highly available topologies, especially ones that include clustering, is a technically complex process that requires a high degree of technical expertise.

The server and software topology that you choose is only part of a high availability solution. You must also introduce redundancy at a number of different levels to reduce or eliminate single points of failure. When determining which topology and implementation to choose, consider the following design factors.

Level of high availability

When considering your topology, evaluate the level of high availability that you require:

- Consider the amount of automation that you need. Must the system take care of failover and recovery automatically, or is a system administrator available to intervene?
- Consider how your needs might differ depending upon how the system is used, such as whether it is a development system, testing system, or production system. For example, how important is it that the development system or testing system is highly available?
- Consider the level of high availability that you require for different software tiers. The tiers are used differently depending upon whether the system is used

primarily for development or in production. The tiers are also used differently depending upon the product modules that you have installed.

Important: For products that have components installed on multiple tiers, a failure on one tier typically renders the entire system nonoperational until the problem is fixed or a failover occurs. Likewise, this would not apply to client-only products, unless the client computer is affected.

Performance and throughput requirements

The amount of scalability that each high availability solution offers differs from solution to solution. When you are choosing a topology and architecture, consider how the system might need to scale to support greater performance and throughput requirements in the future. Other considerations include the configuration of your network, your electrical infrastructure, and your backup, restore, and disaster recovery contingencies. For examples of how the solutions differ, see the scenario comparisons.

See "Capacity planning" on page 54 to analyze your performance, throughput, and storage needs.

Security requirements

Different topologies lend themselves to different security possibilities. Consider how you want to implement firewalls and other security precautions between the different tiers, between the computers in each tier, and between the system and external data sources and targets. These considerations are important for the services tier and the engine tier that must have fast communications paths with external systems.

For more information about security, see the IBM WebSphere Application Server documentation:

Complexity factors

High availability and scalability add layers of complexity to the installation. Eliminating single points of failure requires implementation of complex redundant hardware and software components.

When you design your system, consider the amount of complexity that your IT staff must support. Is there sufficient in-house expertise to support the system that you design? If your high availability features fail, the amount of time you need to bring the system online again might negate any uptime gains that your features provide.

Begin with a relatively simple system that is well within your ability to support. As your IT group becomes familiar with support of the system, implement the high availability features on a staggered schedule until your entire system is in place. This approach also gives the group experience with changing the system and provides them with knowledge for future system scaling.

Maintainability factors

Consider the maintenance costs of the system that you design. Determine what aspects of maintenance to automate, and which aspects to leave in the hands of support personnel.

High availability features can simplify system maintenance. For example, in a clustered system, you can take servers offline for certain updates without rendering the entire system nonoperational.

In an organization where specific departments are responsible for different systems, consider isolating the tiers on separate computers so each department can "own" a tier. For example, in an organization where a group of database administrators is responsible for corporate databases, consider isolating the metadata repository tier on computers that are within their control.

Cost

Implementing high availability typically adds to the initial cost of the system. Extra hardware, software, training, and other costs make the initial outlay larger. However, the productivity gains that the highly available system provides might make up for these costs.

Active-passive topology:

To create a basic level of availability for IBM InfoSphere Information Server, install it on a file system that is shared by two or more computers. The computers are clustered by using high availability (HA) software.

In this topology, two computers share a storage area network (SAN). The metadata repository tier, engine tier, and services tier are all installed on the SAN. One of the computers (the *active server*) hosts the InfoSphere Information Server instance. The other computer (the *passive server*) is started, but does not run any tier software. A high availability software product such as IBM Tivoli System Automation for Multiplatforms is installed on both the active server and the passive server. The HA software maintains a *heartbeat*: a periodic signal that is sent from the active server to the passive server and that indicates that the active server is operational.

The following diagram illustrates this topology. Although the diagram shows only one client tier computer, you can include multiple client tier computers.

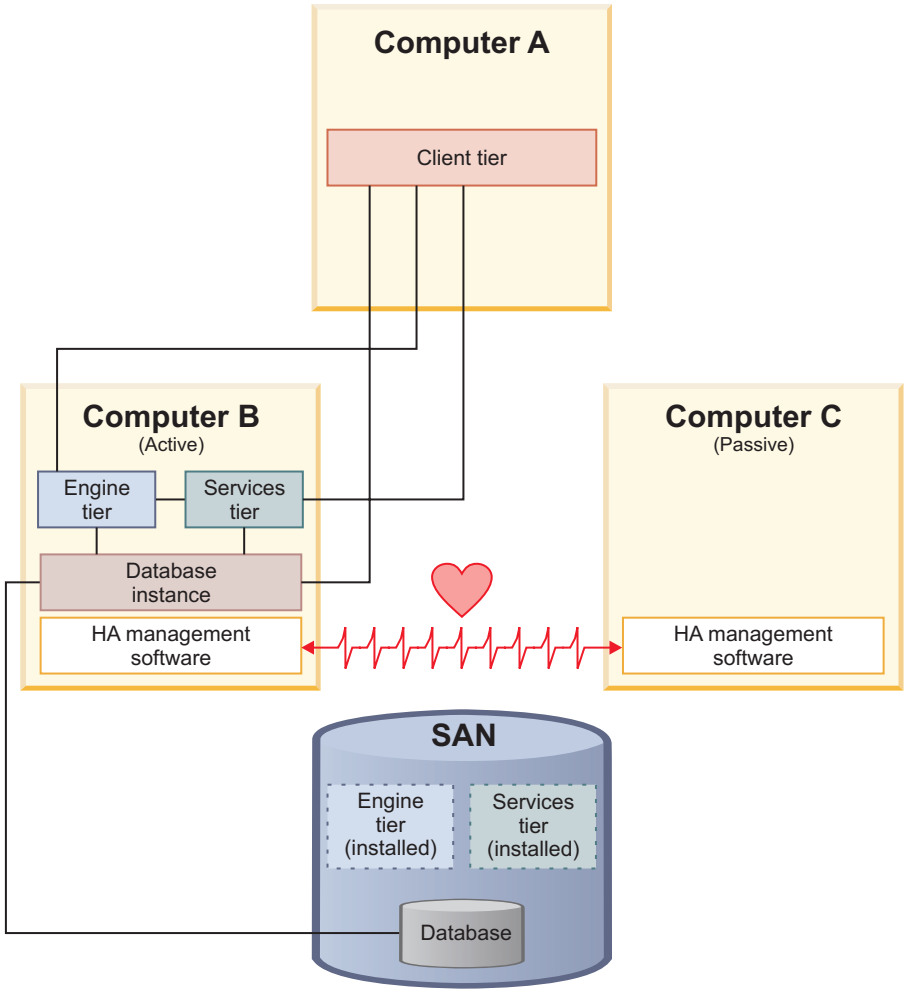


Figure 18. Active-passive topology

If the active server fails, the heartbeat also fails. The HA software restarts all InfoSphere Information Server services on the passive server. This process is called a *failover*. The following diagram illustrates this process.

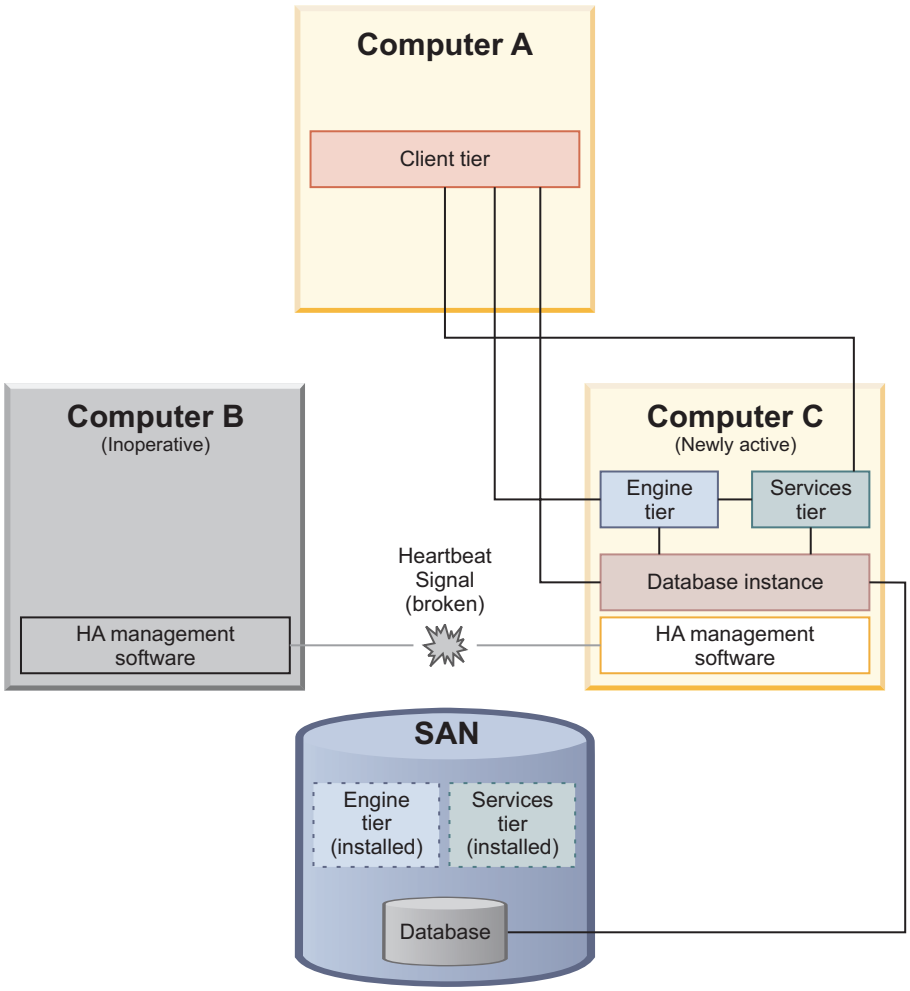


Figure 19. Active-passive server cluster configuration with a failed computer

At installation time, you create a host name alias for the system. This information is associated with the active server. All client programs connect to the system by using this address or host name. If the active server fails, the information is automatically reassociated with the passive server, so client programs do not need to connect to a different address.

However, a failover is not apparent to client users. If the active server fails, all InfoSphere Information Server services are unavailable until the failover is complete and the passive server is operational. The failover process can take several minutes. Also, any jobs that were running fail and must be restarted after the passive server has taken over.

In this configuration, the HA software monitors the health only of the server hardware and operating-system-level processes. The software initiates a failover only if one of these elements fails. Other software monitors the health of individual InfoSphere Information Server processes, but does not trigger a failover if a process fails. Instead, the system administrator is notified. By using this design, a system administrator can correct the problem, instead of allowing the problem to be reproduced on the passive server.

Some other variations of this topology are possible, for example:

- Computer B hosts the services tier and the metadata repository tier. Computer C hosts the engine tier. The computers can fail over to one another. For example, if Computer B fails, Computer C also takes over the services tier and the metadata repository tier. If Computer C fails, Computer B also hosts the engine tier.
- Computer B hosts the services tier and the metadata repository tier. Computer C hosts the engine tier. Another computer (Computer D) serves as backup server for both Computer B and Computer C.

These variations are beyond the scope of this documentation.

To simplify the configuration further, you can eliminate the HA software from your implementation, and have your IT department to trigger a failover manually when necessary.

Topology dimensions and analysis

The active-passive topology is appropriate for small- to medium-sized companies with few concurrent system users. The topology provides a basic level of high availability without incurring significant additional server hardware costs. This configuration is suitable mainly for systems that run scheduled production jobs. The types of activities involved in a development environment make it more likely that a failover would cause dynamic file corruption or synchronization problems.

This topology is not scalable. You can add other passive servers if necessary. However, performance, throughput, and concurrency are not improved because at any specific time, only one instance of each tier component is running in the configuration. Among the small scale InfoSphere Information Server deployments, this topology is popular and rates high on the performance and high availability dimensions while considering the hardware, skills, and complexity dimensions.

Table 4. Analysis: Two-server, active-passive topology

Consideration	Advantages and disadvantages
Level of high availability	• A host failure causes all tiers to be restarted. • When a node fails, there might be up to a 10-minute outage. • If any component fails, there is the potential for job design loss (client disconnected). • Software failure might not cause a failover.
Scalability	• The topology is difficult to scale because adding nodes does not increase performance or throughput. • To increase engine performance, add parallel engine computer nodes.
Security	• Complexity is fairly low. • A security breach on the server might jeopardize the entire system.
Complexity	• Complexity is fairly low. • No IBM WebSphere Application Server or database clustering experience is required. • HA software might be difficult to implement.
Maintainability	By using failover, you can take one server offline for maintenance, while the rest of the system remains operational.
Cost	This topology requires minimal hardware and software.

Clustered topologies:

To maximize the high availability capabilities of IBM InfoSphere Information Server, install it on a file system that is shared by eight or more computers in which all tiers are implemented within clusters.

In this topology, all tiers are configured with the most advanced high availability solutions, and both application server clustering and metadata repository tier clustering are implemented. The cost in hardware is high because a minimum of eight computers are required:

- Two computers for the engine tier, configured with the active-passive failover solution
- Three or more computers for the services tier, configured as an IBM WebSphere Application Server Network Deployment cluster (the WebSphere Application Server Deployment Manager and managed nodes are deployed on dedicated computers)
- One computer for the front-end dispatcher (typically a web server) that sends requests to the IBM WebSphere Application Server Network Deployment cluster
- Two computers for the metadata repository tier that runs on IBM DB2 high availability disaster recovery (HADR) with the Automatic Client Reroute

The following diagram illustrates a topology with a clustered application server and metadata repository tier with clustered DB2.

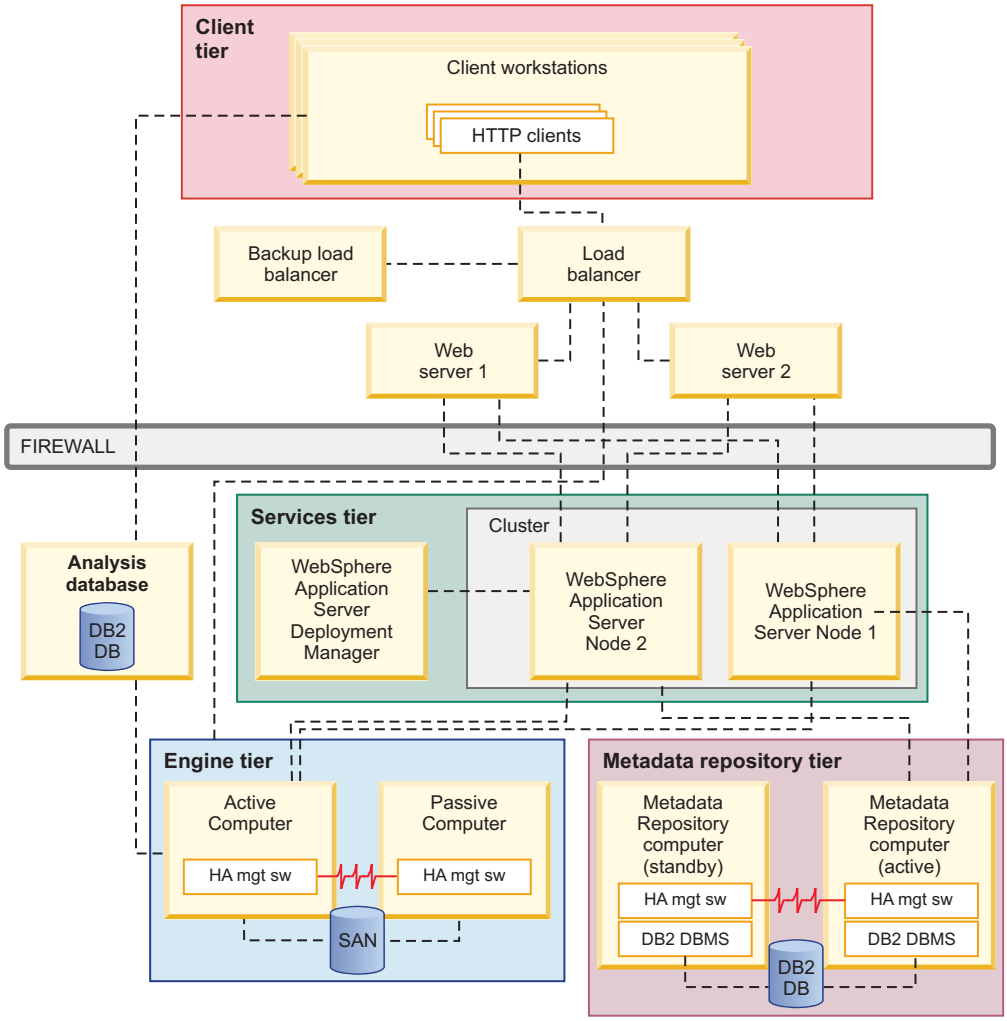


Figure 20. Fully clustered topology

Topology dimensions and analysis

You can adapt this topology as needed to meet available resources and specific requirements. For example, to reduce hardware cost, you can consolidate some of the tiers and components on fewer computers. You can also collocate some of the services tier cluster nodes with the DB2 HADR instances of the metadata repository tier or collocate the front-end web server with one of the Services tier nodes. Similarly, if performance and scalability at the services tier is more important than high availability, then you might consider configuring the metadata repository tier with the active-passive failover solution rather than with the DB2 HADR technology. Additionally, the engine tier and metadata repository tier can be consolidated on the same computers to mitigate the cost of two additional computers.

Table 5. Scenario analysis: Fully clustered topology

Consideration	Advantages and disadvantages
Level of high availability	• The failure of a node in the WebSphere Application Server cluster does not cause service interruption. However, a slight degradation of performance might occur. • The failure of a node in the metadata repository tier cluster causes minimal or no service interruption. • The failure of a node in the engine tier causes up to a 10-minute interruption in engine operation, while failover occurs and the other node takes over. However, activities that do not require the engine tier continue without interruption.
Scalability	• The topology is highly scalable. • The topology supports many concurrent users. • You can add nodes to the WebSphere Application Server cluster to improve the performance and throughput of the application server. • To increase engine performance, add additional parallel engines.
Security	• The topology requires a complex security setup that includes multiple firewalls and other security considerations. • If the security of one node or tier is compromised, the entire system is not necessarily compromised.
Complexity	• The topology is highly complex. • You must understand WebSphere Application Server clustering. • You must understand IBM DB2 clusters if you are using clustered or HADR DB2. • The topology might require the implementation of high-availability cluster-management software, which might be difficult to implement.
Maintainability	• The maintenance of this topology is complex. • The topology is highly flexible because you can take a server offline for maintenance, and the rest of the system remains operational. • The isolation of each tier on one or more separate computers enables groups of users within the organization to be responsible for or use particular tiers.
Cost	The cost is high.

Application server clustering:

An application server runs your J2EE applications and can be stand-alone or managed. If you use IBM WebSphere Application Server Network Deployment, you can cluster your application servers to run the same set of J2EE applications and participate in workload management.

Clusters enable enterprise applications to scale beyond the amount of throughput capable of being achieved with a single application server. Clusters also enable enterprise applications to be highly available because requests are automatically routed to the running servers in the event of a failure. The servers that are members of a cluster can be on different host computers. In contrast, servers that are part of the same node must be located on the same host computer. Clustering is done in a cell that contains a deployment manager and one or more nodes that are managed by the deployment manager. A cell can include one or more clusters.

In a WebSphere Application Server Network Deployment clustered topology, IBM InfoSphere Information Server deploys the services tier applications in each application server.When a member of the cluster fails, other cluster members continue to provide services. Components in other tiers communicate with the cluster, not directly with a specific application server in the cluster. As long as at least one member of the cluster is operational, there is no interruption in service.

Implementing an IBM WebSphere Application Server Network Deployment cluster improves concurrency. *Concurrency* is the number of clients that can use the system simultaneously with a reasonable level of system responsiveness. You can scale a clustered configuration in order to meet increased concurrency requirements.

The following diagram shows a topology with a services tier where application server clustering is implemented across two computers. A third computer hosts the WebSphere Application Server Deployment Manager. In the diagram, the engine tier is still set up in an active-passive configuration. The metadata repository tier is located on a separate computer.

This topology includes a sophisticated front-end configuration that includes dual web servers and a load balancer. These systems are set up outside the firewall. HTTP clients such as the WebSphere Application Server web console access the server-side components through the load balancer. A backup load balancer ensures that the load balancer computer is not a single point of failure.

Important: IBM WebSphere Application Server Network Deployment does not support deployment of firewalls between components within a cell (for example, between the deployment manager and the nodes).

Note: High availability configurations that include WebSphere Application Server clustering typically also include metadata repository clustering. The database cluster ensures that the metadata repository is not a single point of failure.

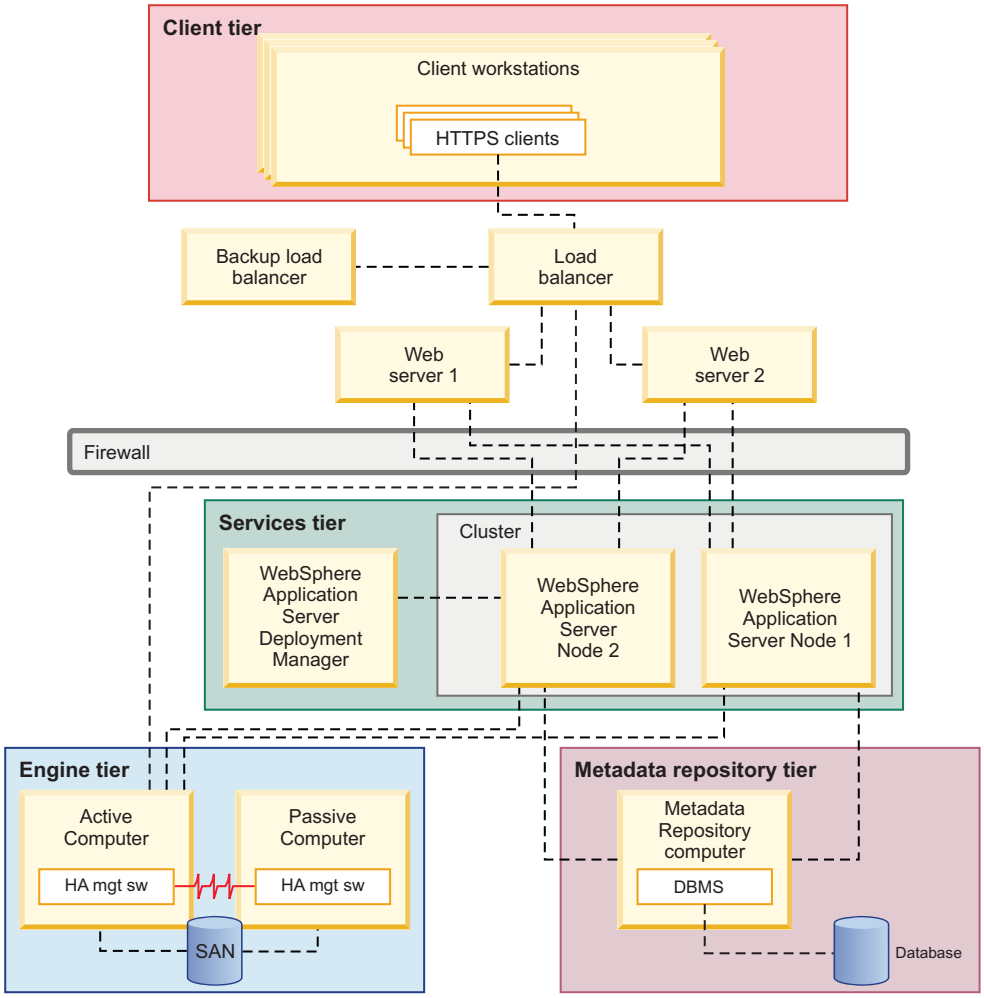


Figure 21. Topology of an application server cluster with a load balancer

Options for installing on a cluster:

When you install IBM InfoSphere Information Server, you have a choice of whether to install by using a front-end web server or a cluster member. This is a new feature as of Version 11.3.

It is recommended to take advantage of this new feature and install InfoSphere Information Server by using a cluster member and configure the front-end web server after installation. By doing so, you avoid potential complications during installation if there are problems with the front-end web server configuration. Some of the items to consider in a front-end web server configuration are as follows:

- All communication with InfoSphere Information Server is through SSL and cell level certificates must be imported into the front-end web server and plugins.
- Web server port numbers must be specified into application server virtual host aliases.
- The front-end web server must be able to access InfoSphere Information Server through the firewall.
- A set of InfoSphere Information Server web applications are deployed in the cluster and some of these applications are accessed through REST APIs. Therefore, the front-end web server must be pre-configured to propagate SSL

requests to the back-end application server. Or, the front-end web server must be fully managed by the application server with automatic update and propagation of the HTTP plugin when new web applications are deployed.

A front-end configuration to an IBM WebSphere Application Server cluster:

If you deploy a WebSphere Application Server cluster within your IBM InfoSphere Information Server implementation, the resulting configuration would typically have a front-end configuration for the cluster unless the cluster has only one member. When you install InfoSphere Information Server, you can install against a front-end configuration or a cluster member. If you install against a cluster member, you would typically configure the front-end configuration for the cluster after installation.

The front-end configuration can consist of any of the following components located upstream of the cluster:

- A front-end web server
- A load balancer
- A load balancer upstream of the web servers (IP sprayer topology)

To maximize security, deploy firewalls at various key points within the topology that you select. To improve delivery response, you can also deploy caching proxy servers upstream of the web servers and load balancers.

Front-end web servers:

If you deploy an IBM WebSphere Application Server cluster within your IBM InfoSphere Information Server implementation, you can provide a front-end web server upstream of the cluster.

InfoSphere Information Server supports IBM HTTP Server and Apache HTTP Server. These software products are not shipped with InfoSphere Information Server.

A front-end web server manages servlet and JSP requests. The web server typically uses an internal workload management (WLM) plug-in to dispatch web requests to the appropriate cluster members and their containers. It performs the dispatches according to policies that are preconfigured in the WLM plug-in. The front-end web server increases performance and security because the web server provides a unique HTTP entry point to the cluster.

To prevent single points of failure at the web server level, deploy a backup dedicated web server to take over in case the active one fails.

The following diagram shows an InfoSphere Information Server cluster with a front-end web server. A backup web server is also deployed.

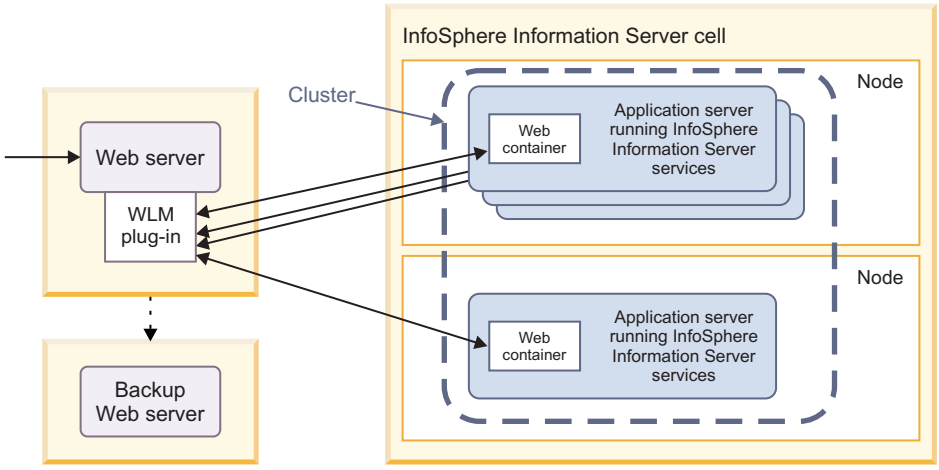


Figure 22. Cluster with front-end web server

If you choose IBM HTTP Server, you can set up the Deployment Manager to manage configuration and administration of the web server. The web server can be defined either as a managed node or an unmanaged node.

- If you want to use the Deployment Manager to manage web servers in the same manner as any other managed nodes in the cell, define the web servers as managed nodes. For example, you can use the Deployment Manager to start or stop IBM HTTP Server or to install the IBM HTTP Server plug-in configuration file that contains the workload management policies that are defined for the cluster.
- If you deploy the web servers upstream of a firewall that protects the cluster members, define the web servers as unmanaged nodes. In this case, you can manage the web servers with the Deployment Manager by enabling the IBM HTTP Server Admin Process on the same computer as the IBM HTTP Server by opening its port through the firewall.

The following diagram shows an IBM HTTP Server instance that is defined as a managed node. The Deployment Manager manages IBM HTTP Server through the local node agent.

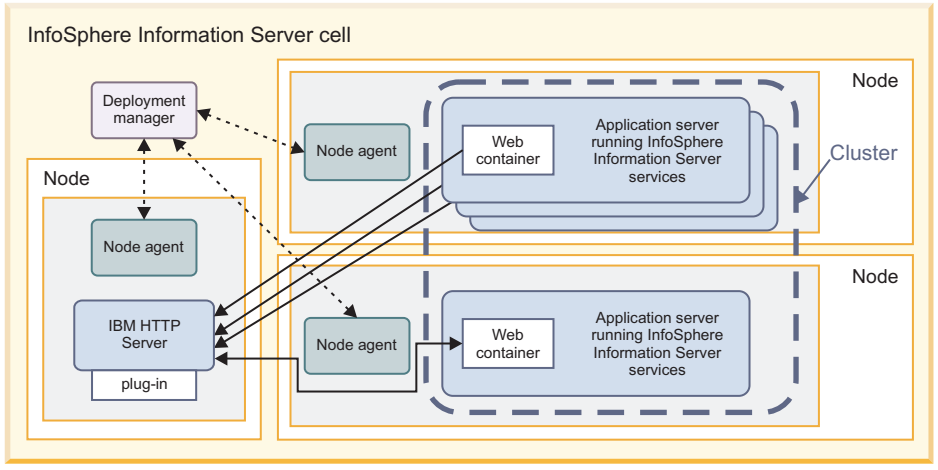


Figure 23. IBM HTTP Server instance that is defined as a managed node

The following diagram shows an IBM HTTP Server instance that is defined as an unmanaged node on a stand-alone server. A firewall separates the server from the cell. The Deployment Manager manages the IBM HTTP Server instance by using the IBM HTTP Server Admin Process.

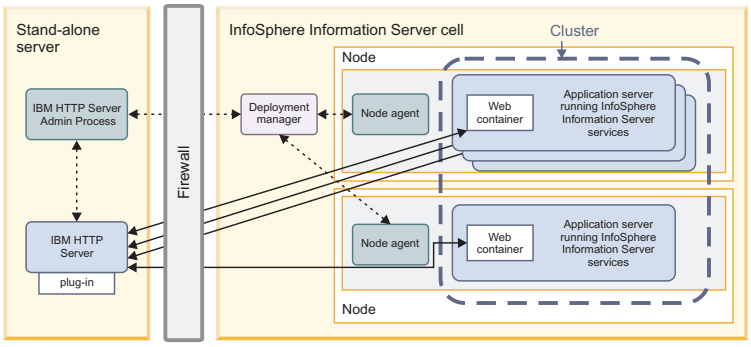


Figure 24. IBM HTTP Server instance that is defined as an unmanaged node on a stand-alone server

If you want to use the Apache HTTP Server outside of a firewall, you can run it as an unmanaged node. If you want the Apache HTTP server to be managed, then it must be on the same machine as a node agent that is managed by the deployment manager (as shown in Figure 2), and thus it cannot be on the other side of any firewall from the deployment manager. If unmanaged, copy the web server plug-in configuration files to each web server manually after services tier installation or a fix pack installation.

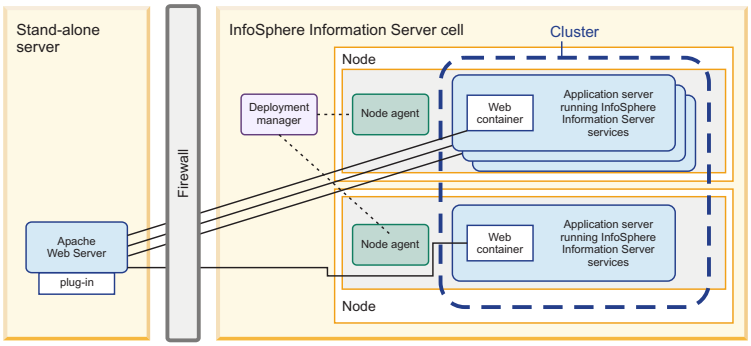


Figure 25. Apache HTTP Server

Load balancing:

If you deploy an IBM WebSphere Application Server cluster within your IBM InfoSphere Information Server implementation, you can provide a load balancer upstream of the cluster.

InfoSphere Information Server supports many load balancer solutions manufactured by IBM and by other vendors.

The following list highlights the advantages and disadvantages of a load balancer approach.

Advantages of deploying a load balancer

- Easier deployment because you do not have to set up a network of web servers and then configure the web server plug-ins.

- More load balancing algorithm alternatives are available. Load balancers often offer more algorithms than the ones provided by web server plug-ins, including advanced load-based balancing strategies that monitor usage levels on individual computers.

Disadvantages of deploying a load balancer

- Additional configuration is required. You must perform additional configuration to maintain persistent connections between clients and servers (web server plug-ins handle persistent connection maintenance automatically). Also, you must re-configure the load balancer whenever the downstream cluster topology changes (for example, when a node is added or removed).
- Hardware-based load balancers typically cost more.

To prevent single points of failure at the load balancer level, consider deploying a backup load balancer to take over in case the active one fails.

The following diagram shows an InfoSphere Information Server cluster with a front-end load balancer. A backup load balancer is also deployed.

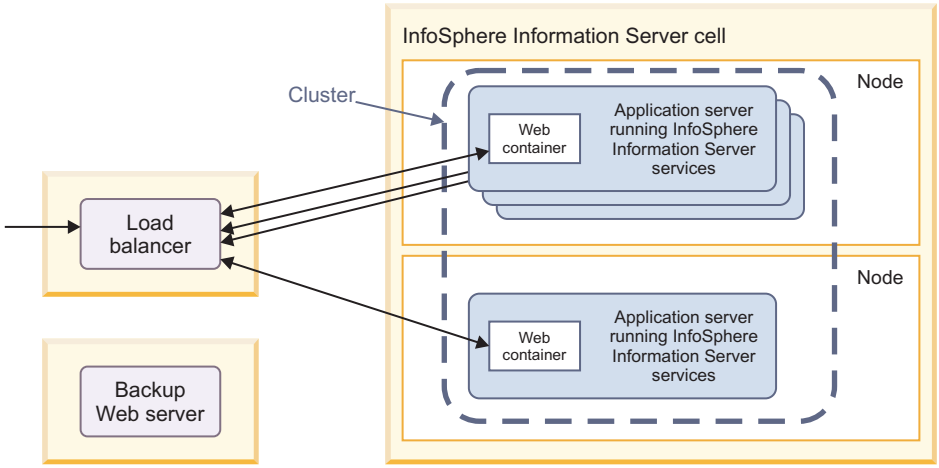


Figure 26. Cluster with a front-end load balancer

You cannot manage the load balancers by using the Deployment Manager.

Make sure to configure session affinity in the load balancer. There are several methods to achieve session affinity. The method you choose depends on the product that you use and the forwarding mechanism that you configure. Mechanisms include media access control (MAC) address forwarding, Network Access Translation/Network Address Port Translation (NAT/NAPT) forwarding, and content-based forwarding. Layer 2 load balancers (load balancers that operate at the MAC layer) generally achieve session affinity by using a "stickiness to source IP address" approach. Layer 3 and higher load balancers (load balancers that operate at the IP layer and above) typically achieve session affinity by using passive cookies or a content-based technique. Session affinity is typically implemented in this manner with HTTP servers. See the load balancer documentation for more information

IP sprayer topology:

To obtain high availability and the maximum balancing of server and workload capacity, deploy a load balancer upstream of a group of web servers.

The load balancer performs intelligent balancing among the web servers based on server availability and workload capacity. Choose this topology to eliminate a single point of failure at the web server level. This arrangement also spreads the load of web requests across multiple web servers. IBM InfoSphere Information Server supports many IP sprayer solutions manufactured by IBM and by other vendors.

To prevent single points of failure at the load balancer level, deploy a backup load balancer to take over in case the active one fails.

The following diagram shows an IP sprayer topology. The diagram shows an InfoSphere Information Server cluster with two front-end web servers. A load balancer is deployed upstream of the web servers. A backup load balancer is also deployed. The web servers and load balancers are installed on separate servers.

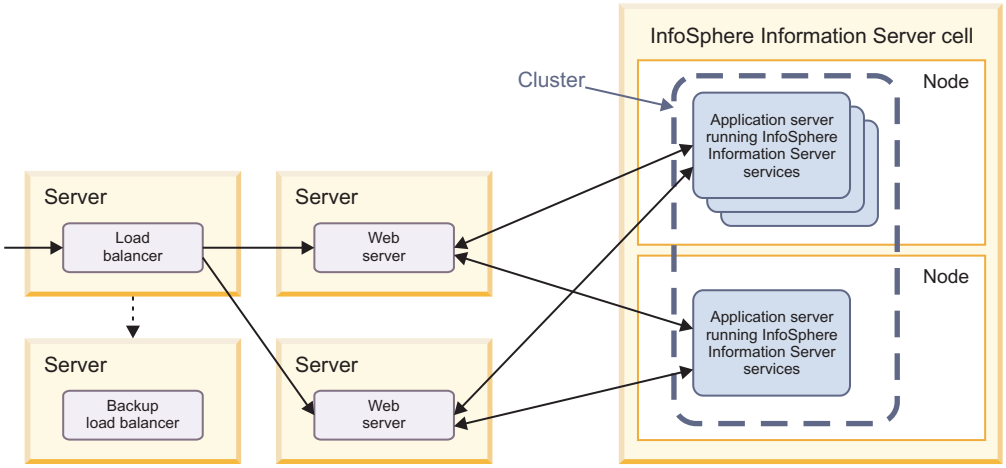


Figure 27. IP sprayer topology

Firewalls:

To secure access to your IBM InfoSphere Information Server cluster, deploy firewalls at various key points in the topology that you select.

For example, add a firewall in front of the caching proxies, load balancers, and web servers. Isolate them within a DMZ to deny direct access to internal components from clients that you do not trust. Alternatively, you can place a firewall between the DMZ and the cluster. You can also add a firewall between the services tier and the metadata repository database.

Important: IBM WebSphere Application Server does not support the deployment of firewalls between components within a cell. For example, the application server does not support firewalls between nodes in a cell, or between nodes and the Deployment Manager in a cell.

The following diagram shows an InfoSphere Information Server cluster behind a DMZ. A firewall is located between the DMZ and the outside world. A firewall is

also located between the cluster and the DMZ. For maximum security, a firewall is also located between the cluster and the metadata repository database.

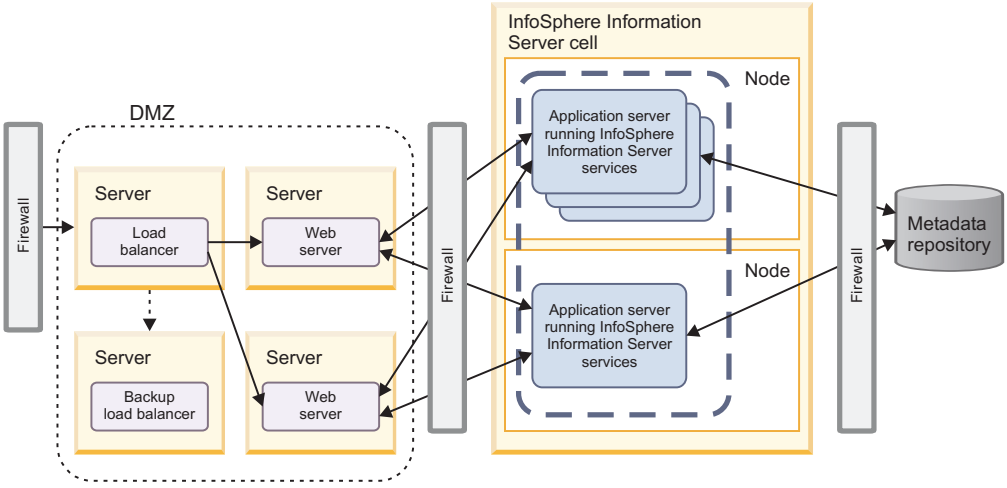


Figure 28. Cluster with firewalls

If you deploy firewalls, make sure to open any ports through the firewalls that are used by InfoSphere Information Server.

Database clustering:

To provide a high availability topology, you can create database clusters across computers. By clustering your databases, you combine the computing power of the servers that are involved to provide higher scalability, more combined computing power, or built-in redundancy that provides higher availability of your information.

DB2 database clustering:

IBM offers several high availability DB2 configurations for the metadata repository database.

IBM InfoSphere Information Server supports these configurations:

- DB2 clustering by using high availability clustering software
- DB2 with high availability disaster recovery (HADR)

In either configuration, you can use DB2 automatic client reroute to enable IBM InfoSphere Information Server processes to reconnect to a standby node when a failover occurs.

Database clustering

To provide a high availability DB2 configuration, you can create a DB2 cluster across computers. In this configuration, the metadata repository database is shared between nodes in the cluster. If a failover occurs, another node in the cluster provides DB2 functionality. To provide high availability, set up your cluster in an active-passive configuration with a single active DB2 instance on one computer and one or more passive instances on the other computers. If the DB2 instance encounters a problem or fails, a passive instance can take over.

To manage this configuration, you can choose between several high availability clustering management software products. This software maintains a "heartbeat" signal between the nodes in the cluster. If the heartbeat fails on the active node, the software initiates failover to another node.

With this configuration, the DB2 failover is automatic, but might take several minutes as the new instance acquires resources, repeats certain transactions, and undoes other transactions. To minimize interruption and manual intervention, configure DB2 automatic client reroute. This function causes other components in the IBM InfoSphere Information Server instance, such as IBM WebSphere Application Server, to automatically reconnect to the new DB2 instance.

This configuration does not provide redundancy for the database itself. Instead, it provides high availability for database client processes and smooths the reconnection to the new node. To provide redundancy for the database itself, implement high availability disaster recovery (HADR).

The following diagram shows a topology that includes a clustered metadata repository tier. In this case, HA cluster management software on the metadata repository computers monitors DB2 processes specific to DB2. The software also monitors the health of the computer hardware and the network.

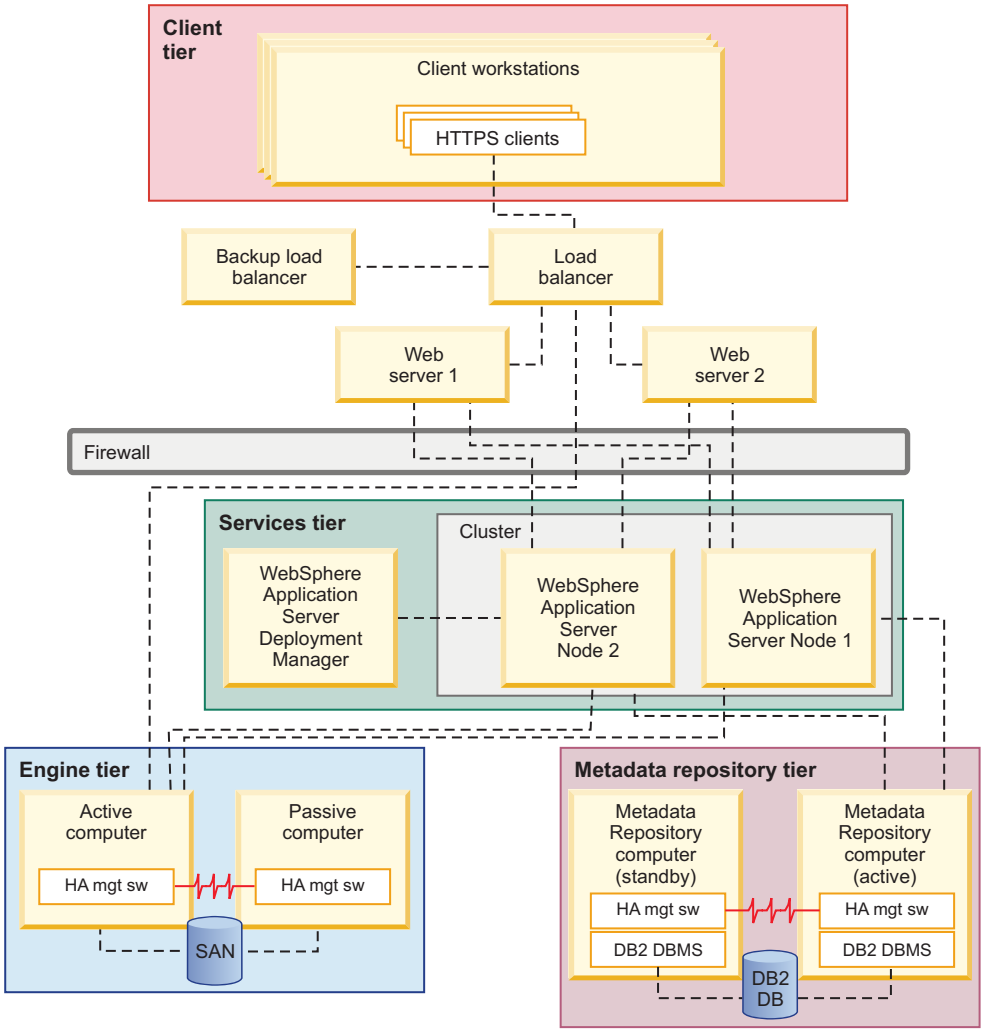


Figure 29. Topology with a clustered metadata repository tier

High availability disaster recovery (HADR)

To provide high availability at the DB2 database level, you can create a DB2 high availability disaster recovery (HADR) configuration. In this configuration, a complete, separate copy of the database is maintained on a standby node at a local or remote location. The primary DB2 database processes transactions. It uses internal HADR processes to replicate transaction logs to the standby node, where the logs are stored. A process at the standby node then replays the logged transactions directly to the standby database. The two copies are maintained in synchronization or near-synchronization with each other.

DB2 processes normally access only the primary copy of the database. However, if the primary copy fails, an administrator triggers the standby copy to take over the transactional workload. You can set up automatic client reroute to make the failover behavior nearly transparent. If the primary copy becomes unavailable, automatic client reroute tries the connection again to the failed primary copy. If the reconnection fails, automatic client reroute determines whether the standby copy is available. If the server is available, automatic client reroute reroutes the application server connection to the standby copy. Pending transactions are rolled back and then reissued on the standby copy. Automatic client reroute can also be configured

to reconnect to an alternate standby copy if the first standby copy is unavailable. Failover is quick (normally 10 seconds to 15 seconds if the servers are on the same LAN segment).

If the unavailable database becomes available again, it is automatically reintegrated as the new standby database and is resynchronized.

The following diagram shows an IBM InfoSphere Information Server implementation in which the metadata repository tier is set up in an HADR configuration.

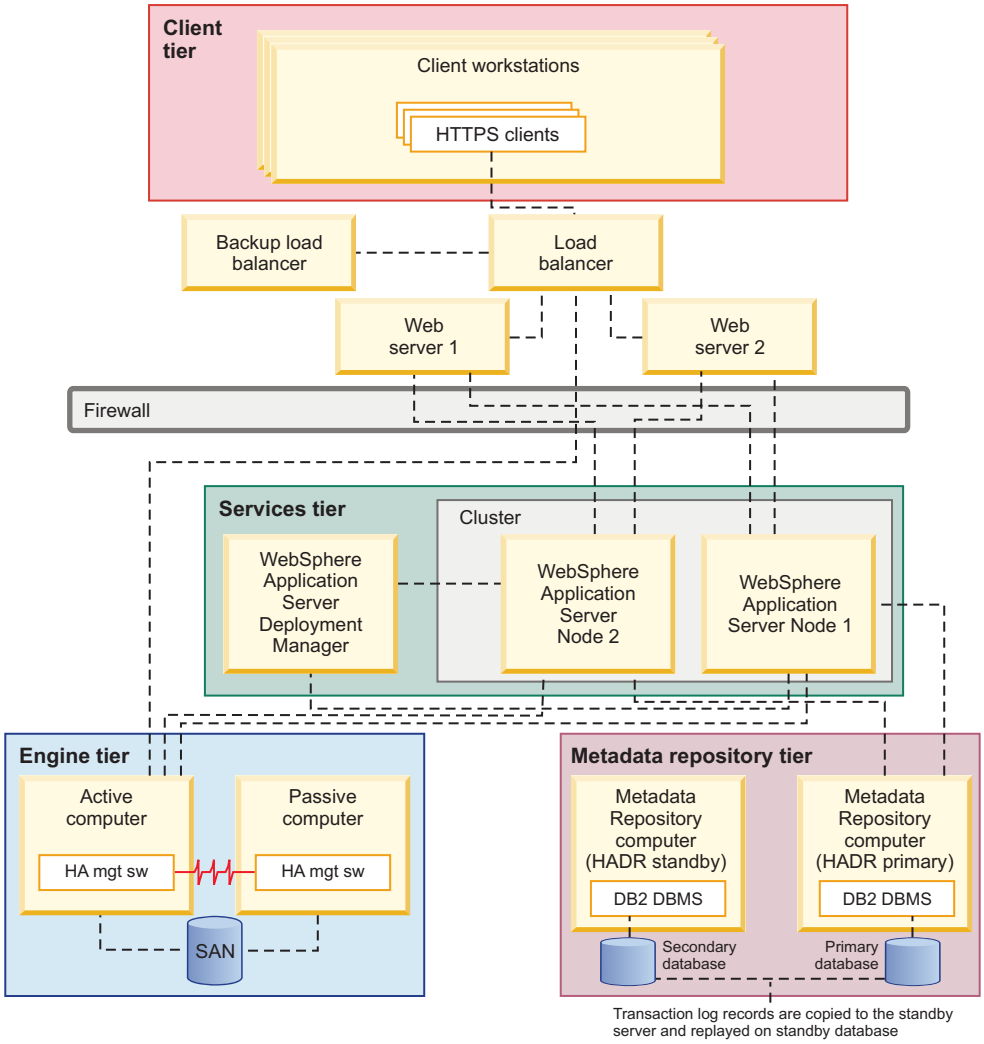


Figure 30. HADR configuration

Parallel processing and grid topologies:

To maximize the IBM InfoSphere Information Server engine efficiency and throughput, create a parallel processing or grid topology.

In these topologies, the workload for each job is distributed across several processors. The processors can be located on the same computer or on different computers that are connected over a network. The parallel engine uses the

processors and computers concurrently to complete each job. The parallel engine and the InfoSphere Information Server engine make up the engine tier.

You can configure the InfoSphere Information Server parallel framework for grid computing processing. Install a supported resource manager to enable the dynamic assignment of available computer nodes (physical computers) at run time yo.

Parallel processing topologies:

In a parallel processing topology, the workload for each job is distributed across several processors.

In IBM InfoSphere DataStage, you design and run jobs to process data. Normally, a job extracts data from one or more data sources, transforms the data, and loads it into one or more new locations.

In a parallel processing topology, the workload for each job is distributed across several processors on one or more computers, called *compute nodes*. Within InfoSphere DataStage, the user modifies a configuration file to define multiple processing nodes. These nodes work concurrently to complete each job quickly and efficiently. A *conductor node* computer orchestrates the work.

Parallel processing environments are categorized as *symmetric multiprocessing* (*SMP*) or *massively parallel processing* (*MPP*) systems.

Symmetric multiprocessing (SMP) systems

In a symmetric multiprocessing (SMP) environment, multiple processors share other hardware resources. In the following diagram, multiple processors share the same memory and disk space, but use a single operating system.

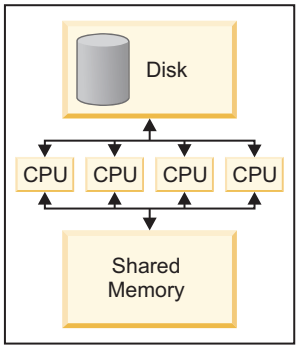


Figure 31. Symmetric multiprocessing (SMP) system

The workload for a parallel job is distributed across the processors in the system. The actual speed at which the job completes might be limited by the shared resources in the system. To scale the system, you can increase the number of processors, add memory, or increase storage. The scalability strategy that you implement depends on how your job is limited within your current system.

Massively parallel processing (MPP) systems

In a massively parallel processing (MPP) system, many computers are physically housed in the same chassis, as shown in the following diagram:

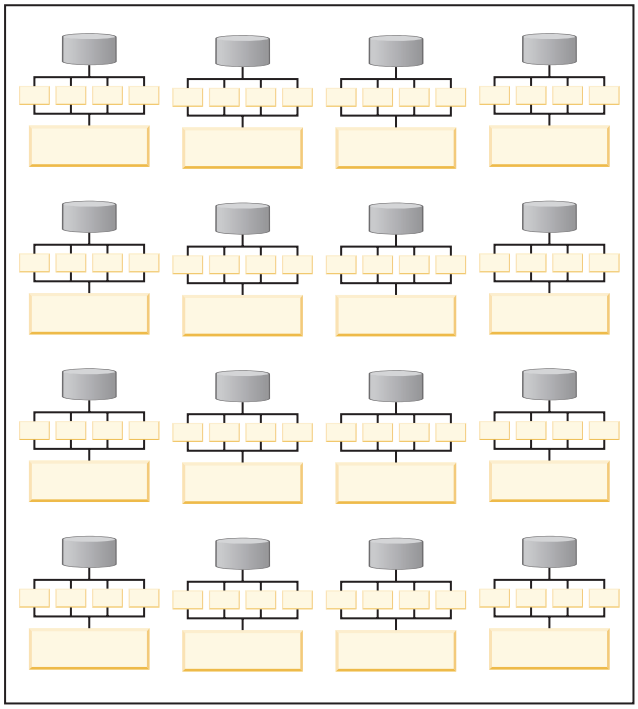


Figure 32. Massively parallel processing (MPP) system

An MPP system is physically dispersed, as shown in the following diagram:

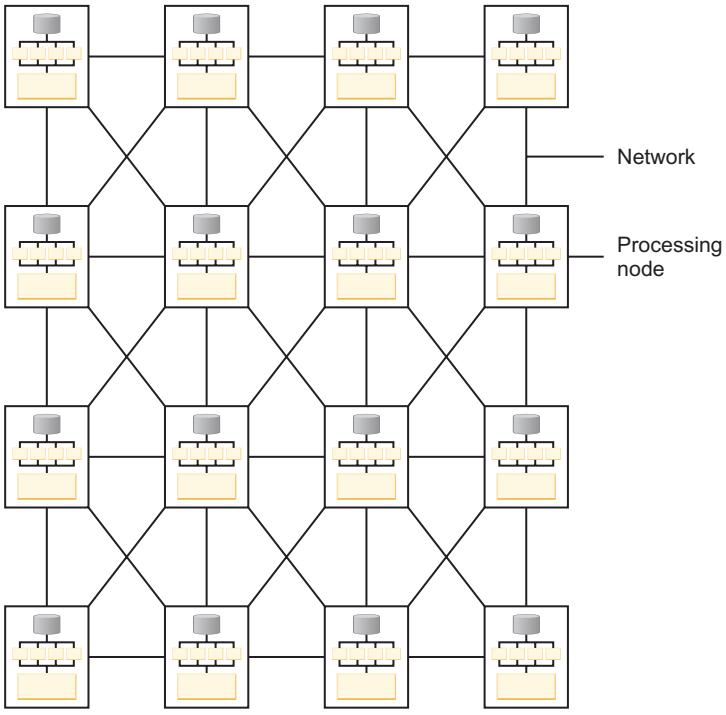


Figure 33. MPP system

In an MPP environment, performance is improved because no resources must be shared among physical computers. To scale the system, you can add computers and associated memory and disk resources.

In an MPP system, a file system is commonly shared across the network. In this configuration, program files can be shared instead of installed on individual nodes in the system.

Grid topologies:

In a grid topology, a resource manager dynamically allocates resources for job processing.

The standard InfoSphere Information Server suite supports grid implementations on AIX and Linux, using the Grid Toolkit, with Platform LSF® as the resource manager. Each computer in the grid must use the same operating system and release level. The computers must be connected and accessible within the same high-speed local area network. You can set up this implementation by using the instructions in this documentation.

In an IBM InfoSphere Information Server parallel processing topology, you modify a configuration file to assign resources to a job. To change the degree of parallelism, or the servers on which a job runs, you modify the configuration file to include the new number of nodes and associated servers. This architecture makes the association of nodes used by a job static.

With a grid implementation, a resource manager manages the pool of available resources and dynamically allocates them as jobs are submitted for processing. For each job, the resource manager waits until resources are available and then generates a configuration file that assigns resources to the job. The system minimizes the idle time for each processor and helps ensure that resources are used in the most efficient manner.

Grid topologies vary widely. They depend in part on the storage implementation such as network-attached storage (NAS) or storage area networks (SAN). However, in any grid topology, one server is designated as the conductor node, and other servers are designated as compute nodes. All jobs are initiated from the conductor node, and processed on the compute nodes.

A grid can be implemented on relatively low-cost hardware, with minimum additional costs. To scale the grid, add additional nodes and register them with the resource manager. You can take nodes out of service without halting processing of jobs within the grid.

The following diagram shows a typical grid topology. The diagram shows an IBM BladeCenter® chassis with 14 blades. External disk storage is a network-attached storage (NAS) system. The storage is accessible from all blades through the NFS server installed in the first node (node 0).

In the following diagram, the metadata repository tier and services tier are installed on the first blade (node 1). The server engine, resource manager, and one parallel engine are installed on the second blade (node 2). All other blades (nodes 3 through 14) are compute nodes that provide processing power for the grid. The client tier is installed on a workstation that is connected to the blade system (this computer is not shown).

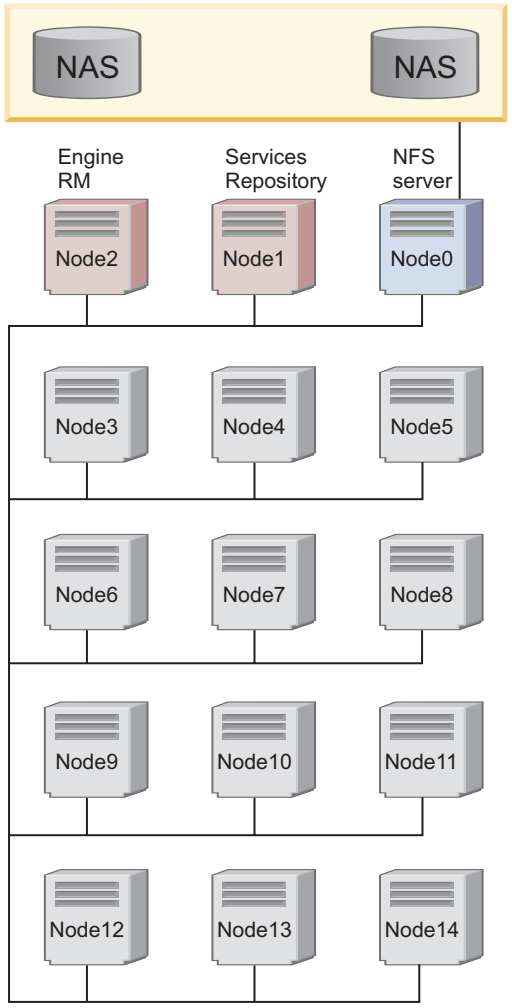


Figure 34. Grid topology

Capacity planning

You plan the use of disk, volume group, and file system resources to create an optimal operating environment for IBM InfoSphere Information Server.

As part of capacity planning for your installation, size your file systems and databases to accommodate your anticipated needs. After the installation, continue monitoring the file systems to ensure that sufficient space is available.

Before you change your computers, back up those computers where you plan to install InfoSphere Information Server.

Attention: If the file system runs out of space, InfoSphere Information Server stops functioning correctly.

File systems

Plan for the following separate file systems:

Temporary storage
Directories for scratch data that is used during processing and later discarded.

Program code
Program code directories for the InfoSphere Information Server installation directory, database server installation directory, and the IBM WebSphere Application Server installation directory.

Storage directories
Database storage directories for the metadata repository database, and for the IBM InfoSphere Information Analyzer databases and IBM InfoSphere QualityStage Match Designer database (if you are installing these product modules).

Isolate each functional subgroup or components in its own separate file system to accommodate different performance and capacity characteristics and backup requirements. File systems need to be expandable without requiring destruction or recreation, if the operating system that you are installing on permits it.

RAID and SAN configurations

Performance of a job (for IBM InfoSphere DataStage, IBM InfoSphere QualityStage, and IBM InfoSphere Information Analyzer) depends on all components being optimized. For RAID (Redundant Array of Independent Disks) and SAN (system area network) configurations, maximum performance is a combination of maximum bandwidth (controllers and disk) and minimized contention.

The following guidelines can assist in the configuration of RAID or SAN technologies:

- Ensure that your database is optimally tuned.
- Ensure that computers where you install the repository tier and services tier are on a high-speed network.
- Minimize contention between temporary systems (scratch, buffer, and sort) and data file systems. For temporary storage, do not use a network file system (NFS). Consider using the local disk for temporary storage.
- Minimize contention between the disks and controllers that are associated with InfoSphere Information Server file systems and other applications or servers.
- Consider isolating multiple mount points to separate high-speed disk interconnects and controllers.
- Consider the trade-off between the granularity of file systems and the underlying configuration in comparison to the available, unused storage.
- Do not create overly complex device configurations. These configurations can be difficult to administer and might not offer noticeable performance improvement.
- If possible, check your configuration with experts in hardware, storage, operating systems, and InfoSphere Information Server.
- Systems with parallelism require different expertise than systems with engines that do not feature parallelism.

The optimal disk configuration balances cost, complexity, ease of administration, and ultimate performance.

Analysis database sizing:

An analysis database is a component that IBM InfoSphere Information Analyzer uses when it runs analysis jobs.

The extended analysis information is stored in the analysis databases. The extended analysis information includes the high-volume, detailed analysis results, such as column analysis, primary key analysis, and domain analysis. Additionally, the metadata repository contains the information analysis projects that contain the analysis results.

Before you create the analysis databases, review the quantity of data to be analyzed. This review helps you to determine an appropriate storage size, location, and configuration of the analysis databases.

When you plan for the size of your databases, consider these factors that affect the size of each database:

- Number of tables to be analyzed
- Number of columns in tables to be analyzed
- Number of unique records within these tables
- Number of `char` and `varchar` columns
- Types of analysis to be done

Unless you use sampled analysis, an analysis database might be larger than the combined size of all the analyzed data sources.

Match Designer database sizing:

After you create the Match Designer database, use automatic storage management to allocate database space as needed. To estimate initial space requirements and plan for the appropriate hardware, you must make some assumptions about the type and volume of matching that plan to accomplish.

Consider the following factors:

- Estimated number of match specifications
- Average number of match passes per match specification
- Average number of input data records per match specification
- Average number of bytes per input data record
- Average number of input frequency records per match specification

When you calculate the number of bytes in a Match Designer input record, assume that varchar columns contain maximum width values. Also, keep in mind that most Match Designer input records are preprocessed through an InfoSphere QualityStage Standardization Stage. This stage adds 500 bytes to a source record or 1000 bytes if the source record is double-byte.

When calculating frequency record volume, consider whether the data values for the source columns that participate in your matches typically have high or low rates of duplication. For data that is highly unique, your frequency record count might be similar to your data record count. For data that has a high rate of duplication, your frequency record count might be significantly less than your data record count.

Review the factors described in Table 2. Then use the factors in the following formula to estimate how many megabytes of space the Match Designer results database is likely to require.

```
(number of match specifications x (match size + (match size x 10%) +
frequency size) x replication factor) / 1,000,000
```

Table 6. Space requirement factors and calculations

Factor	Description	Calculation
Match specification count	Match specifications define and test criteria for matching data.	Approximate number of match specifications you expect to retain in your database
Match pass count	Match passes define the column-level matching criteria applied during a single pass through match specification input.	Average number of match passes per match specification (used in the Match Size calculation)
Data record count	Data records define the content and format of match specification input.	Average number of input data records per match specification (used in the Match Size calculation)
Data record length	Data records define the content and format of match specification input. When calculating data record length, assume sizing scenarios that require the most space.	Average number of bytes per match specification input data record (used in the Match Size calculation)
Match size (in bytes)	Match size aggregates the match pass and data record count and length factors into the estimated space requirements for the data component of a single match specification.	(Match pass count) x (data record count) x (data record length)
Frequency record count	Frequency records indicate how often a particular value occurs in a particular column. For frequency record count, consider high or low rates of duplication for the data values in source columns.	Average number of frequency records per match specification
Frequency size (in bytes)	Frequency size estimates the space requirement for the frequency component of a single match specification.	(Frequency record count per match) x 360

Table 6. Space requirement factors and calculations (continued)

Factor	Description	Calculation
Replication factor	Replication factors account for significant differences in space requirements of different match types: • One-source matches take input from one source, then group and match the data. • Two-source one-to-one matches compare a record on the data source to a record on the reference source. • Two-source many-to-one matches can match any reference source record with many data source records.	Select the replication factor for the match type you expect to run most often: • For one-source matches, use a factor of **1** • For two-source one-to-one matches, use a factor of **2** • For two-source many-to-one matches, use a factor of **5**

Exceptions database sizing:

The exceptions database is used by the IBM InfoSphere Data Quality Console.

The exception stages run on the engine to produce exception descriptors that are stored in the exceptions database as a set of related rows and tables. Each job run that generates an exception adds a record to the descriptor table and creates a new table to hold the exception details for that run. The database will grow as jobs are run. The database content also shrinks when exceptions are deleted. The exceptions database can be managed using the `ESDBAdmin` tool. In general, the size requirements are similar to that of the metadata repository.

Standardization Rules Designer database sizing:

The Standardization Rules Designer database is used by IBM InfoSphere QualityStage.

This data store holds the Standardization Rules Designer rule set revision records. When a rule set is put under revision in the Designer, the database is populated with rule set artifacts retrieved from the metadata repository. The database grows as the number of opened revisions grow. When you delete a revision, the rule set artifacts are deleted from the database. When you publish the revision back to the metadata repository, you have the choice to keep the revision open or delete it. In general, the size requirements are similar to that of the metadata repository.

Installation scenarios

The following installation scenarios describe and illustrate how companies might implement different configurations of IBM InfoSphere Information Server based on the needs of their organization. Use these scenarios to better understand how you plan to use InfoSphere Information Server.

Scenario A: Basic client-server configuration:

Scenario A is an installation on a single computer, plus client workstations. There are no provisions for high availability. This configuration can be used quite successfully for deployment systems.

A small retail chain wants to integrate its forecasting, distribution, replenishment, and inventory management processes. As part of the integration, they want to migrate financial reporting data from several systems to a single system of record. They decide to create two separate IBM InfoSphere Information Server installations: a small development system where they build and test projects and jobs, and a larger production system (see "Scenario B: Isolated engine tier and parallel engine configuration"). For the development system, a simple topology can be used where a single server with a storage mechanism (Host1 in the following diagram) services five client workstations.

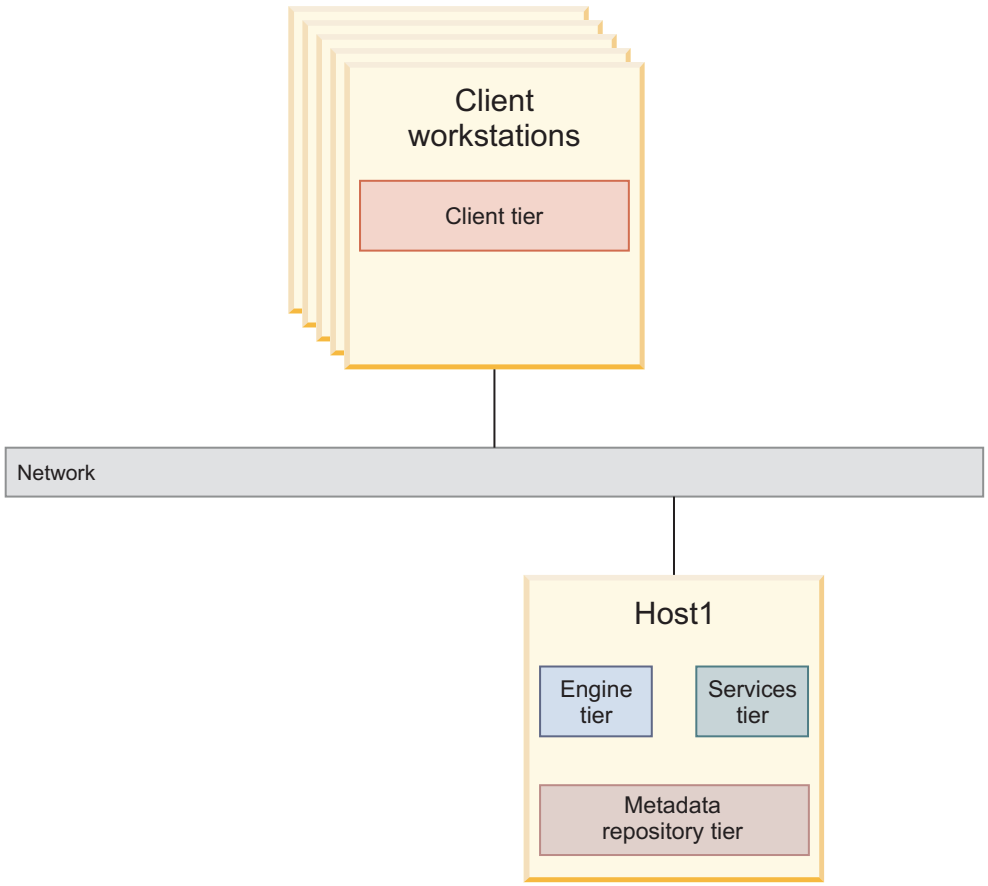


Figure 35. Scenario A configuration

Scenario B: Isolated engine tier and parallel engine configuration:

Scenario B is an installation with a computer that is dedicated to the IBM InfoSphere Information Server engine.

The retail chain that built the Scenario A development system also wants to build a larger production system. They are expecting a relatively high volume of job throughput. They decide to create a configuration in which the engine tier is built on a dedicated server (Host2 in the following diagram). To provide additional

capacity, they set up a separate computer (Host3) to serve as a parallel processing compute node. The repository tier and services tier share a server (Host1).

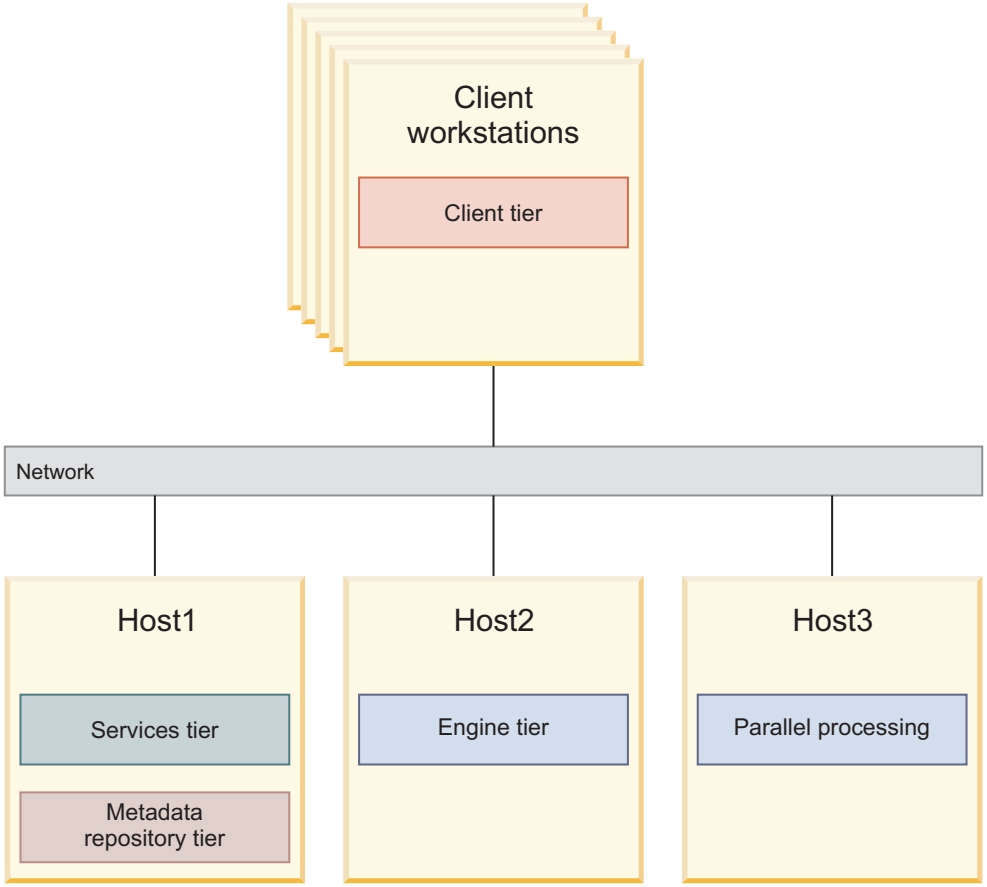


Figure 36. Scenario B configuration

Scenario C: Two-server active-passive high availability configuration:

Scenario C is a configuration that includes two computers that can fail over to one another in case of failure.

An industrial company has purchased IBM InfoSphere Information Server and plans to use IBM InfoSphere DataStage, IBM InfoSphere QualityStage, and IBM InfoSphere Information Analyzer for data integration. They expect to have three or four developers for the project.

Even if they are a small company that does not expect much volume on the production system, the company cannot afford much downtime of the system. The system must be highly-available. To minimize hardware costs, they decide to implement a two-server active-passive topology for this system. All components are to run on one server and fail over to the other server if a hardware, network, or operating system failure occurs, or if an administrator forces a failover for maintenance purposes. While the failover is occurring, the system is nonoperational.

They decide to create two separate IBM InfoSphere Information Server installations:

- A small development system where projects and jobs are built and tested. The configuration is identical to “Scenario A: Basic client-server configuration” on page 59.
- A more highly available production system, using the two-server active-passive configuration with a storage mechanism.

Several client workstations will access both systems.

To increase processing throughput, the passive node will be used as a parallel engine compute node while the primary node is active.

The following diagram illustrates this topology.

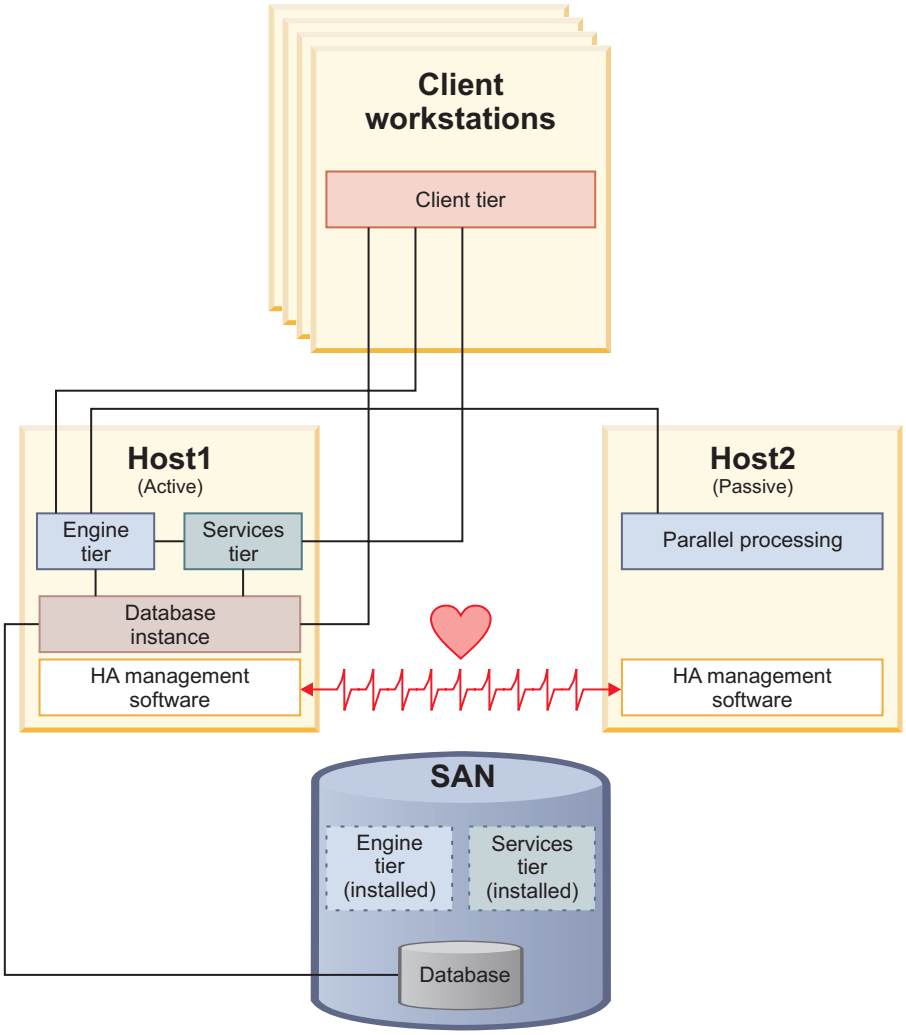


Figure 37. Scenario C configuration

Scenario D: Clustered metadata repository tier and services tier configuration:

Scenario D is a configuration that includes IBM WebSphere Application Server clustering and IBM DB2 clustering for high availability.

A large company uses both IBM WebSphere Application Server Network Deployment (ND) and DB2 for other applications. They recently purchased IBM InfoSphere Information Server. They plan to use the full suite of components for future data integration projects.

The company has existing WebSphere Application Server and IBM DB2 installations. They are each managed within different support groups within the company IT department. Each group understands how to provide high availability solutions for the component under their supervision.

The large number of concurrent users requires a clustered WebSphere Application Server topology that permits almost no downtime. This topology also allows for future expansion to accommodate additional capacity. The WebSphere Application Server group plans to create a WebSphere Application Server cluster for InfoSphere Information Server. The DB2 group will provide clustering of the metadata repository database and IBM InfoSphere Information Analyzer database to minimize downtime.

The engine tier will be configured in an active-passive server cluster topology. The cluster management software will be IBM Tivoli System Automation for Multiplatforms. To increase processing throughput, the passive node will be used as a parallel engine node while the primary node is active.

They are using UNIX computers for the servers and Microsoft Windows workstations for the clients.

The following diagram illustrates the configuration that they will build.

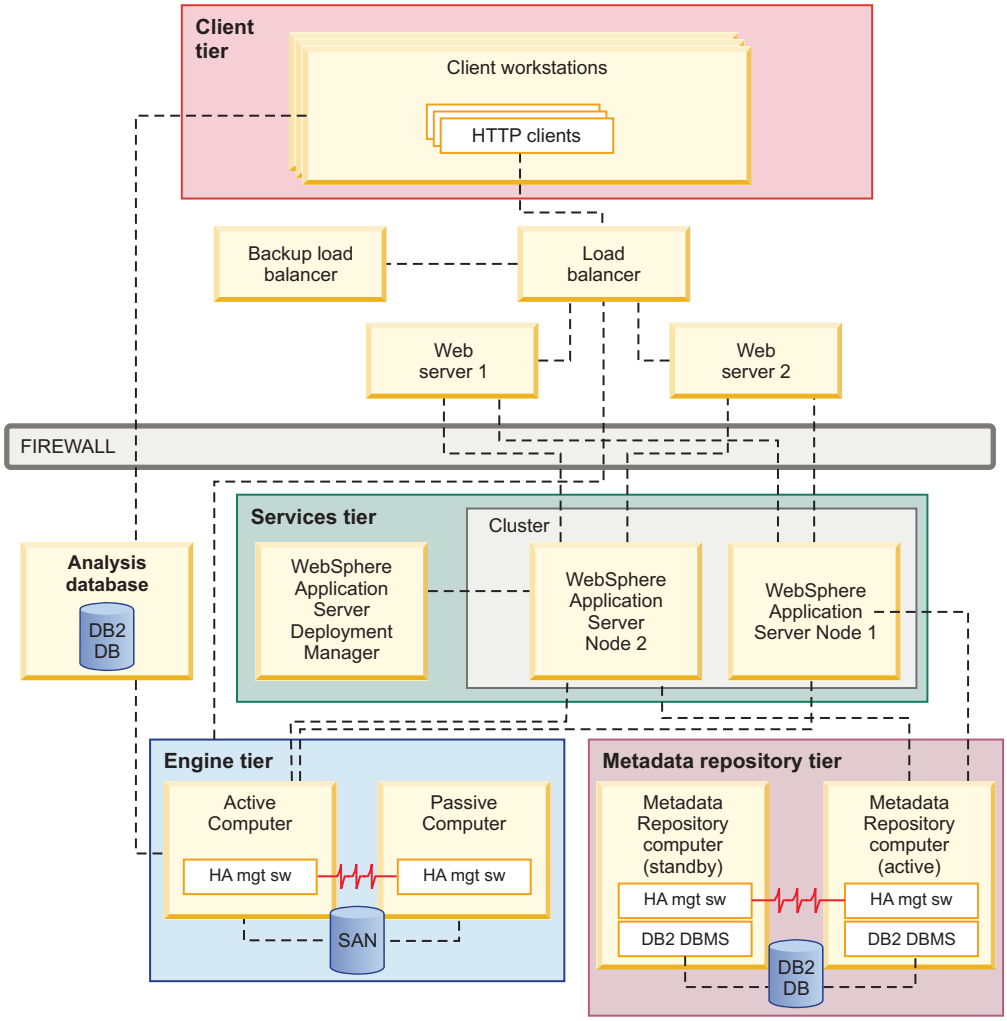


Figure 38. Scenario D configuration

Reviewing installation checklists

The following checklists include procedures for installing and configuring different implementations of IBM InfoSphere Information Server. Use these checklists to guide you through the installation process and to ensure that your environment is configured correctly.

Basic installation checklist

Use this checklist to complete a new installation of a configuration where there is no redundancy for high availability.

Complete the following steps:

1. Plan and build your topology
2. Install and configure each tier
3. Perform post-installation configuration tasks

To complete these steps, use the task lists in the following sections.

Plan and build your topology

To plan and build your topology, follow the steps in the following task list.

Table 7. Topology planning and building tasks

Description	For more information	Record your information
Review the system requirements.	"Reviewing system requirements and release notes" on page 94	
Learn about tiers and components.	Tiers and computers	
Design the topology of your installation.	"Designing an installation topology" on page 17	
Determine whether to preinstall your database system.	"Options for installing the database software" on page 13	
Determine whether to preinstall IBM WebSphere Application Server or preconfigure an existing installation of WebSphere Application Server.	"Options for installing the application server software" on page 14	
Build and test the hardware infrastructure.	Refer to the hardware documentation.	

Prepare to install the software

To prepare your computers to install the product software, complete the steps in the following task list.

Table 8. Software planning tasks

Description	For more information	Record your information
Obtain the InfoSphere Information Server software and any applicable fix packs.	"Obtaining IBM InfoSphere Information Server software" on page 193	
On all tier computers, verify name resolution, and open TCP ports for InfoSphere Information Server across any applicable firewalls.	"Configuring your network" on page 96	
AIX Update the kernel parameters and user limits on each computer.	Modifying kernel parameters and user limits	
Linux Update the kernel parameters and user limits on each computer.	Modifying kernel parameters and user limits	
Prepare each computer to run the InfoSphere Information Server installation program, and back up any critical data on the computers.	"Preparing to run the installation program" on page 197	

Install and configure each tier

For each tier, complete the computer preparation tasks. Then install IBM InfoSphere Information Server components, and perform post-installation tasks.

Install and configure the tiers in this order:

- Repository tier
- Services tier
- Engine tier
- Client tier

You can install more than one software tier on the same computer. In this case, you complete the planning and preparation tasks for the tiers on the computer. Then run the installation program only once for that computer, installing the software for the tiers. Then complete the post-installation tasks for the tiers on the computer.

For example, if the repository tier and services tier components are installed on computer Host1, and the engine tier and client tier are each installed on computers Host2 and Host3, complete the tasks in this order.

1. Repository tier and services tier preparation tasks on computer Host1.
2. Installation of repository tier and services tier components on computer Host1 by using the installation program.
3. Engine tier preparation tasks on computer Host2.
4. Installation of engine tier components on computer Host2, by using the installation program.
5. Client tier preparation tasks on computer Host3.
6. Installation of client tier components on computer Host3 by using the installation program.

Not all product modules require an engine tier. IBM InfoSphere FastTrack, IBM InfoSphere Information Governance Catalog, and IBM InfoSphere Glossary Anywhere do not require an engine tier.

Repository tier installation task list

To configure the repository tier, follow the steps in the following task list.

Table 9. Planning tasks: repository tier

Description	For more information	Record your information
Create the following administration accounts on the repository tier computer. Log in to each account to confirm that it functions properly: • Owner of the metadata repository • (If installing InfoSphere Information Analyzer) The owner of the analysis database • If installing the DB2 database system: **Windows** DB2 administrator user **AIX Linux** DB2 instance owner DB2 fenced user	"Setting up administrator and owner accounts" on page 107	
If you decide to preinstall DB2 or another supported database system, install the database system now.	"Pre-installing the metadata repository database" on page 114	

Table 9. Planning tasks: repository tier (continued)

Description	For more information	Record your information
If you preinstalled DB2, configure the databases within DB2 by using scripts. Alternatively, configure the databases when you run the InfoSphere Information Server installation program. If you preinstalled another database system, configure the databases by using scripts.	Configuring an existing installation of DB2 database software by using the installation program	
Run the InfoSphere Information Server installation program if necessary.	Run the InfoSphere Information Server installation program for this tier in either of the following situations: • To install the DB2 database system on the target computer and create the metadata repository and analysis databases within it. See Part 2, "Installing IBM InfoSphere Information Server software," on page 191. • To create the metadata repository and analysis database within a local preinstalled supported DB2 database system instance. See Configuring an existing installation of DB2 database software by using the installation program. If you preinstall the DB2 database system or another database system and create the databases within it by using scripts that are provided with the InfoSphere Information Server software, do not run the installation program for this tier.	

Services tier installation task list

To configure the services tier, follow the steps in the following task list.

Table 10. Installation tasks: services tier

Description	For more information	Record your information
Preinstall WebSphere Application Server, if you decide not to use the InfoSphere Information Server installation program to install it. When creating a profile during the installation, make sure to disable administrative security.	"Preinstalling WebSphere Application Server Network Deployment" on page 148	
Preconfigure an existing WebSphere Application Server installation if necessary.	"Setting up an existing non-clustered installation of WebSphere Application Server" on page 150	

Table 10. Installation tasks: services tier (continued)

Description	For more information	Record your information
Run the InfoSphere Information Server installation program.	Part 2, "Installing IBM InfoSphere Information Server software," on page 191	

Engine tier installation task list

If your installation requires an engine tier, follow the steps in the following task list to plan and configure the tier.

Note: If you are creating a configuration with more than one server engine, repeat these steps for each server engine.

Table 11. Installation tasks: engine tier

Description	For more information	Record your information
AIX For non-English locales, make sure that the locale and encoding are set correctly. The installation program uses the locale and encoding information of the computer to determine the appropriate locale to configure the installation to.	"Setting the locale and character encoding for the engine tier target computer" on page 163	
Linux For non-English locales, make sure that the locale and encoding are set correctly.	"Setting the locale and character encoding for the engine tier target computer" on page 163	
Create an administrator account for the InfoSphere Information Server engine. Log in to the account to activate the password.	"Setting up administrator and owner accounts" on page 107	
Run the InfoSphere Information Server installation program.	Part 2, "Installing IBM InfoSphere Information Server software," on page 191	

Client tier installation task list

To configure your client tier, follow the steps in the following task list.

Note: If you are creating a configuration with more than one client computer, repeat these steps for each client computer.

Table 12. Installation tasks: client tier

Description	For more information	Record your information
Run the installation program. In the Tier Selection page, choose **Client** only.	Part 2, "Installing IBM InfoSphere Information Server software," on page 191	

Test the installation

After you install the software, complete the following tasks to test and validate your installation.

Table 13. Testing the installation tasks

Description	For more information	Record your information
Test the IBM InfoSphere Information Services Director installation (if installed).	"Testing the IBM InfoSphere Information Services Director installation" on page 254	
Run the IBM Support Assistant Lite for InfoSphere Information Server tool to verify the installation.	"Running the ISALite tool to verify the installation and troubleshoot" on page 251	

Perform post-installation configuration tasks

After you run the InfoSphere Information Server installation program for each tier, perform these tasks:

Table 14. Post-installation configuration tasks (all tiers)

Description	For more information	Record your information
Windows If necessary, set up explicit firewall exceptions for client-side executable files that require network access.	"Setting up firewall exceptions" on page 306	
Configure security within InfoSphere Information Server: 1. Choose a user registry and configure it for InfoSphere Information Server. 2. Create users and groups in the user registry. 3. Assign security roles to users and groups. 4. Configure InfoSphere Information Server engine security. Map credentials if necessary. 5. Assign project roles if required by the installed suite components. Optionally, also: • Configure IBM WebSphere Application Server for non-root administration. • Configure InfoSphere Information Server agents for non-root administration. • Configure the Auditing service.	"Configuring security for IBM InfoSphere Information Server" on page 279	

Table 14. Post-installation configuration tasks (all tiers) (continued)

Description	For more information	Record your information
Configure IBM InfoSphere DataStage and IBM InfoSphere QualityStage (if installed): 1. If you installed globalization support (NLS), ensure that IBM InfoSphere DataStage users have the correct settings specified for localization. 2. Configure access to ODBC data sources. 3. For systems that have large numbers of simultaneously connected users or large numbers of simultaneously running jobs, tune the server engine. Optionally, also: • Migrate jobs that you created in an earlier version of InfoSphere DataStage to the current version. • Configure the IBM InfoSphere QualityStage Match Designer. • Set up the application server to run SQA or investigation reports. • Test the installation.	"Configuring IBM InfoSphere DataStage and IBM InfoSphere QualityStage" on page 337	
Configure IBM InfoSphere Information Analyzer (if installed).	Configuring IBM InfoSphere Information Analyzer	
Configure IBM InfoSphere Information Governance Catalog (if installed).	"Configuring IBM InfoSphere Information Governance Catalog" on page 349	

Basic isolated engine tier and parallel engine installation checklist

Use this checklist to complete a new installation of a configuration in which the engine tier is built on a dedicated server with a separate computer to serve as a parallel processing compute node.

Complete the following steps:

1. Plan and build your topology
2. Prepare to install the software
3. Install and configure the repository tier and the services tier on one computer
4. Install and configure the engine tier on a separate computer
5. Install and configure the client tier on each client workstation
6. Test the installation
7. Configure the parallel processing engine
8. Perform post-installation configuration tasks

To complete these steps, use the task lists in the following sections.

Plan and build your topology

To plan and build your topology, follow the steps in the following task list.

Table 15. Topology planning and building tasks

Description	For more information	Record your information
Review the system requirements.	"Reviewing system requirements and release notes" on page 94	
Learn about tiers and components.	Tiers and computers	
Design the topology of your installation.	"Designing an installation topology" on page 17	
Determine whether to preinstall your database system.	"Options for installing the database software" on page 13	
Determine whether to preinstall IBM WebSphere Application Server or preconfigure an existing installation of WebSphere Application Server.	"Options for installing the application server software" on page 14	
Build and test the hardware infrastructure.	Refer to the hardware documentation.	

Prepare to install the software

To prepare your computers to install the product software, complete the steps in the following task list.

Table 16. Software planning tasks

Description	For more information	Record your information
Obtain the InfoSphere Information Server software and any applicable fix packs.	"Obtaining IBM InfoSphere Information Server software" on page 193	
On all tier computers, verify name resolution, and open TCP ports for InfoSphere Information Server across any applicable firewalls.	"Configuring your network" on page 96	
AIX Update the kernel parameters and user limits on each computer.	Modifying kernel parameters and user limits	
Linux Update the kernel parameters and user limits on each computer.	Modifying kernel parameters and user limits	
Prepare each computer to run the InfoSphere Information Server installation program, and back up any critical data on the computers.	"Preparing to run the installation program" on page 197	

Install and configure the repository tier, the services tier, and the client tier

For each tier, complete the computer preparation tasks. Then install IBM InfoSphere Information Server components, and perform post-installation tasks.

Install and configure the tiers in this order:

- Repository tier
- Services tier
- Engine tier
- Client tier
- Parallel processing engine

Complete the planning and preparation tasks for the repository tier and services tier on a single computer. Then, run the installation program on that computer to install both of these tiers. You can then perform the planning, preparation, and installation tasks for the engine tier on a separate computer. Then complete the post-installation tasks for the tiers on the computer.

For example, if the repository tier and services tier components are installed on computer Host1, the engine tier is installed on computer Host2, the client tier is installed on separate client computers, and the parallel processing engine is installed on Host3, complete the tasks in the following order.

1. Repository tier and services tier preparation tasks on computer Host1.
2. Installation of repository tier and services tier components on computer Host1 by using the installation program.
3. Engine tier preparation tasks on computer Host2.
4. Installation of engine tier components on computer Host2 by using the installation program.
5. Client tier preparation tasks on each client computer.
6. Installation of client tier components on each client computer by using the installation program.
7. Parallel processing engine preparation tasks on computer Host 3.
8. Installation of parallel processing engine on Host 3 by using the installation program.

Not all product modules require an engine tier. IBM InfoSphere FastTrack, IBM InfoSphere Information Governance Catalog, and IBM InfoSphere Glossary Anywhere do not require an engine tier.

Repository tier installation task list

Complete the steps in the following task list to prepare and install the repository tier.

Table 17. Planning tasks: repository tier

Description	For more information	Record your information
Create the following administration accounts on the repository tier computer. Log in to each account to confirm that it functions properly: • Owner of the metadata repository • (If installing InfoSphere Information Analyzer) The owner of the analysis database • If installing the DB2 database system: **Windows** DB2 administrator user **AIX Linux** DB2 instance owner DB2 fenced user	"Setting up administrator and owner accounts" on page 107	
If you decide to preinstall DB2 or another supported database system, install the database system now.	"Pre-installing the metadata repository database" on page 114	
If you preinstalled DB2, configure the databases within DB2 by using scripts. Alternatively, configure the databases when you run the InfoSphere Information Server installation program. If you preinstalled another database system, configure the databases by using scripts.	Configuring an existing installation of DB2 database software by using the installation program	

Table 17. Planning tasks: repository tier (continued)

Description	For more information	Record your information
Run the InfoSphere Information Server installation program if necessary.	Run the InfoSphere Information Server installation program for this tier in either of the following situations: • To install the DB2 database system on the target computer and create the metadata repository and analysis databases within it. See Part 2, "Installing IBM InfoSphere Information Server software," on page 191. • To create the metadata repository and analysis database within a local preinstalled supported DB2 database system instance. See Configuring an existing installation of DB2 database software by using the installation program. If you preinstall the DB2 database system or another database system and create the databases within it by using scripts that are provided with the InfoSphere Information Server software, do not run the installation program for this tier.	

Services tier installation task list

To configure the services tier, complete the steps in the following task list.

Table 18. Installation tasks: services tier

Description	For more information	Record your information
Preinstall WebSphere Application Server, if you decide not to use the InfoSphere Information Server installation program to install it. When creating a profile during the installation, make sure to disable administrative security.	"Preinstalling WebSphere Application Server Network Deployment" on page 148	
Preconfigure an existing WebSphere Application Server installation if necessary.	"Setting up an existing non-clustered installation of WebSphere Application Server" on page 150	
Run the InfoSphere Information Server installation program.	Part 2, "Installing IBM InfoSphere Information Server software," on page 191	

Install the engine tier on a separate computer

Complete the following task list to plan, configure, and install the engine tier. Do not create a parallel engine configuration yet. Parallel engine configuration is completed at the end of the installation procedure.

Note: If you are creating a configuration with more than one server engine, repeat these steps for each server engine.

Table 19. Installation tasks: engine tier

Description	For more information	Record your information
AIX For non-English locales, make sure that the locale and encoding are set correctly.	"Setting the locale and character encoding for the engine tier target computer" on page 163	
Linux For non-English locales, make sure that the locale and encoding are set correctly.	"Setting the locale and character encoding for the engine tier target computer" on page 163	
Create an administrator account for the InfoSphere Information Server engine. Log in to the account to activate the password.	"Setting up administrator and owner accounts" on page 107	
Run the InfoSphere Information Server installation program.	Part 2, "Installing IBM InfoSphere Information Server software," on page 191	

Install the client tier on each client workstation

To configure your client tier, follow the steps in the following task list.

Note: If you are creating a configuration with more than one client computer, repeat these steps for each client computer.

Table 20. Installation tasks: client tier

Description	For more information	Record your information
Run the installation program. In the Tier Selection page, choose **Client** only.	Part 2, "Installing IBM InfoSphere Information Server software," on page 191	

Test the installation

After you install the software, complete the following tasks to test and validate your installation.

Table 21. Testing the installation tasks

Description	For more information	Record your information
Test the IBM InfoSphere Information Services Director installation (if installed).	"Testing the IBM InfoSphere Information Services Director installation" on page 254	
Run the IBM Support Assistant Lite for InfoSphere Information Server tool to verify the installation.	"Running the ISALite tool to verify the installation and troubleshoot" on page 251	

Configure the parallel processing engine

Configure the parallel processing engine for IBM InfoSphere DataStage and IBM InfoSphere QualityStage.

Table 22. Post-installation configuration tasks (all tiers)

Description	For more information	Record your information
• Add computers to share engine processing • Set up a C++ compiler • Set up environment variables for the parallel engine • Configure the parallel engine • Configure database connectivity for the parallel engine	Chapter 7, "Configuring a parallel processing environment," on page 309	

Perform post-installation configuration tasks

After you run the InfoSphere Information Server installation program for each tier, perform these tasks:

Table 23. Post-installation configuration tasks (all tiers)

Description	For more information	Record your information
Windows If necessary, set up explicit firewall exceptions for client-side executable files that require network access.	"Setting up firewall exceptions" on page 306	
Configure security within InfoSphere Information Server: 1. Choose a user registry and configure it for InfoSphere Information Server. 2. Create users and groups in the user registry. 3. Assign security roles to users and groups. 4. Configure InfoSphere Information Server engine security. Map credentials if necessary. 5. Assign project roles if required by the installed suite components. Optionally, also: • Configure IBM WebSphere Application Server for non-root administration. • Configure InfoSphere Information Server agents for non-root administration. • Configure the Auditing service.	"Configuring security for IBM InfoSphere Information Server" on page 279	

Table 23. Post-installation configuration tasks (all tiers) (continued)

Description	For more information	Record your information
Configure IBM InfoSphere DataStage and IBM InfoSphere QualityStage (if installed): 1. If you installed globalization support (NLS), ensure that IBM InfoSphere DataStage users have the correct settings specified for localization. 2. Configure access to ODBC data sources. 3. For systems that have large numbers of simultaneously connected users or large numbers of simultaneously running jobs, tune the server engine. Optionally, also: • Migrate jobs that you created in an earlier version of InfoSphere DataStage to the current version. • Configure the IBM InfoSphere QualityStage Match Designer. • Set up the application server to run SQA or investigation reports. • Test the installation.	"Configuring IBM InfoSphere DataStage and IBM InfoSphere QualityStage" on page 337	
Configure IBM InfoSphere Information Analyzer (if installed).	Configuring IBM InfoSphere Information Analyzer	
Configure IBM InfoSphere Information Governance Catalog (if installed).	"Configuring IBM InfoSphere Information Governance Catalog" on page 349	

High availability (active-passive configuration) installation checklist

Use this checklist to complete a new installation of a two-server active-passive configuration where all components can fail over from one server to the other.

Complete the following steps:

1. Plan and build your topology
2. Prepare to install the software
3. Install and configure high availability cluster management software
4. Install and configure the metadata repository tier, services tier, and engine tier within the two-server configuration
5. Install the client tier
6. Test the installation
7. Perform postinstallation configuration tasks

To complete these steps, use the task lists in the following sections.

Plan and build your topology

To plan and build your topology, follow the steps in the following task list.

Table 24. Topology planning and building tasks

Description	For more information	Record your information
Review the system requirements.	"Reviewing system requirements and release notes" on page 94	
Learn about tiers and components.	Tiers and computers	
Learn about the topology of the active-passive high availability configuration.	"Active-passive topology" on page 33	
Determine whether to preinstall your database system.	"Options for installing the database software" on page 13	
Determine whether to preinstall IBM WebSphere Application Server or preconfigure an existing installation of WebSphere Application Server.	"Options for installing the application server software" on page 14	

Prepare to install the software

To prepare your computers to install the product software, complete the steps in the following task list.

Table 25. Software planning tasks

Description	For more information	Record your information
Configure the servers, the network, and the storage device.	**AIX** "Setting up the servers, network, and storage (Tivoli System Automation for Multiplatforms example)" on page 172 **Linux** "Setting up the servers, network, and storage (Tivoli System Automation for Multiplatforms example)" on page 172 **Windows** "Setting up the servers, network, and storage (Microsoft example)" on page 186	
Obtain the InfoSphere Information Server software and any applicable fix packs.	"Obtaining IBM InfoSphere Information Server software" on page 193	
On both computers, verify name resolution, and open TCP ports for InfoSphere Information Server across any applicable firewalls.	"Configuring your network" on page 96	
AIX Update the kernel parameters and user limits on both computers.	Modifying kernel parameters and user limits (AIX)	

Table 25. Software planning tasks (continued)

Description	For more information	Record your information
Linux Update the kernel parameters and user limits on both computers.	Modifying kernel parameters and user limits (Linux)	
Prepare each tier computer to run the InfoSphere Information Server installation program, and back up any critical data on the computers.	"Preparing to run the installation program" on page 197	

Preparing and installing the tiers

The servers will accommodate the metadata repository tier, services tier, and engine tier. For this reason, run the IBM InfoSphere Information Server installation program only once to install the software on SAN storage. Then run the installation program on each client tier computer to install the client tier software.

To manage the cluster and enable failover, install high availability software such as IBM Tivoli System Automation for Multiplatforms.

InfoSphere Information Server metadata repository, services, and engine tier software installation task list

To prepare your tier software for installing the product, follow the steps in the following task list.

Table 26. Preparation tasks: metadata repository, services, and engine tiers

Description	For more information	Record your information
Acquire and install high availability cluster management software on both engine tier servers, and set up the cluster.	Refer to the documentation for your high availability cluster management software. For example configurations, see the following topics: **AIX** "Installing the high availability software (Tivoli System Automation for Multiplatforms example)" on page 173 **Linux** "Installing the high availability software (Tivoli System Automation for Multiplatforms example)" on page 173 **Windows** "Setting up the cluster (Microsoft example)" on page 186	

Table 26. Preparation tasks: metadata repository, services, and engine tiers (continued)

Description	For more information	Record your information
Set the service IP address.	Refer to the documentation for your high availability cluster management software. For example configurations, see the following topics: **AIX** "Setting up the virtual IP address (Tivoli System Automation for Multiplatforms example)" on page 174 **Linux** "Setting up the virtual IP address (Tivoli System Automation for Multiplatforms example)" on page 174 **Windows** "Setting up the virtual IP address (Microsoft example)" on page 187	
AIX Linux Define and start a resource group.	Refer to the documentation for your high availability cluster management software. See "Defining and starting the resource group (Tivoli System Automation for Multiplatforms example)" on page 175.	
AIX Linux For non-English locales, make sure that the locale and encoding are set correctly.	Setting the locale and character encoding for the engine tier target computer	
Set up administrator and component owner accounts on the shared storage device.	"Setting up administrator and owner accounts" on page 107	
If you decide to preinstall DB2 or another supported database system, install the database system now.	"Pre-installing the metadata repository database" on page 114	
If you preinstalled DB2, configure the databases within DB2 by using scripts. Alternatively, configure the databases when you run the InfoSphere Information Server installation program. If you preinstalled another database system, configure the databases by using scripts.	Configuring an existing installation of DB2 database software by using the installation program	

Table 26. Preparation tasks: metadata repository, services, and engine tiers (continued)

Description	For more information	Record your information
Preinstall WebSphere Application Server Network Deployment on the shared storage device, if you decide not to use the InfoSphere Information Server installation program to install it. When creating a profile during the installation, make sure to disable administrative security.	"Preinstalling WebSphere Application Server Network Deployment" on page 148	
Preconfigure the IBM InfoSphere QualityStage Match Designer database.	"Creating the IBM InfoSphere QualityStage Match Designer database" on page 143	
Run the installation program: • In the High Availability Server Cluster Configuration page, click **Server cluster configuration**. In the **Virtual host name** field, enter the virtual host name that was created when the HA software was configured. • Whenever prompted for an installation directory, specify a path on the shared storage device.	For example configurations, see the following topics: **AIX Linux** "Installing IBM InfoSphere Information Server in the cluster (Tivoli System Automation for Multiplatforms example)" on page 178 **Windows** "Installing IBM InfoSphere Information Server in the cluster (Microsoft example)" on page 187	
Set up the resource agent scripts.	Go to the documentation for your high availability cluster management software. For example configurations, see the following topics: **AIX Linux** "Setting up the high availability control scripts (Tivoli System Automation for Multiplatforms example)" on page 179 **Windows** "Setting up the engine and services tier resources for high availability (Microsoft example)" on page 188	
Test the high availability configuration.	Simulate active cluster node failures and network failures and observe the results. For an example based on Tivoli System Automation for Multiplatforms, see "Testing the configuration (Tivoli System Automation for Multiplatforms example)" on page 180.	

Client tier installation task list

To configure your client tier, follow the steps in the following task list.

Note: If you are creating a configuration with more than one client computer, repeat these steps for each client computer.

Table 27. Installation tasks: client tier

Description	For more information	Record your information
Run the installation program. In the Tier Selection page, choose **Client** only.	Chapter 4, "Installing IBM InfoSphere Information Server software," on page 201	

Test the installation

After you install the software, complete the following tasks to test and validate your installation.

Table 28. Testing the installation tasks

Description	For more information	Record your information
Test the IBM InfoSphere Information Services Director installation (if installed).	"Testing the IBM InfoSphere Information Services Director installation" on page 254	
Run the IBM Support Assistant Lite for InfoSphere Information Server tool to verify the installation.	See "Running the ISALite tool to verify the installation and troubleshoot" on page 251	
Test the engine tier failover function.	Go to the documentation for your high availability cluster management software. For example configurations, see the following topics: **AIX Linux** "Testing the configuration (Tivoli System Automation for Multiplatforms example)" on page 180 **Windows** "Testing the configuration (Microsoft example)" on page 189	

Perform post-installation configuration tasks

After you run the InfoSphere Information Server installation program for each tier, perform these tasks:

Table 29. Postinstallation configuration tasks (all tiers)

Description	For more information	Record your information
Windows If necessary, set up explicit firewall exceptions for client-side executable files that require network access.	"Setting up firewall exceptions" on page 306	

Table 29. Postinstallation configuration tasks (all tiers) (continued)

Description	For more information	Record your information
Configure security within InfoSphere Information Server: 1. Choose a user registry and configure it for InfoSphere Information Server. 2. Create users and groups in the user registry. 3. Assign security roles to users and groups. 4. Configure InfoSphere Information Server engine security. Map credentials if necessary. 5. Assign project roles if required by the installed suite components. Optionally, also: • Configure IBM WebSphere Application Server for non-root administration. • Configure InfoSphere Information Server agents for non-root administration. • Configure the Auditing service.	"Configuring security for IBM InfoSphere Information Server" on page 279	
Configure IBM InfoSphere DataStage and IBM InfoSphere QualityStage (if installed): 1. If you installed globalization support (NLS), ensure that IBM InfoSphere DataStage users have the correct settings specified for localization. 2. Configure access to ODBC data resources. 3. **Windows**: For systems that have large numbers of simultaneously connected users or large numbers of simultaneously running jobs, tune the server engine. Optional: • Configure the parallel processing environment. • Migrate jobs that you created in an earlier version of InfoSphere DataStage to the current version. • Configure the IBM InfoSphere QualityStage Match Designer. • Set up the application server to run SQA or investigation reports. • Test the installation.	"Configuring IBM InfoSphere DataStage and IBM InfoSphere QualityStage" on page 337	
Configure IBM InfoSphere Information Analyzer (if installed).	Configuring IBM InfoSphere Information Analyzer	

Table 29. Postinstallation configuration tasks (all tiers) (continued)

Description	For more information	Record your information
Configure IBM InfoSphere Information Governance Catalog (if installed).	"Configuring IBM InfoSphere Information Governance Catalog" on page 349	

High availability (clustered configuration) installation checklist

Use this checklist to complete a new installation where the metadata repository database system and services tier (IBM WebSphere Application Server) are set up in a clustered topology. The engine tier is set up in an active-passive high availability configuration.

Complete the following steps:

1. Plan the topology
2. Prepare the database cluster
3. Prepare the IBM InfoSphere Information Server databases within the cluster
4. Prepare the WebSphere Application Server cluster
5. Run the InfoSphere Information Server installation program to install InfoSphere Information Server software for the services tier and cause it to communicate with the database cluster
6. Configure WebSphere Application Server to communicate with the database cluster
7. Install and configure high-availability cluster management software on the engine tier servers
8. Install and configure the engine tier
9. Install and configure the client tier
10. Test the installation
11. Configure the cluster
12. Perform postinstallation configuration tasks

If you are setting up an IBM DB2 high-availability disaster recovery (HADR) topology, configure HADR after you install InfoSphere Information Server. See "Setting up the metadata repository database in an IBM DB2 database HADR configuration" on page 129.

To complete these steps, use the task lists in the following sections.

Plan and build your topology

To plan and build your topology, follow the steps in the following task list.

Table 30. Topology planning and building tasks

Description	For more information	Record your information
Review the system requirements.	"Reviewing system requirements and release notes" on page 94	
Learn about tiers and components.	Tiers and computers	
Design the topology of your installation.	"High availability topologies" on page 30	
Determine whether to preinstall your database system.	"Options for installing the database software" on page 13	

Prepare to install the software

To prepare your computers to install the product software, complete the steps in the following task list.

Table 31. Software planning tasks

Description	For more information	Record your information
Determine your cluster architecture.	"Options for installing on a cluster" on page 41	
Build and test the hardware infrastructure.	Metadata repository tier: • DB2 clustered installation: "Preparing the servers, file system, and storage" on page 119 • DB2 HADR installation: "Preparing the servers, storage, and network" on page 130 Services tier: • "Options for installing on a cluster" on page 41 Engine tier: **AIX Linux** "Setting up the servers, network, and storage (Tivoli System Automation for Multiplatforms example)" on page 172 **Windows** "Setting up the servers, network, and storage (Microsoft example)" on page 186	
Configure the servers, the network, and the storage device.	**AIX Linux** "Setting up the servers, network, and storage (Tivoli System Automation for Multiplatforms example)" on page 172 **Windows** "Setting up the servers, network, and storage (Microsoft example)" on page 186	
Obtain the InfoSphere Information Server software and any applicable fix packs.	"Obtaining IBM InfoSphere Information Server software" on page 193	
On both computers, verify name resolution, and open TCP ports for InfoSphere Information Server across any applicable firewalls.	"Configuring your network" on page 96	

Table 31. Software planning tasks (continued)

Description	For more information	Record your information
AIX Linux Update the kernel parameters and user limits on both computers.	Modifying kernel parameters and user limits	
Prepare each tier computer to run the InfoSphere Information Server installation program, and back up any critical data on the computers.	"Preparing to run the installation program" on page 197	

Install and configure the tiers

For each tier, complete planning tasks and preinstallation tasks. Then install IBM InfoSphere Information Server components, and perform postinstallation tasks.

Install and configure the tiers in this order:

- Metadata repository tier
- Services tier
- Engine tier
- Client tier

Metadata repository tier installation task list

To configure your metadata repository tier, follow the steps in the following task list.

Table 32. Installation tasks: metadata repository tier

Description	For more information	Record your information
Configure two servers with identical internal storage, memory, and processing power. Designate one server as "primary" and the other as "standby". These servers will support the metadata repository database and analysis database in a nonpartitioned topology. Configure a high-speed network between the servers, and a separate, isolated network to carry a heartbeat signal for the high availability software.	"Preparing the servers, file system, and storage" on page 119	
(Optional, for an IBM DB2 clustered topology.) Install high-availability (HA) software to monitor the health of the servers and to manage failover.	"Installing high availability software on DB2 database system servers" on page 121	

Table 32. Installation tasks: metadata repository tier (continued)

Description	For more information	Record your information
Create the following administration accounts on each computer that will host the database system. The accounts on each computer must be identical. Log in to each account to confirm that it functions properly: • Owner of the metadata repository • (If installing IBM InfoSphere Information Analyzer) The owner of the analysis database • If installing the DB2 database system: **Windows** DB2 Administrator user **AIX Linux** DB2 instance owner DB2 fenced user	"Setting up administrator and owner accounts" on page 107	
Install the database system, create the database cluster, and preconfigure the database by using scripts that are supplied with the installation media.	• DB2 clustered installation: "Installing the DB2 database system and creating the cluster and metadata repository database" on page 122 • DB2 HADR installation: "Installing the DB2 database system and creating the metadata repository database" on page 131	

Note: Do not run the InfoSphere Information Server installation program at this time. Run the installation program when you configure the services tier.

Services tier installation task list

Table 33. Installation tasks: services tier

Description	For more information	Record your information
Install WebSphere Application Server Network Deployment on all the computers that will host the services tier. This installation includes the deployment manager system and all computers that will host a managed node in the cluster. Also install any required WebSphere Application Server fix packs and patches.	"Installing IBM WebSphere Application Server Network Deployment for a cluster installation" on page 152	
On the computer that hosts the deployment manager, create a deployment manager profile that contains the deployment manager process. Make sure that Administrative Security is enabled.	"Creating a deployment manager profile" on page 153	

Table 33. Installation tasks: services tier (continued)

Description	For more information	Record your information
On the computer that hosts the managed node, create a custom profile. This profile will contain the node agent process and the managed server process that will be part of the cluster.	"Creating a custom profile" on page 154	
Use the WebSphere Application Server administrative console to create a cluster with one cluster member definition. **Note:** Start with a minimal cluster topology that contains a single node. After the installation procedure is complete, expand the cluster.	"Creating a cluster and cluster members for the application server" on page 155	
Install a front-end web server.	"Installing a front-end web server" on page 156	
(Optional) Install backup web servers, load balancers, firewalls, and other architectural components.	"Installing a front-end load balancer (optional)" on page 158	

Table 33. Installation tasks: services tier (continued)

Description	For more information	Record your information
Run the installation program on the computer where the WebSphere Application Server deployment manager is installed: 1. From the `is-suite` directory on the installation media, run the following command to start the installation program: `./setup -verbose` 2. Using a supported web browser, open the URL that the **./setup** command displays. 3. In the installation program, accept all defaults, except these items: a. In the Tier Selection page, choose **Services** only. b. In the Application Server page, choose **Use an existing WebSphere Application Server installation**. c. In the WebSphere Application Server Profile page, choose the Deployment Manager profile. d. In the WebSphere Application Server Cluster Information page, specify a cluster member or the front-end HTTP dispatcher for the WebSphere Application Server cluster. e. In the Database Management System Selection page, choose **Use existing database management system installation**. f. In the Database Connection page, specify the metadata repository tier connection and database information. 4. Complete the rest of the installation steps.	"Running the installation program (clustered installation)" on page 160	

Engine tier installation task list

To configure your engine tier, follow the steps in the following task list.

Note: If you are creating a topology with more than one server engine, repeat these steps for each server engine.

Table 34. Installation tasks: engine tier

Description	For more information	Record your information
Acquire and install high-availability cluster management software on both engine tier servers, and set up the cluster.	Go to the documentation for your high-availability cluster management software. For example configurations, see the following topics: **AIX Linux** "Installing the high availability software (Tivoli System Automation for Multiplatforms example)" on page 173 **Windows** "Setting up the cluster (Microsoft example)" on page 186	
Set the service IP address.	Refer to the documentation for your high availability cluster management software. For example configurations, see the following topics: **AIX Linux** "Setting up the virtual IP address (Tivoli System Automation for Multiplatforms example)" on page 174 **Windows** "Setting up the virtual IP address (Microsoft example)" on page 187	
AIX Linux Define and start a resource group.	Refer to the documentation for your high availability cluster management software. See "Defining and starting the resource group (Tivoli System Automation for Multiplatforms example)" on page 175.	
AIX Linux For non-English locales, make sure that the locale and encoding are set correctly.	"Setting the locale and character encoding for the engine tier target computer" on page 163	
Create an administrator account for the InfoSphere Information Server engine. Log in to the account to activate the password.	"Setting up administrator and owner accounts" on page 107	

Table 34. Installation tasks: engine tier (continued)

Description	For more information	Record your information
Run the installation program: • In the Tier Selection page, choose **Engine** only. • In the High Availability Server Cluster Configuration page, click **Server cluster configuration**. In the **Virtual host name** field, enter the virtual host name that was created when the HA software was configured. • Whenever prompted for an installation directory, specify a path on the shared storage device. • In the IBM InfoSphere Information Server Location page, specify a cluster member or front-end HTTP server for the WebSphere Application Server cluster. Enter the InfoSphere Information Server administrator account name and password.	For example configurations, see the following topics: **AIX Linux** "Installing IBM InfoSphere Information Server in the cluster (Tivoli System Automation for Multiplatforms example)" on page 178 **Windows** "Installing IBM InfoSphere Information Server in the cluster (Microsoft example)" on page 187	
Set up the resource agent scripts.	Go to the documentation for your high availability cluster management software. For example configurations, see the following topics: **AIX Linux** "Setting up the high availability control scripts (Tivoli System Automation for Multiplatforms example)" on page 179 **Windows** "Setting up the engine and services tier resources for high availability (Microsoft example)" on page 188	

Client tier installation task list

To configure your client tier, follow the steps in the following task list.

Note: If you are creating a topology with more than one client computer, repeat these steps for each client computer.

Table 35. Installation tasks: client tier

Description	For more information	Record your information
Run the installation program. In the Tier Selection page, choose **Client** only.	Chapter 4, "Installing IBM InfoSphere Information Server software," on page 201	

Test the installation

After you install the software, complete the following tasks to test and validate your installation.

Table 36. Testing the installation tasks

Description	For more information	Record your information
Test the IBM InfoSphere Information Services Director installation (if installed).	"Testing the IBM InfoSphere Information Services Director installation" on page 254	
Run the IBM Support Assistant Lite for InfoSphere Information Server tool to verify the installation.	See "Running the ISALite tool to verify the installation and troubleshoot" on page 251	
Test the engine tier failover function.	Go to the documentation for your high availability cluster management software. For example configurations, see the following topics: **AIX Linux** "Testing the configuration (Tivoli System Automation for Multiplatforms example)" on page 180 **Windows** "Testing the configuration (Microsoft example)" on page 189	

Configure the cluster

After you install the software and test the installation, complete the following tasks to configure the cluster.

Table 37. Configuring the cluster tasks

Description	For more information	Record your information
Windows (Optional) Configure the cluster to run as a service.	"Configuring a cluster to run as a service (Windows)" on page 159	
DB2 clustered and HADR installations: Configure the primary and secondary servers.	• DB2 cluster installation: "Configuring automatic client reroute" on page 124 • DB2 HADR installation: – "Configuring the HADR primary server" on page 136 – "Configuring the HADR standby server" on page 138	

Table 37. Configuring the cluster tasks (continued)

Description	For more information	Record your information
Configure WebSphere Application Server to work with the metadata repository tier cluster.	• DB2 clustered installation: "WebSphere Application Server Network Deployment: Specifying the alternate address and port for a DB2 clustered database" on page 125 • DB2 HADR installation: Specifying the alternate address and port for a DB2 HADR database to connect to IBM WebSphere Application Server	
DB2 HADR installation: Start HADR.	"Starting HADR" on page 140	

Perform post-installation configuration tasks

After you run the InfoSphere Information Server installation program for each tier, perform these tasks:

Table 38. Postinstallation configuration tasks (all tiers)

Description	For more information	Record your information
(Optional) Add additional WebSphere Application Server nodes or other components to the cluster.	"Adding a new cluster member" in the *IBM InfoSphere Information Server Administration Guide*	
To communicate the WebSphere Application Server cluster configuration to the front-end Web server, run the WebSphere Application Server administrative console to generate the `plugin-cfg.xml` file, and then manually copy it to each Web server. Then restart the Web server. If you are using IBM HTTP Server, you can automatically propagate the file to your Web servers by using the WebSphere Application Server administrative console.	"Propagating the plugin-cfg.xml file to the front-end Web server" in the *IBM InfoSphere Information Server Administration Guide*	

Table 38. Postinstallation configuration tasks (all tiers) (continued)

Description	For more information	Record your information
Configure security within InfoSphere Information Server: 1. Choose a user registry and configure it for InfoSphere Information Server. 2. Create users and groups in the user registry. 3. Assign security roles to users and groups. 4. Configure InfoSphere Information Server engine security. Map credentials if necessary. 5. Assign project roles if required by the installed suite components. Optionally, also: • ConfigureWebSphere Application Server for non-root administration. • Configure InfoSphere Information Server agents for non-root administration. • Configure the Auditing service.	"Configuring security for IBM InfoSphere Information Server" on page 279	
Configure IBM InfoSphere DataStage and IBM InfoSphere QualityStage (if installed): 1. If you installed globalization support (NLS), ensure that InfoSphere DataStage users have the correct settings specified for localization. 2. Configure access to ODBC data sources. 3. **Windows**: For systems that have large numbers of simultaneously connected users or large numbers of simultaneously running jobs, tune the server engine. Optionally, also: • Configure the parallel processing environment. • Migrate jobs that you created in an earlier version of InfoSphere DataStage to the current version. • Configure the InfoSphere QualityStage Match Designer. • Set up the application server to run SQA or investigation reports. • Test the installation.	"Configuring IBM InfoSphere DataStage and IBM InfoSphere QualityStage" on page 337	
Configure InfoSphere Information Analyzer (if installed).	Configuring IBM InfoSphere Information Analyzer	

Table 38. Postinstallation configuration tasks (all tiers) (continued)

Description	For more information	Record your information
Configure IBM InfoSphere Information Governance Catalog (if installed).	"Configuring IBM InfoSphere Information Governance Catalog" on page 349	

Setting up the tier computers

Complete the following preparation steps for each computer in your installation.

You install InfoSphere Information Server tier by tier. When you run the installation program, it prompts you to select the product modules to install, to specify which tier or tiers to install on the computer, and it installs the appropriate product module components in the tiers that you specify.

Although the installation program completes a prerequisite check before installing software, you must check some additional prerequisites identified by IBM Software Support. See the following technote for more information regarding prerequisites: www.ibm.com/support/docview.wss?uid=swg27019888.

Reviewing system requirements and release notes

Ensure that all computers meet the hardware and software requirements for the product modules that you want to install.

To obtain the most up-to-date information about the installation requirements, see these resources:

- The system requirements at publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/softwareReqsForProduct.html.
- The release notes at http://www.ibm.com/support/knowledgecenter/SSZJPZ_11.3.0/com.ibm.swg.im.iis.productization.iisinfsv.relinfo.doc/topics/iisihrinfo_infsv_rnote_v91.html.

Running ISA Lite to check prerequisites

To determine if your system is ready to install IBM InfoSphere Information Server, run the IBM Support Assistant Lite for InfoSphere Information Server (ISA Lite) tool. The Prerequisite Checker that is included with ISA Lite evaluates your system to determine whether required packages are installed, if kernel parameters are too low for a successful installation, or if potential problems exist with environment variables.

Before you begin

Download and install the ISA Lite tool from the support site (www.ibm.com/support/docview.wss?uid=swg24022700).

About this task

For a full explanation of how to complete the following steps, refer to the documentation that is available under the `\doc` folder of the ISA Lite installation.

Procedure

1. Run the tool in console mode to create a response file. You can use this response file to run the same script in the future without reentering all of your information.
2. Run the Prerequisites Checker to evaluate your system. You can run the Prerequisites Checker before starting a new installation of InfoSphere Information Server or before you start an upgrade to a new version.
3. View the generated reports to see the prerequisites that passed and failed. Correct any failed prerequisites and then run the tool again.

Directories that must be present on all tier computers before installation

The IBM InfoSphere Information Server installation program requires certain directories be present on all tier computers before you run it.

The temporary directories will contain temporary files that are required by the installation program. The files are not needed after installation and can be removed safely.

Temporary directories

You do not have to back up temporary directories.

The root user, the user who runs the installation program, and all other InfoSphere Information Server users must have read, write, and execute permissions to the temporary directories. Also, the sticky bit must be unset:

```
chmod -t temp_dir_path
```

Put the temporary directories on a file system that has high-speed read-write performance.

The following table provides operating system-specific information about the temporary directories.

Table 39. Temporary directories: operating system-specific information

Operating system	Default directory path	Additional requirements
AIX Linux	`/tmp` `/var/tmp`	The file system that contains the temporary directory must meet the system requirements for temporary space on its physical partition. Use the **df** command to see the amount of available disk space per file system. If the command results do not list the directory or its parent directory, use the amount on the root file system (/).
Windows	The directory specified by the TEMP environment variable. This directory is typically the `%USERPROFILE%\AppData\Local\Temp` directory, where `%USERPROFILE%` is the home directory of the user doing the installation.	The total number of characters in the temporary directory path and the path of the files that the installation program extracts into the path cannot exceed 256 characters. Long path names might cause the installation to fail.

Home directory for user IDs

The home directory contains existing user IDs and (by default) the IDs that the installation program creates.

The Administrator ID or the root user ID that performs the installation must be able to write to the home directory and all subdirectories.

Unlike temporary directories, these directories should be backed up regularly.

The following table provides operating system-specific information about the home directory.

Table 40. Home directory: operating system-specific information

Operating system	Default directory path	Additional requirements
AIX	`/home/`*`user_name`*	Owners must have at least 755 permissions for their directories and parent directories, and have read/write/execute access to all subdirectories. If the home directory is on a network file system (NFS), export the file system with the `no_root_squash` option enabled and the `nosetuid` option off.
Linux	`/home/`*`user_name`*	Owners must have at least 755 permissions for their directories and parent directories, and have read/write/execute access to all subdirectories. If the home directory is on a network file system (NFS), export the file system with the `no_root_squash` option enabled and the `nosetuid` option off.
Windows	`%USERPROFILE%`, typically `C:\Users\`*`user_name`*	None

Configuring your network

IBM InfoSphere Information Server exchanges high volumes of data across several network ports, which requires that you carefully configure the host names and ports for your network.

Before you begin

Ensure that a TCP/IP network is used for all communications in the installation.

About this task

For performance reasons, do not install separate tiers of InfoSphere Information Server across a wide area network (WAN) because of the data exchanges that are

required between them. Although clients can work over a WAN if the latency is low, when the network latency increases, it has a detrimental effect on the clients.

The port numbers in the following tables might vary from those that are used in your system because the installation program automatically computes the port value based on port availability. For example, for the DSRPC port, the installation program starts with port 31538 and increments by 1 until it locates an open port. Ensure that you select the correct ports in your system to open and make available. The port values that are used in your system are stored in the response file that the installation program generates. Before you install the product, check the Response File Review panel in the installation program for the location, name, and contents of your response file.

Procedure

1. Set up the name resolution.
   a. Verify that the computers in the installation can resolve all the other computers by the name that is returned by the `hostname` command.
   b. Ensure that each of the following tier computers can resolve the URI addresses of the computers that are listed for that tier.

Table 41. Tiers and the computers that they must connect to

Tier computer to connect from	Tier computers to connect to
Repository tier	Repository tier computers (in addition to the loopback address `127.0.0.1`)
Client tier	Services tier computers Engine tier computers Repository tier computers (For IBM InfoSphere Information Analyzer, if the IADB is hosted on the same computer as the metadata repository.)
Services tier	Services tier computers (in addition to the loopback address `127.0.0.1`) Engine tier computers Repository tier computers
Engine tier	Engine tier computers (in addition to the loopback address `127.0.0.1`) Services tier computers Repository tier computers (depending on the DSODB configuration)

   c. Ensure that only `localhost` is mapped to the loopback IP address. The IPv4 version of this address is `127.0.0.1` and the IPv6 version is `::1/128`. No other entries can map `localhost` to the IP address of the local computer. The following example shows how these entries are configured in the following files:

Operating system	File
AIX	`/etc/hosts`
Linux	`/etc/hosts`

Operating system	File
Windows	C:\Windows\system32\drivers\etc\hosts

IPv4:

```
127.0.0.1 localhost.localdomain localhost
xx.xx.xx.xx longhostname shorthostname
```

IPv6:

```
::1/128 localhost.localdomain localhost
xx:xx::xx:xx longhostname shorthostname
```

Where:

localdomain
The domain name.

xx.xx.xx.xx **or** *xx:xx::xx:xx*
The unique IP address for the computer. IPv4 or IPv6 format.

longhostname shorthostname
The long and short names for the computer.

2. For the repository tier, open and make TCP ports available through firewalls.

Table 42. Port numbers for the repository tier. These are the default port numbers for the database management systems. Choose the ports that are configured for the type of database management system that you are using for the specified uses that are applicable for your configuration.

Component	Default port numbers	Tiers that use the port
IBM DB2 database for the metadata repository	50000	Services
IBM DB2 database for the analysis database	50000	Services, engine, client

3. For the services tier, open and make TCP ports available through firewalls. Refer to the table that corresponds to your IBM WebSphere Application Server configuration.

   In addition to the listed ports below, refer to Port number settings.

Table 43. Port numbers for the services tier (basic configuration with IBM WebSphere Application Server Liberty Profile)

Component	Default port numbers for basic configuration	Tiers that use the port
InfoSphere Information Server web-based clients (HTTPS)	Liberty HTTPS port (9443)	Client **Note:** The Liberty HTTPS port is set during the installation of InfoSphere Information Server.

Table 44. Port numbers for the services tier (basic configuration with IBM WebSphere Application Server Network Deployment)

Component	Default port numbers for basic configuration	Tiers that use the port
InfoSphere Information Server web-based clients (HTTPS)	WC_defaulthost_secure (9443)	Client
IBM WebSphere Application Server administrative console (redirects to HTTPS)	WC_adminhost (9060)	Client
IBM WebSphere Application Server administrative console (HTTPS)	WC_adminhost_secure (9043)	Client

Table 44. Port numbers for the services tier (basic configuration with IBM WebSphere Application Server Network Deployment) (continued)

Component	Default port numbers for basic configuration	Tiers that use the port
IBM InfoSphere Information Services Director services with JMS bindings	SIB_ENDPOINT_ADDRESS (7276) SIB_ENDPOINT_SECURE_ ADDRESS (7286) SIB_MQ_ENDPOINT_ ADDRESS (5558) SIB_MQ_ENDPOINT_ SECURE_ADDRESS (5578)	Only clients of the Information Services Director services require access to these ports.
IBM InfoSphere Information Services Director services with EJB bindings	BOOTSTRAP_ADDRESS (2809) ORB_LISTENER_ADDRESS (9100)	Only clients of the Information Services Director services require access to these ports.

Table 45. Port numbers for the services tier (clustered configuration with IBM WebSphere Application Server Network Deployment)

Component	Default port numbers for clustered configuration	Tiers that use the port
InfoSphere Information Server web-based clients (HTTPS)	Firewall between the client and dispatcher (if any): usually 443. This value varies depending on the configuration of the front-end dispatcher. Firewall between the dispatcher and cluster members (if any): WC_defaulthost_secure of all cluster members (for example, 9443, 9444, and so on)	Client
IBM WebSphere Application Server administrative console (redirects to HTTPS)	WC_adminhost (9060)	Client
IBM WebSphere Application Server administrative console (HTTPS)	WC_adminhost_secure (9043)	Client
IBM InfoSphere Information Services Director services with JMS bindings	SIB_ENDPOINT_ADDRESS of all cluster members (7277, 7278, and so on) SIB_ENDPOINT_SECURE_ ADDRESS of all cluster members (7287, 7288, and so on) SIB_MQ_ENDPOINT_ ADDRESS of all cluster members (5559, 5560, and so on) SIB_MQ_ENDPOINT_ SECURE_ADDRESS of all cluster members (5579, 5580, and so on)	Only clients of the Information Services Director services require access to these ports.

Table 45. Port numbers for the services tier (clustered configuration with IBM WebSphere Application Server Network Deployment) (continued)

Component	Default port numbers for clustered configuration	Tiers that use the port
IBM InfoSphere Information Services Director services with EJB bindings	BOOTSTRAP_ADDRESS ORB_LISTENER_ADDRESS	Only clients of the Information Services Director services require access to these ports.

4. For the engine tier, open and make TCP ports available through firewalls.

Table 46. Port numbers for the engine tier

Component	Default port numbers	Tiers that use the port
InfoSphere Information Server ASB agent	31531, and a random port number greater than 1024 Instead of using a random port number, you can fix this to a specific port by specifying `agent.objectport=#` in the *`IS_install_path`*`/ASBNode/conf/agent.properties` file after you complete the installation. After designating a specific port, restart the ASB agent so that the change takes effect.	Services
IBM InfoSphere DataStage service	DSRPC (31538)	Services, client
Parallel job monitors	13400 (port 1) and 13401 (port 2) Access to port 1 is required only from the conductor node. Access to port 2 is required from the conductor node and the node where the InfoSphere Information Server engine is installed, if that node is different from the conductor node.	Engine
Parallel engine	APT_PM_STARTUP_PORT (multiple ports, uses a port number of 10000 or greater)	Engine
Parallel engine remote process startup	**Remote Shell and Secure Shell** 22 **Multiple nodes only** 514	Engine
Parallel engine (multiple nodes only)	APT_PLAYER_CONNECTION_PORT (multiple ports, uses a port number of 11000 or greater)	Engine

5. For the client tier, open and make TCP ports available through firewalls.

Table 47. Port numbers for the client tier

Component	Default port numbers	Tiers that use the port
IBM InfoSphere Metadata Asset Manager, when bridges are installed.	19443	Services

6. For the installation program, open and make TCP ports available through firewalls.

Table 48. Port numbers for the installation program

Component	Default port numbers	Tiers that use the port
The installation program	8443 (HTTPS)	Any web browser that connects to the computer running the installation program.
The SOAP Connector Port (Stand-alone environment)	SOAP_CONNECTOR_ADDRESS (8880)	Services **Note:** The SOAP Connector Port is used during suite install and most Patch and FixPack installations to the Services tier. It is only used from the Services tier computer to the WebSphere Application Server running locally on the same tier (The Services tier computer connects to itself).
The SOAP Connector Port for the deployment manager (Clustered environment)	SOAP_CONNECTOR_ADDRESS (8879)	Services **Note:** The SOAP Connector Port is used during suite install and most Patch and FixPack installations to the Services tier. In a clustered WebSphere Application Server environment, it also used for communicating between the deployment manager and remote nodes. It is only used from the Services tier computer to the WebSphere deployment manager running locally on the same tier (The Services tier computer connects to itself). The default value for SOAP_CONNECTOR_ADDRESS for a cluster member (not the deployment manager) is 8880.

7. For connecting to external data sources, see the vendor-supplied documentation for additional port assignments.

Modifying kernel parameters and user limits (AIX Linux)

IBM InfoSphere Information Server requires specific resources for each tier. Modify the kernel parameters and user limits, and restart the computer before you start the installation program.

The kernel parameter and user limit requirements are minimum requirements for InfoSphere Information Server. The optimal value for each parameter might be higher depending on your topology and the operational data load. Additional kernel parameter and user limit changes might be required. When more than one value is provided for your configuration, use the largest minimum value.

All changes must be made by a trained system administrator. Make a backup copy of the kernel parameters and user limits before you modify them. After you modify the kernel parameters and user limits with the suggested values, restart the computer, and then install InfoSphere Information Server.

Minimum kernel parameters

You might need to configure the kernel parameters and user limits for your application server. For more information, refer to the IBM WebSphere Application Server documentation.

Minimum kernel parameter values for IBM DB2 databases

The minimum metadata repository kernel parameter requirements apply only to DB2 database installations. Before you install DB2 as part of InfoSphere Information Server, ensure that your computer meets the following minimum requirements for kernel parameter for DB2:

- **AIX**: The default kernel parameter values are sufficient.
- **Linux**: The minimum Linux kernel parameter values are described in the following table:

Table 49. Minimum Linux kernel parameters and user limits for the metadata repository computer on DB2

Parameter	Value
MSGMAX	65535*
MSGMNB	65535*
MSGMNI	1024
RANDOMIZE_VA_SPACE	0
SEMMNI	1024
SEMMNS	256000
SEMMSL	250
SEMOPM	32
SHMALL	• For the default 4096 **PAGE_SIZE**, use 4294967296 • For a **PAGE_SIZE** larger than 4096, divide 17592186044416 by **PAGE_SIZE** (17592186044416/**PAGE_SIZE**)
SHMMAX	Set this value to be equal to 100% of the physical memory (in bytes) on your computer.
SHMMNI	4096

* Required by DB2.

When the DB2 kernel requirements are different from the InfoSphere Information Server kernel parameter requirements, use the largest minimum value. If you use a different database system for the metadata repository, consult the product documentation for that database system to find information about kernel parameter requirements.

Minimum AIX kernel parameters

Run the **smit** or **smitty** command to modify kernel parameters on AIX. The following table contains the minimum kernel parameters and user limits for the engine and services tiers.

Table 50. Minimum AIX kernel parameters and user limits (engine and services tiers)

Parameter	Engine and services tiers
MAXUPROC	4096 or unlimited
NOFILES	10240

Minimum Linux kernel parameters

The following table contains the minimum kernel parameters and user limits for the engine and services tiers.

Table 51. Minimum Linux kernel parameters and user limits (engine and services tiers)

Parameter	Engine and services tiers
MSGMAX	8192
MSGMNB	16384
MSGMNI	No requirement
SEMMNI	1024
SEMMNS	128000
SEMMSL	250
SEMOPM	32
SHMALL	2511724800 Use the smaller value of 837241600 if the following conditions apply: • The computer has only the minimum memory requirement of 4 GB. • The computer does not have the metadata repository with DB2 installed. • The page size is 4096.
SHMMAX	307200000
SHMMNI	2000
nproc	1000 or unlimited
nofile	10240

Directory and user permissions

Ensure that the users who install IBM InfoSphere Information Server have read and write permissions for the InfoSphere Information Server installation directories.

The following table shows the default InfoSphere Information Server installation directories that require read and write permissions by the specified users.

Table 52. Required directory permissions

Installation components	Directories	Permissions	Users
InfoSphere Information Server	Installation parent directories: **AIX** • /opt • /opt/IBM **Linux** • /opt • /opt/IBM **Windows** • C:\ • C:\IBM • %WinDir%: the Windows directory. Typically C:\WINDOWS	Read or read and write	**AIX** Root user **Linux** Root user **Windows** Local administrator or domain administrator with local administrator privileges
InfoSphere Information Server	Installation directory: **AIX** /opt/IBM/InformationServer **Linux** /opt/IBM/InformationServer **Windows** C:\IBM\InformationServer	Read and write	**AIX** Root user **Linux** Root user **Windows** Local administrator or domain administrator with local administrator privileges
Engine tier and IBM InfoSphere DataStage	Installation parent directories: **AIX** • /opt • /opt/IBM **Linux** • /opt • /opt/IBM **Windows** • C:\ • C:\IBM • %WinDir%: the Windows directory. Typically C:\WINDOWS	Read or read and write	InfoSphere DataStage administrator (default user name is dsadm)

Table 52. Required directory permissions (continued)

Installation components	Directories	Permissions	Users
Engine tier and InfoSphere DataStage	Installation directory: **AIX** `/opt/IBM/InformationServer` **Linux** `/opt/IBM/InformationServer` **Windows** `C:\IBM\InformationServer`	Read and write	InfoSphere DataStage administrator (default user name is dsadm)
InfoSphere Information Server	Temporary directory: **AIX** `/tmp` **Linux** `/tmp` **Windows** This directory is typically the *x*`:\Documents and Settings\`*user_name*`\Local Settings\temp` directory, where *user_name* is the MicrosoftWindows login name for the user who installs the product. On Windows Server 2008 and Windows 7, this directory is typically the *x*`:\Users\`*user_name*`\AppData\Local\Temp` directory.	Read and write	**AIX** • Root user • Database instance owner (default user name is db2inst1) **Linux** • Root user • Database instance owner (default user name is db2inst1) **Windows** • Local administrator or domain administrator with local administrator privileges • Database owner (default user name is db2admin)

Table 52. Required directory permissions (continued)

Installation components	Directories	Permissions	Users
Metadata repository	Database server directories for IBM DB2: **AIX** • `/opt` • `/opt/IBM` • `/opt/IBM/db2/V10.5/` **Linux** • `/opt` • `/opt/IBM` • `/opt/IBM/db2/V10.5/` **Windows** • `C:\` • `C:\IBM` • `C:\IBM\SQLLIB` • `%WinDir%`: the Windows directory. Typically `C:\WINDOWS` If you are using another database system, refer to the documentation for the database system for directory permissions information.	Read or read and write	**AIX** Database instance owner (default user name is db2inst1) **Linux** Database instance owner (default user name is db2inst1) **Windows** Database owner (default user name is db2admin)
Metadata repository	Database storage directories: **AIX** `/opt/IBM/InformationServer/Repository` **Linux** `/opt/IBM/InformationServer/Repository` **Windows** `C:\DB2`	Read and write	**AIX** Database instance owner (default user name is db2inst1) **Linux** Database instance owner (default user name is db2inst1) **Windows** Database owner (default user name is db2admin)

Table 52. Required directory permissions (continued)

Installation components	Directories	Permissions	Users
IBM WebSphere Application Server	Installation directory: **AIX** `/usr/IBM/WebSphere/AppServer` **Linux** `/opt/IBM/WebSphere/AppServer` **Windows** `C:\IBM\WebSphere\AppServer`	Read and write	**AIX** Root user **Linux** Root user **Windows** Local administrator or domain administrator with local administrator privileges

Setting up administrator and owner accounts

When you run the IBM InfoSphere Information Server installation program, you must specify the user names and passwords for administrator and owner accounts for the program. In some circumstances, the installation program can create these accounts for you, but depending on your system administration policies, you might want to create these user IDs and passwords on each of the tier computers before you install InfoSphere Information Server.

Before you begin

Make sure that you are familiar with the naming restrictions described in "Naming restrictions for user IDs and passwords" on page 110.

About this task

Creating administrator and owner accounts before you install InfoSphere Information Server ensures that operating system policies, such as requiring a password change at the first login, do not interfere with the installation program use of the accounts.

Create the accounts as local operating system accounts, lightweight directory access protocol (LDAP), accounts, or NIS accounts.

Important: If you use LDAP, PAM, or NIS accounts, you must create these accounts before you run the InfoSphere Information Server installation program. The installation program cannot create these accounts.

When you create the accounts, record the user ID and password for each user. If you choose to not use the default user IDs, the names that you choose must be distinct from each other in the same way that the default user IDs are.

Procedure

1. On the operating systems where you install InfoSphere Information Server, make sure that the user who installs InfoSphere Information Server can log in using the following accounts.

   The user must log in using one of these accounts to install InfoSphere Information Server.

Table 53. User accounts that are required to install the product

Operating system	Accounts
AIX **Linux**	The root user account.
Windows	A user account in the local Administrators group. This account must be one of the following types: • A local account that is directly assigned to the Windows local Administrators group. • A domain account that is directly assigned to the Windows local Administrators group.

2. **AIX Linux**: On the operating systems where you install each engine tier, create the following additional accounts and record the passwords.

   Note: If your installation does not include an engine tier, you can skip this step.

Table 54. Accounts that you create for an engine tier (AIX Linux)

Account description	Default user name	Primary group	Secondary group	Notes®
Administrator for the InfoSphere Information Server engine	`dsadm`	`dstage`	You customize this group to apply to your security model.	This user performs administrative tasks, such as adding new projects and defining their properties. To run jobs, create a separate user with the developer role. This account must be an account on the local operating system. During installation time, you cannot specify a Lightweight Directory Access Protocol (LDAP) account as the engine administrator account.

3. Create the following database schema accounts based on your product set and record the passwords.

   Note: When using a DB2 database, these accounts are operating system accounts. When you use a DB2 database for the repository, the DB2 database uses the operating system authentication for connection requests. There are no specific group membership requirements for these accounts.

Table 55. Database schema accounts that you create on the metadata repository tier

Products	Account description	Default user name	Notes
Required for all products	Metadata repository owner	`xmeta`	**Important:** Do not specify the DB2 instance owner or administration server user for this user name.

Table 55. Database schema accounts that you create on the metadata repository tier (continued)

Products	Account description	Default user name	Notes
Required for all products	Staging area user	`xmetasr`	**Important:** This user name must be different from the user name that you specify for the owner of the metadata repository (xmeta). The xmetasr schema must be in the same database as the xmeta schema. Do not specify the DB2 instance owner or administration server user for this user name.
Information Analyzer	Owner of the Information Analyzer analysis database	`iauser`	**Important:** Do not specify the DB2 instance owner or administration server user for this user name.
QualityStage	Owner of the QualityStage Standardization Rules Designer database	`srduser`	**Important:** Do not specify the DB2 instance owner or administration server user for this user name. If you are installing the repository tier, this ID must be created before or during installation. If you use the Database Support scripts to create your metadata repository database, do not create this user ID before installing InfoSphere Information Server. Provide the user ID and password for the srduser during installation. Scripts to configure this user and its tables will be generated by the installation program and will be put in the *IIS_install_path*/`SQLScripts` directory.
QualityStage (optional)	QualityStage Matching Database	`No default user name`	The InfoSphere Information Server installation program does not create the Match Designer results database. You create the database before or after the installation, as long as the database is configured and accessible when you use the Match Designer.
DataStage (optional)	Owner of the operations database for the DataStage Operations Console	`dsodb`	Do not specify the DB2 instance owner or administration server user for this user name. If you are installing the repository tier, this ID must be created before or during installation if you want the installation program to create the operations database. If you use the Database Support scripts to create your metadata repository database, scripts to configure this user and its tables will be generated during installation and are put in the *IIS_install_path*/`SQLScripts` directory.

4. If you choose to install DB2 as part of the InfoSphere Information Server installation, create the following additional accounts on the operating system where you install the metadata repository tier and record the passwords.

Table 56. Accounts that you create on the metadata repository tier before installing DB2

Account description	Default user name	Primary group	Secondary group	Notes
Windows DB2 administration server user	db2admin	db2admns	None	Create the operating system account for DB2 database only.
AIX **Linux** DB2 instance owner	db2inst1	db2iadm1	None	Create the operating system account for DB2 database only. **AIX Linux:** This account must have read, write, and execute permission to the following directories: • /tmp • /var/tmp
AIX **Linux** DB2 fenced user	db2fenc1	db2fadm1	None	Create the operating system account for DB2 database only.

AIX Linux: The InfoSphere Information Server installation program must have root access. Root must have write permissions to the following directories:

- /tmp
- /var
- /var/tmp
- /usr/local/bin

5. Verify permissions for all users including root. If any directories are mounted from network file systems such as AFS™, permissions might not be adequate to perform a successful installation.
6. Determine the names and passwords to use for the following administrator accounts:

Table 57. Administrator accounts set up by the installation program

Account description	Default user name
Administrator for IBM WebSphere Application Server Network Deployment	wasadmin
Administrator for InfoSphere Information Server	isadmin

7. Log in to each account to confirm that it functions correctly. Some operating systems are set up to force a password change during the first login. Therefore, a manual login is required to activate the accounts appropriately.

Naming restrictions for user IDs and passwords

Some restrictions apply to the names that you choose for user IDs and passwords. These restrictions are imposed by the underlying components that IBM InfoSphere Information Server uses.

All user IDs and passwords

The following characters are not permitted:

(white space)
! (exclamation point)
" (double quotation mark)
(number sign)

% (percent sign)

& (ampersand)

' (single quotation mark)

: (colon)

{ (left brace)

} (right brace)

Leading or trailing white space, such as spaces, tabs, and line continuation characters, is removed by the installation program. Any existing user IDs and passwords cannot have leading or trailing white space characters as part of their definitions.

Internal user registry user IDs and group IDs

Only alphanumeric characters and the following characters are permitted:

- (hyphen)

. (period)

= (equal sign)

_ (underscore)

\ (backslash)†

For the Administrator user ID (default isadmin), alphanumeric characters and the following characters are permitted:

$ (dollar sign)

- (hyphen)

, (comma)

. (period)

= (equal sign)

_ (underscore)

@ (commercial at)

\ (backslash)†

†When searching in the IBM InfoSphere Information Server Web administration console for a user or group name that contains a backslash (\), the backslash must be escaped with another backslash (for example, '\\').

IBM WebSphere Application Server administrator password

The following characters are not permitted:

- Space characters
- ^ (caret)

IBM DB2 database user names

- You can use these characters in IDs if permitted by the operating system:

  A through Z

  0 through 9

  ! (exclamation point)

  # (number sign)

  $ (dollar sign. Must not be the first character.)

  % (percent sign)

  ((opening parenthesis)

) (closing parenthesis)

AIX Linux : User names and passwords cannot start with a numeric digit.

- **AIX Linux** : The DB2 database requires these user accounts: DB2 instance owner and DB2 fenced user. These account names cannot begin with these characters:

  SQL

  SYS

  IBM
- In addition, these account names cannot use these words:

  USERS

  ADMINS

  GUESTS

  PUBLIC

  LOCAL

  Any SQL reserved word in the SQL reference for the database system
- Names cannot include accented characters.
- **AIX Linux :** Users, groups, and instance names must be in lowercase letters.

IBM WebSphere Application Server trust store password

The password can be made up of only US-ASCII characters and must be a minimum of six characters in length.

Character length limits

- **AIX:** Operating system user IDs: 8
- WebSphere Application Server IDs: 60
- DB2 user IDs and passwords:
  - **AIX Linux**: 8
  - **Microsoft Windows**: 30

Project naming conventions

Use these conventions when you name your projects. IBM InfoSphere DataStage, IBM InfoSphere QualityStage, and IBM InfoSphere Information Analyzer store jobs, analyses, and logs in these projects.

You need to choose the name of one InfoSphere DataStage project during the initial installation. The name of a project is limited to a maximum of 40 characters. The project name can contain alphanumeric characters, including underscores (_). Project names cannot use these reserved letter combinations:

- ds
- DS
- uv
- UV

Example

A useful naming scheme has project names that you maintain in unison with the source code control. As projects are promoted through the source control, the name of the phase and the project reflect the version in this form:

Phase_ProjectName_version

where *Phase* corresponds to the phase in the application development life cycle:

dev Development

it Integration test

uat User acceptance test

prod Production

Pre-installing databases

The IBM InfoSphere Information Server installation program can install the metadata repository database and analysis database for you (along with other product module databases in the repository tier) if you choose to install the InfoSphere Information Server repository tier on a local computer. Or, you can choose to install the metadata repository database and analysis database yourself on a new or existing database system before you install InfoSphere Information Server. The following tasks describe how to install the databases before you install InfoSphere Information Server, if that is what you choose to do.

Preparing database system software

Whether you are installing new database system software or using an existing database system, you must prepare the database systems to be able to support IBM InfoSphere Information Server databases.

Installing DB2 database software:

Preinstall the DB2 database system if necessary before you run the IBM InfoSphere Information Server installation program.

About this task

To preinstall DB2 before you run the InfoSphere Information Server installation program, run the DB2 database system installation program that is included in the InfoSphere Information Server installation media. To learn what DB2 installation packages are included with your InfoSphere Information Server installation package, see the system requirements on the product support site at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.

Preinstall the DB2 database system in the following situations:

- To use an existing installation of DB2 Enterprise Server Edition Version 10.5, rather than the one that can be installed with InfoSphere Information Server
- To build a DB2 database cluster or high availability disaster recovery (HADR) architecture for high availability

Procedure

1. Obtain a DB2 database system installation package from IBM. On the installation media, the DB2 database system installation packages are in the following directories:

Operating system	Directory
AIX	*installable_image_directory*/payloads/Shared/DB2/payload
Linux	*installable_image_directory*/payloads/Shared/DB2/payload

Operating system	Directory
Windows	*installable_image_directory*\payloads\Shared\DB2\payload

2. Install the DB2 database system as described in the DB2 database system documentation:

Database version	Link
IBM DB2 Database, Version 10.5	publib.boulder.ibm.com/infocenter/db2luw/v10r5/index.jsp

3. Optional: If you are using DB2 diagnostics, run the following command to set the DB2 Database Monitor Heap size to 192 or greater:

   ```
   db2 update dbm cfg using MON_HEAP_SZ 192 immediate
   ```

   Restart DB2 for changes to take effect.

   Note: The DB2 Database Monitor Heap is used for various database event monitors and snapshots. These monitoring tools do not work under the default size setting.

What to do next

Prepare your database system for repository tier databases. See "Setting up an existing installation of DB2 database software."

Setting up an existing installation of DB2 database software:

Ensure that your DB2 database system installation is set up to support IBM InfoSphere Information Server repository tier databases.

Procedure

1. The database system must be a supported version. See the system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.
2. The database system must be located on a computer that is network accessible by the computers where the services tier and engine tiers are installed.
3. The database system cannot be set up as a partitioned database environment. This type of environment and the fast communications manager (FCM) are not supported for use with InfoSphere Information Server.
4. If the database system is on the same computer as InfoSphere Information Server, ensure that DB2 Fault Monitoring components are enabled. Ensure that the `/etc/init/db2fmcd.conf` file exists.
5. Ensure that the `/etc/services` file contains a port entry for the instance:

   ```
   db2c_db2_instance_owner    port_number/tcp
   ```

   For example:

   ```
   db2c_db2inst1    50000/tcp
   ```

Pre-installing the metadata repository database

You can install the metadata repository database with the IBM InfoSphere Information Server installation program. Alternatively, you can use scripts to manually create the metadata repository database in an existing database instance.

The metadata repository is a shared component that stores design-time, runtime, glossary, and other metadata for product modules in the InfoSphere Information Server suite. The metadata repository database includes both the metadata repository and a staging area. The staging area is a schema within the metadata repository database that stores staged imports for InfoSphere Metadata Asset Manager.

Important: IBM certifies and validates InfoSphere Information Server using the provided database creation scripts. If you modify the scripts beyond the changes that are specified here, IBM will not support software failures that might result from the modifications. Modifications to these scripts that do not have a functional impact might require that you make similar changes in the future to scripts that are generated by other administrative tools that depend on the information in the scripts here. You are responsible for making those changes.

Pre-installing the metadata repository database in a DB2 database system by using scripts (Windows):

You can use an existing installation of a DB2 database to create a metadata repository for InfoSphere Information Server on Microsoft Windows.

About this task

Use the scripts in the following situations:

- You are configuring a DB2 cluster for the databases. See "Metadata repository database clustering for high availability" on page 118.
- You cannot run the installation program locally on the target computer.
- You do not want to use the InfoSphere Information Server installation program to configure the database.

Attention: Do not change the name of the tablespace named XMETA. This name is used internally by InfoSphere Information Server.

Attention: The scripts run the following commands when configuring the DB2 database system. These changes affect all databases in the DB2 instance:

```
db2set DB2_INLIST_TO_NLJN=YES
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON
db2set DB2_REDUCED_OPTIMIZATION=YES
```

Procedure

1. Navigate to the *install_dir*\DatabaseSupport directory of the InfoSphere Information Server installation media, where *install_dir* is the directory where you extracted the InfoSphere Information Server media.

Database version	Directory
DB2 10.5	*install_dir*\DatabaseSupport\Windows\MetadataRepository\DB2_LUW_10_5

2. Open the script file in a text editor and make the following changes. On 64-bit platforms, you must edit both of the following scripts:

DB2 instance version	Script files
32-bit	create_xmeta_db.sql

DB2 instance version	Script files
64-bit	create_xmeta_db.sql configure_xmeta_db_64bit.sql

a. Replace all occurrences of @DATABASE_NAME@ (including the @ symbols) with the name of the database that you want to create. For example, replace @DATABASE_NAME@ with your database name:

   xmeta

b. Replace all occurrences of @INSTALL_ROOT@ (including the @ symbols) with the path to the root directory where your DB2 database is installed. The default root directory is C:\. For example, replace @INSTALL_ROOT@ with the path to the root directory for your DB2 database:

   C:\

   When you run the script, DB2 creates a directory under the root directory where your database is installed. This directory contains the data files for your database. The default directory is C:\db2\NODE0000.

c. Replace all occurrences of @DATABASE_ALIAS@ (including the @ symbols) with the name of the database that you want to create. For example, replace @DATABASE_ALIAS@ with your database name:

   xmeta

d. Replace all occurrences of @DATABASE_USERNAME@ (including the @ symbols) with the name of the operating system user that you want to have access to the database. For example, replace @DATABASE_USERNAME@ with the operating system user name:

   xmeta

3. Run the operating system script to invoke the SQL script:
   a. Open a DB2 Command Window.
   b. From the DB2 Command Window, change directory to the same location as the edited script and run the following command:

      create_xmeta_db.cmd [*logfile*]

      logfile is the name of the file where you want to save the results of running the command. This argument is optional.
4. Open the db2_grant_Permissions.sql file in a text editor and make the following changes:
   a. Replace all occurrences of @DATABASE_NAME@ (including the @ symbols) with the name of the database that you created. For example, replace @DATABASE_NAME@ with your database name:

      xmeta

   b. Replace all occurrences of @DATABASE_USERNAME@ (including the @ symbols) with the name of the user that you want to grant access to the staging area schema. This user must be different from the operating system user that you specified in step 2d. For example, replace @DATABASE_USERNAME@ with the user name for the staging area schema:

      xmetasr

5. Run the operating system script to invoke the SQL script:
   a. Open a DB2 Command Window.
   b. From the DB2 Command Window, change directory to the same location as the edited script and run the following command:

      configure_staging_area.cmd [*logfile*]

logfile is the name of the file where you want to save the results of running the command. This argument is optional.

6. Restart DB2.

Pre-installing the metadata repository database in a DB2 database system by using scripts (AIX Linux):

You can use an existing installation of a DB2 database to create a metadata repository for InfoSphere Information Server.

About this task

The scripts run the following commands when configuring the DB2 database system. These changes affect all databases in the DB2 instance:

```
db2set DB2_INLIST_TO_NLJN=YES
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON
db2set DB2_REDUCED_OPTIMIZATION=YES
```

Procedure

1. Navigate to the *install_dir*\DatabaseSupport directory of the InfoSphere Information Server installation media, where *install_dir* is the directory where you extracted the InfoSphere Information Server installation media.

Database version	Directory
DB2 10.5	*install_dir*\DatabaseSupport\UNIX_Linux\MetadataRepository\DB2_LUW_10_5

2. Open the script file in a text editor and make the following changes. On 64-bit platforms, you must edit both of the listed scripts:

DB2 instance version	Script files
32-bit	create_xmeta_db.sql
64-bit	create_xmeta_db.sql configure_xmeta_db_64bit.sql

Important: IBM certifies and validates InfoSphere Information Server using the provided database creation scripts. If you modify the scripts beyond the changes that are specified here, IBM will not support software failures that might result from the modifications. Modifications to these scripts that do not have a functional impact might require that you make similar changes in the future to scripts that are generated by other administrative tools that depend on the information in the scripts here. You are responsible for making those changes.

a. Replace all occurrences of @DATABASE_NAME@ (including the @ symbols) with the name of the database that you want to create. For example, replace @DATABASE_NAME@ with your database name:

   xmeta

b. Replace all occurrences of @INSTALL_ROOT@ (including the @ symbols) with the path to the root directory where your DB2 database is installed. For example, replace @INSTALL_ROOT@ with the path to the root directory for your DB2 database:

   /home/db2inst1/db/xmeta

c. Replace all occurrences of @DATABASE_ALIAS@ (including the @ symbols) with the name of the database that you want to create. For example, replace @DATABASE_ALIAS@ with your database name:

   xmeta

d. Replace all occurrences of @DATABASE_USERNAME@ (including the @ symbols) with the name of the operating system user that you want to have access to the database. For example, replace @DATABASE_USERNAME@ with the operating system user name:

   xmeta

3. Run the operating system script to invoke the SQL script:
   a. Log on as the DB2 instance owner.
   b. Change to the directory where the edited script is located, and run the following command:

      create_xmeta_db.sh [*logfile*]

      logfile is the name of the file where you want to save the results of running the command. This argument is optional.
4. Open the db2_grant_Permissions.sql file in a text editor and make the following changes:
   a. Replace all occurrences of @DATABASE_NAME@ (including the @ symbols) with the name of the database that you created. For example, replace @DATABASE_NAME@ with your database name:

      xmeta

   b. Replace all occurrences of @DATABASE_USERNAME@ (including the @ symbols) with the name of the user that you want to grant access to the staging area schema. This user must be different from the operating system user that you specified in the create scripts. For example, replace @DATABASE_USERNAME@ with the user name for the staging area schema:

      xmetasr

5. Run the operating system script to invoke the SQL script:
   a. Log on as the DB2 instance owner.
   b. Change to the directory where the edited script is located, and run the following command:

      configure_staging_area.sh [*logfile*]

      logfile is the name of the file where you want to save the results of running the command. This argument is optional.
6. If you use DB2 diagnostics, run the following command to set the DB2 Database Monitor Heap size to 192 or greater:

   Note: The DB2 Database Monitor Heap is used for various database event monitors and snapshots. These monitoring tools do not work under the default size setting. See http://www.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.admin.config.doc/doc/r0000306.html.

   db2 update dbm cfg using MON_HEAP_SZ 192 immediate

Metadata repository database clustering for high availability:

To provide high availability for your metadata repository database, you can set up the database within a database cluster.

IBM InfoSphere Information Server does not support clustered configurations of the IBM InfoSphere Information Analyzer database or the IBM InfoSphere QualityStage Match Designer database.

Setting up the metadata repository database in an IBM DB2 cluster:

Use this road map to create an IBM InfoSphere Information Server configuration in which the metadata repository database is set up within an IBM DB2 cluster.

If you want to create a new DB2 database cluster installation, complete the tasks in the following table. If you want to integrate InfoSphere Information Server into an existing DB2 database cluster, read through the task topics to make sure that your cluster configuration is compatible with InfoSphere Information Server.

Table 58. Setting up an IBM DB2 database cluster

Task	Link for more information
To set up a DB2 database cluster, first configure the servers, storage, and network.	"Preparing the servers, file system, and storage"
(Optional) Install high availability (HA) software to monitor the health of the servers and to manage failover.	"Installing high availability software on DB2 database system servers" on page 121
Before you install InfoSphere Information Server, install the DB2 database system, create a cluster, and create the metadata repository database on the primary node in the cluster.	"Installing the DB2 database system and creating the cluster and metadata repository database" on page 122
Create users and groups. Run the InfoSphere Information Server installation program.	"Running the installation program for a setup that includes DB2 clustering" on page 123
To allow IBM WebSphere Application Server to reconnect to the standby server smoothly, configure Automatic Client Reroute.	"Configuring automatic client reroute" on page 124
To enable WebSphere Application Server to communicate with more than one server, use the WebSphere Application Server administrative console to specify an alternate database address and port.	Updating IBM WebSphere Application Server configuration for Oracle RAC
Verify the state of the cluster.	"Verifying the IBM DB2 database cluster" on page 129

InfoSphere Information Server does not support cluster configurations of the IBM InfoSphere Information Analyzer analysis database or IBM InfoSphere QualityStage Match Designer database.

To deploy a clustered topology, you must have a solid understanding of network technologies such as protocols, layers, devices, and so on. Clustered topologies are technically complex, and require expertise to implement successfully.

Preparing the servers, file system, and storage:

To set up an IBM DB2 database cluster, first configure the servers, storage, and network.

For high availability, the cluster must include at least two physically separate computers. If possible, use identical hardware with identical internal storage, memory, and processing power.

The metadata repository database cannot be partitioned. They must exist in a non-partitioned environment. For this reason, you cannot set up the cluster in a "mutual takeover" configuration. You must set up an "idle standby" configuration. In an idle standby configuration, a primary server processes all the workload. A secondary server is idle, or in standby mode, ready to take over the workload if there is a failure on the primary system.

The servers must be able to communicate with each other over high-speed networks. To minimize single points of failure, each server must have redundant connections to at least two physically separate high-speed networks.

To support failover when a problem occurs, the HA software creates a heartbeat signal between the servers. Set up at least two more physically separate networks for this signal. These networks can be serial or other networks. The networks that carry the heartbeat signal must be physically separate and isolated from the high-speed networks.

The following diagram shows a two-server active-passive architecture with minimum network redundancy. The diagram shows two servers (Node 1 and Node 2). Each server has four network interfaces (en0, en1, en2, and en3). There are two redundant high-speed networks (Network A and Network B). On each server, two network interfaces are connected to Network A, and two network interfaces are connected to Network B. This arrangement creates redundancy at the network interfaces and within the overall high-speed network architecture.

Each server also has two serial interfaces (tty1 and tty2). Each serial interface on one server is connected to a serial interface on the other. This arrangement creates a redundant serial "ring" network for the heartbeat signal.

Create the DB2 database system file system on disk partitions that are shared by all cluster nodes. The disk partitions must be on highly available media such as a storage area network (SAN).

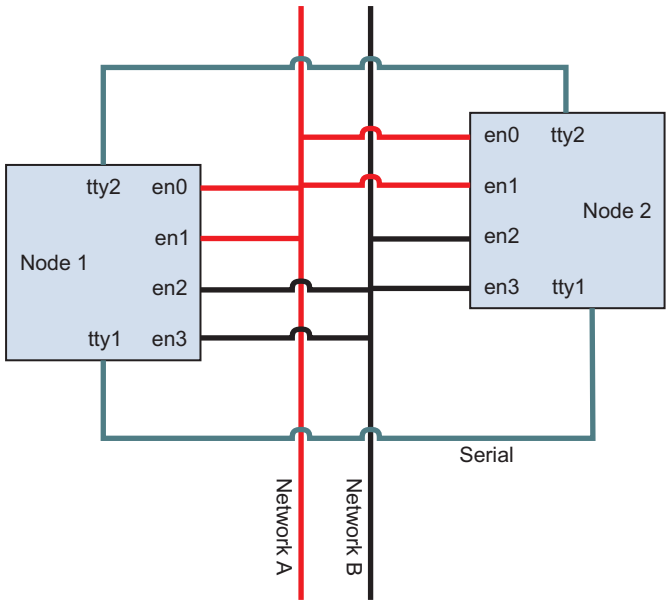


Figure 39. Two servers with redundant connections to networks

For more specific information about DB2 database cluster servers, storage, or network considerations, see the IBM DB2 database system documentation:

- IBM DB2 Database, Version 10.5 for Linux, UNIX, and Windows "Database fundamentals" in the KnowledgeCenter: http://www.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.container.doc/doc/c0052964.html

Installing high availability software on DB2 database system servers:

You can optionally install and configure IBM Tivoli System Automation for Multiplatforms (SA MP) to monitor the health of the IBM DB2 database system servers and to manage failover.

Tivoli System Automation for Multiplatforms software is shipped as part of the DB2 database system installation package. This high availability (HA) software monitors the health of the servers, the network, the operating system, and selected software components. If a critical item fails, the software triggers a failover to the standby server. Then the automatic client reroute feature causes IBM WebSphere Application Server to connect to another server in the cluster.

If you plan to implement SA MP, install it before you install and configure IBM InfoSphere Information Server. When you install SA MP, you configure a floating IP address that is shared between the servers. When you configure automatic client reroute, you provide this IP address as the "alternate server host name".

You can also choose not to install SA MP. If you do not install the software, if a failure occurs on the primary server, the database administrator must manually detect the failure and perform a failover. Then automatic client reroute redirects WebSphere Application Server to the new server.

To install and configure the software, see the DB2 database system documentation:

- ‚IBMDB2 10.5 for Linux, UNIX, and Windows "Installing and upgrading SA MP with the DB2 installer" in the Knowledge Center: http://www.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.qb.server.doc/doc/c0051342.html

Installing the DB2 database system and creating the cluster and metadata repository database:

Before you install IBM InfoSphere Information Server, install the IBM DB2 database system, create a cluster, and create the metadata repository database on the primary node.

Before you begin

A separate copy of the DB2 database system must be installed on each server.

Procedure

1. Install the DB2 database system on the computers that are to make up the cluster.
   - Each DB2 database system instance must be the same version of the DB2 database system.
   - The installations must have the same container file system and installation path, for example: `/home/db2inst1/sqllib`.
   - Create identical user names and passwords for each DB2 database system instance.
   - Within the installation media, the DB2 database system installation packages are in the following directories:

Operating system	Directory
AIX	`payloads/Shared/DB2/payload`
Linux	`payloads/Shared/DB2/payload`
Windows	`payloads\Shared\DB2\payload`

   For information about installing the DB2 database system, see the DB2 database system documentation:
   - DB2 10.5 for Linux, UNIX, and Windows in the Knowledge Center.

   The IBM InfoSphere Information Server installation package includes installation packages for DB2. To preinstall DB2 before you run the InfoSphere Information Server installation program, you can run the DB2 database system installation program that is included in the InfoSphere Information Server installation media. To learn what DB2 installation packages are included with your InfoSphere Information Server installation package, see the system requirements at publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/softwareReqsForProduct.html.
2. Set up the cluster.

   The metadata repository database cannot be partitioned. The database must exist in a non-partitioned environment. For this reason, you cannot set up the system in a “mutual takeover” configuration. You must set up an “idle standby” configuration. In this configuration, a primary system processes all the workload. A secondary system is idle, or in standby mode, ready to take over the workload if there is a failure on the primary system.

For information about creating a DB2 database cluster, search for "Configuring a clustered environment for high availability" in the DB2 database system documentation:

- DB2 10.5 for Linux, UNIX, and Windows in the Knowledge Center.

3. On the primary node, create the metadata repository database within the DB2 database system. To create the database, use the scripts that are supplied with the InfoSphere Information Server installation media. Locate the database file system on a disk partition that is shared by all nodes in the cluster. The disk partition must be on highly available media such as a storage area network (SAN).

Operating system	Procedure
AIX	Configuring an existing installation of DB2 database software by using scripts (AIX)
Linux	Configuring an existing installation of DB2 database software by using scripts (Linux)
Windows	"Pre-installing the metadata repository database in a DB2 database system by using scripts (Windows)" on page 115

Running the installation program for a setup that includes DB2 clustering:

Create administrator and component owner accounts and groups. Run the IBM InfoSphere Information Server installation program.

Procedure

1. Create identical administrator and component owner accounts and user groups on all servers in the cluster. Make sure that the group IDs and user IDs are identical on both servers. The following sample commands create some of the accounts and groups:

Operating system	Commands
AIX	`/usr/bin/mkgroup id='202' db2iadm1` `/usr/bin/mkgroup id='203' db2fadm1` `/usr/bin/mkgroup id='205' xmeta` `/usr/bin/mkgroup id='206' iauser` `/usr/bin/mkgroup id='207' dstage` `/usr/bin/mkuser pgrp=db2iadm1` `  groups=db2iadm` `  home=/home/db2inst1 db2inst1` `/usr/bin/mkuser pgrp=db2fadm1` `  groups=db2fadm` `  home=/home/db2fenc1 db2fenc1` `/usr/bin/mkuser pgrp=xmeta` `  groups=xmeta` `  home=/home/xmeta  xmeta` `/usr/bin/mkuser prgp=iauser` `  groups=iauser` `  home=/home/iauser  iauser` `/usr/bin/mkuser pgrp=dstage` `  groups=dstage` `  home=/home/dsadm  dsadm`

Operating system	Commands
Linux	/usr/sbin/groupadd db2iadm1 /usr/sbin/groupadd db2fadm1 /usr/sbin/groupadd xmeta /usr/sbin/groupadd iauser /usr/sbin/groupadd dstage /usr/sbin/useradd -g db2iadm1 -d /home/db2inst1 db2inst1 /usr/sbin/useradd -g db2fadm1 -d /home/db2fenc1 db2fenc1 /usr/sbin/useradd -g xmeta -d /home/xmeta xmeta /usr/sbin/useradd -g iauser -d /home/iauser iauser /usr/sbin/useradd -g dstage -d /home/dsadm dsadm

For details on creating administrator users and groups, see "Setting up administrator and owner accounts" on page 107.

2. Run the InfoSphere Information Server installation program in graphical mode. Complete the entire InfoSphere Information Server installation as described in "Scenario D: Clustered metadata repository tier and services tier configuration" on page 61.

Configuring automatic client reroute:

For IBM WebSphere Application Server to reconnect to the standby server smoothly in a DB2 clustered or HADR configuration, configure automatic client reroute.

About this task

In an IBM InfoSphere Information Server installation with a clustered DB2 database system setup, automatic client reroute can redirect WebSphere Application Server to connect to an alternate server when the primary server fails. The alternate server is preregistered within automatic client reroute. When the failure occurs, automatic client reroute first tries to reconnect to the primary database. If the connection fails, automatic client reroute establishes a connection with the standby server.

If an update, deletion, or insertion transaction was taking place and the failure occurred before the transaction was committed, the user receives an error message. When the user tries the transaction again, the standby server handles the transaction. If a read transaction was taking place, the read transaction is automatically run on the standby server, and no error message is displayed.

Procedure

1. Log in to the primary node.
2. Run the following command to configure automatic client reroute for the database:

   ```
   db2 update alternate server for database database using
     hostname standby_IP port port
   ```

   The *standby_IP* can be an IPv4 address or an IPv6 address, as shown in the following examples:

IPv4:

```
db2 update alternate server for database xmeta using
  hostname 192.0.2.7 port 60000
```

IPv6:

```
db2 update alternate server for database xmeta using
  hostname ::ffff:192.0.2.7 port 60000
```

WebSphere Application Server Network Deployment: Specifying the alternate address and port for a DB2 clustered database:

For IBM WebSphere Application Server Network Deployment to communicate with more than one server, specify an alternate address and port for a DB2 database.

Procedure

1. Start the WebSphere Application Server administrative console.
2. Log in as the WebSphere Application Server administrator (default: wasadmin), or as another user with administrative privileges.
3. In the navigation pane, expand **Resources**, expand **JDBC**, and click **JDBC Providers**. The JDBC providers page opens.
4. Complete the following steps for each of these JDBC providers that might be listed in the Administrative console:
   - **ASB JDBC Provider**
   - **ASB XA JDBC Provider**
   - **QSSRD Repository JDBC Provider**
   - **QSSRD Repository XA JDBC Provider**

   a. Click the JDBC provider link. The JDBC provider page opens.

   b. Under **Additional Properties**, click **Data sources**

   c. Complete the following steps for each of these data sources that might be listed in the Administrative console:
      - **ASB JDBC DataSource**
      - **ASB JDBC XA DataSource**
      - **ASB JDBC Staging XA DataSource**
      - **JReport JDBC DataSource**
      - **QSSRD DataSource**
      - **QSSRD Global XA DataSource**

      1) Click the data source.
      2) Under Additional Properties, click **Connection Pool Properties**.
      3) In the Purge policy field, select **EntirePool**.
      4) Click **OK**.
      5) Under Additional Properties, click **Custom Properties**.
      6) Click **clientRerouteAlternateServerName**. If this property is not present, click **New** and add the property and value.
      7) In the **Value** field, type the alternate server address:
         - If IBM Tivoli System Automation for Multiplatforms (SA MP) is running in the cluster, type the service IP address that the cluster nodes share.
         - If SA MP is this not running in the cluster, type the IP address of the standby server.
      8) Click **OK**.

9) Click **clientRerouteAlternatePortNumber**. If this property is not present, click **New** and add the property and value.
10) In the **Value** field, type the alternate server port number:
    - If SA MP is running in the cluster, type the service port number that the cluster nodes share.
    - If SA MP is not running in the cluster, type the port number of the standby server.
11) Click **enableSeamlessFailover**. If this property is not present, you can click **New** and add the property and value.
12) In the **Value** field, type `1`.

    This property controls whether the server returns an error that indicates that failover occurred. For more information, see the DB2 database system documentation.
13) Click **OK**.
14) Click **New**.
15) In the **Name** field, type `maxRetriesForClientReroute`
16) In the Value field, type `2`. This parameter sets the number of connection retry attempts.
17) Click **Apply**.
18) In the **Name** field, type `retryIntervalForClientReroute` This parameter sets the sleep interval between retry attempts.
19) In the Value file, type `15`.

    Note: You might want to set these values differently. The values you choose depend upon the hardware and the topology of your configuration.
20) Click **OK**.
21) Click **Apply**.
22) In the **Messages** box at the top of the page, click **Save** to apply the changes to the master configuration.
23) Click the data source name in the line at the top of the page to return to the data source page.
24) Click **Test Connection** to test the connection.
25) Repeat for all the above data sources that are listed.

d. Repeat for all the above JDBC providers that are listed in the Administrative console.

5. If your metadata repository is configured for active-passive failover, but you did not indicate the virtual host name for your metadata repository database in the IBM InfoSphere Information Server installation program, complete the following steps:
   a. Click **Resources** > **JDBC** > **Data sources**.
   b. In the Data sources panel, click the first data source that is listed.
   c. Under **Common and required data source properties**, update the server name to the virtual host name.
   d. Click **Apply**, then click **OK**.
   e. Repeat each step for every data source that is listed.
   f. Click **Test connection**. Ensure that you see a message that says, `The test connection operation for data source` *`data_source_name`* `on server` *`server_name`* `at node` *`node_name`* `was successful.`

g. If a message displays that asks you to complete synchronization, follow the instructions.

6. Log out of the console.

WebSphere Application Server Liberty Profile: Specifying the alternate address and port for a DB2 clustered database:

For IBM WebSphere Application Server Liberty Profile to communicate with more than one server, specify an alternate address and port for a DB2 database.

Procedure

1. Edit the following file:

   AIX Linux
   IS_install_path/wlp/usr/servers/iis/datasources.xml

   Windows
   IS_install_path\wlp\usr\servers\iis\datasources.xml

2. For each <dataSource/> element in the file, add the following information in bold as attributes to the <properties.db2.jcc/> child element, substituting the alternate address and port information with your values.

```
<dataSource jndiName="jdbc/ASBDataSource"
    id="DataSource_ASBDataSource" jdbcDriverRef="db2Driver"
     isolationLevel="TRANSACTION_READ_COMMITTED"
     beginTranForResultSetScrollingAPIs="true"
    connectionManagerRef="ConnectionPool_ASBDataSource">
   <properties.db2.jcc databaseName="${iis.db.name}" portNumber="${iis.db.port}"
     driverType="4" serverName="${iis.db.host}"
     password="${iis.db.password}" user="${iis.db.user}"
     clientRerouteAlternateServerName="<alternate_host_name>"
     clientRerouteAlternatePortNumber="<alternate_port>"
     maxRetriesForClientReroute="2"
     retryIntervalForClientReroute="15"
     enableSeamlessFailover="1"
  >
</dataSource>
```

Updating the database properties URL with the alternate host and port:

To update the alternate host and port names for your metadata repository database, you specify the values in the database.properties file located on the services tier computer and run the **iisAdmin -deploy** command to deploy the changes.

Procedure

1. On the services tier computer, open the database.properties file that is used by the application server in a text editor. Ensure you edit the correct file because there are two copies of the file. The file to edit first is located in one of the following directories:
   - **For IBM WebSphere Application Server Network Deployment:**

     AIX Linux
     IS_install_path/ASBServer/apps/lib/iis/classes

     Windows
     IS_install_path\ASBServer\apps\lib\iis\classes
   - **For IBM WebSphere Application Server Liberty Profile:**

AIX Linux

IS_install_path/ASBServer/apps/lib/iis/15properties

Windows

IS_install_path\ASBServer\apps\lib\iis\15properties

2. In the file, locate the `url` keyword and change the values of the alternate host name and port to match the values of the alternate metadata repository host.

Table 59. URL keyword and value format for different database systems

Database system	`url` keyword and value
IBM DB2 cluster or HADR database system	Enter the following on one line. The → symbol indicates a line continuation. `url=jdbc\:db2\://host\:port/dbname;→` `clientRerouteAlternateServerName\=alternate_host;→` `clientRerouteAlternatePortNumber\=alternate_host_port;→` `queryCloseImplicit\=2;`

3. On the services tier computer, back up the following file and copy the file that you changed in step 2 to this one, overwriting it.

AIX Linux

IS_install_path/ASBServer/conf/database.properties

Windows

IS_install_path\ASBServer\conf\database.properties

4. Propagate the changes through the application server by entering one of the following commands.
   - **For WebSphere Application Server Liberty Profile:**

     AIX Linux

     IS_install_path/ASBServer/bin/iisAdmin.sh -deploy -srcDir *IS_install_path*/ASBServer/apps/lib/iis/15properties -libName iis -libModule 15properties

     Windows

     IS_install_path/ASBServer\bin\iisAdmin.bat -deploy -srcDir *IS_install_path*\ASBServer\apps\lib\iis\15properties -libName iis -libModule 15properties

   - **For WebSphere Application Server Network Deployment stand-alone configuration:**

     AIX Linux

     IS_install_path/ASBServer/bin/iisAdmin.sh -deploy -srcDir *IS_install_path*/ASBServer/apps/lib/iis/classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password*

     Windows

     IS_install_path/ASBServer\bin\iisAdmin.bat -deploy -srcDir *IS_install_path*\ASBServer\apps\lib\iis\classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password*

   - **For WebSphere Application Server Network Deployment clustered configuration:**

     AIX Linux

     IS_install_path/ASBServer/bin/iisAdmin.sh -deploy -srcDir *IS_install_path*/ASBServer/apps/lib/iis/classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password* -cluster *cluster_name*

Windows

```
IS_install_path/ASBServer\bin\iisAdmin.bat -deploy -srcDir
IS_install_path\ASBServer\apps\lib\iis\classes -libName iis
-libModule classes -user wasadmin -password wasadmin_password
-cluster cluster_name
```

5. Restart the application server.

Application server configuration	Procedure
Stand-alone	Restart WebSphere Application Server
Clustered	Restart the WebSphere Application Server cluster members

Verifying the IBM DB2 database cluster:

After you complete the cluster configuration and IBM InfoSphere Information Server installation, verify the state of the cluster.

Procedure

1. Log in to the primary node in the DB2 database cluster.
2. Run the **lssam** command to validate the state of the cluster. For example:

   ```
   lssam
   ```

   The command displays a report. For example:

   ```
   Online IBM.ResourceGroup:db2_db2inst1_0-rg Nominal=Online
           |- Online IBM.Application:db2_db2inst1_0-rs
                   |- Online IBM.Application:db2_db2inst1_0-rs:myprimary
                   '- Offline IBM.Application:db2_db2inst1_0-rs:mystandby
           '- Online IBM.ServiceIP:db2ip_192_0_2_16-rs
                   |- Online IBM.ServiceIP:db2ip_192_0_2_16-rs:myprimary
                   '- Offline IBM.ServiceIP:db2ip_192_0_2_16-rs:mystandby
   ```

   This report shows a two-node cluster. One node (`myprimary`) is active (online). The other node (`mystandby`) is offline.

Setting up the metadata repository database in an IBM DB2 database HADR configuration:

Use this road map to create an IBM InfoSphere Information Server implementation in which the metadata repository database is set up within a DB2 database high availability disaster recovery (HADR) configuration.

If you want to create a DB2 database HADR installation, complete the tasks in the following table. If you want to integrate InfoSphere Information Server into an existing HADR configuration, read through the task topics to make sure that your HADR configuration is compatible.

Table 60. Configuring the metadata repository database in a DB2 database HADR configuration.

Task	Link for more information
Configure the servers, the network, and the storage.	"Preparing the servers, storage, and network" on page 130
Install DB2 database system on both servers in the configuration. Use the scripts in the installation media to configure the database on the primary server.	"Installing the DB2 database system and creating the metadata repository database" on page 131

Table 60. Configuring the metadata repository database in a DB2 database HADR configuration. (continued)

Task	Link for more information
Create administrator and component owner accounts and groups, and complete the InfoSphere Information Server installation and configuration.	"Running the installation program for a HADR configuration" on page 132
On the HADR primary server, configure the HADR variables and verify your settings.	"Configuring the HADR primary server" on page 136
Back up the database from the primary server and restore it on the standby server. Then set HADR variables on the standby server.	"Configuring the HADR standby server" on page 138
To enable IBM WebSphere Application Server to communicate with either server in the HADR setup, use the WebSphere Application Server administrative console to specify an alternate database address and port.	Updating IBM WebSphere Application Server configuration for Oracle RAC
Start the HADR software on the standby server and then on the primary server.	"Starting HADR" on page 140
Validate the HADR configuration and test the takeover function.	"Validating the HADR configuration" on page 141

InfoSphere Information Server does not support HADR configurations of the IBM InfoSphere Information Analyzer analysis database or the IBM InfoSphere QualityStage Match Designer database.

To deploy a clustered or HADR topology, you must have a solid understanding of network technologies such as protocols, layers, devices, and so on. Clustered topologies are technically complex, and require expertise to implement successfully.

Preparing the servers, storage, and network:

To set up an IBM DB2 database HADR configuration, first configure the servers, storage, and network.

The high availability disaster recovery (HADR) configuration must include at least two physically separate computers. Use identical hardware with identical storage, memory, and processing power. The configuration must satisfy the following requirements:

- The computers must have identical operating system versions.
- The computers must run identical versions of DB2 database system.
- The computers must each have separate storage that is not shared between them.
- They must have the same container file system and installation path for DB2 database system.
- The servers must be able to communicate with each other using TCP/IP over a high-speed network.
- The IBM WebSphere Application Server installation must be able to communicate with both servers over the network.
- On each server, the `services` file must specify the ports for HADR. The port name is user-defined. The port number can be any number, provided there is no

conflict. The ports on the primary and standby servers do not need to be identical. However, if the ports are identical, configuration is easier.

The `services` file can be found in the following locations:

Operating system	Directory
AIX	`/etc/services`
Linux	`/etc/services`
Windows	`c:\windows\system32\drivers\etc\services`

The following example shows HADR ports specified within the `services` file:

```
# HADR ports assigned by user
ha_myhadr      51012/tcp
ha_myhadr_int  51013/tcp
```

- Use fully-qualified domain names when specifying host names while configuring HADR in the steps that follow.

For more specific information about DB2 database HADR servers, storage, or network considerations, search for these terms in the IBM DB2 database system documentation:

- DB2 10.5 for Linux, UNIX, and Windows in the Knowledge Center.

Installing the DB2 database system and creating the metadata repository database:

Before you install IBM InfoSphere Information Server, install the IBM DB2 database system on both servers and create the InfoSphere Information Server metadata repository database on the primary node.

Before you begin

A separate copy of the DB2 database system must be installed on each server.

Procedure

1. Install the DB2 database system on both computers.
   - Each DB2 database system instance must be the same version of the DB2 database system.
   - Each instance must use a separate file system for storage.
   - The installations must have the same container file system and installation path. For example: `/home/db2inst1/sqllib`.
   - Create identical user names and passwords for each DB2 database system instance.

   For information about installing the DB2 database system, see the DB2 database system documentation:

   - IBM DB2 Database, Version 10.5 for Linux, UNIX, and Windows

   The IBM InfoSphere Information Server installation package includes installation packages for DB2. To preinstall DB2 before you run the InfoSphere Information Server installation program, you can run the DB2 database system installation program that is included in the InfoSphere Information Server installation media. To learn what DB2 installation packages are included with your InfoSphere Information Server installation package, see the system requirements at publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/softwareReqsForProduct.html.

Within the installation media, the DB2 database system installation packages are in the following directories:

Operating system	Directory
AIX	`payloads/Shared/DB2/payload`
Linux	`payloads/Shared/DB2/payload`
Windows	`payloads\Shared\DB2\payload`

2. On the primary node, create the metadata repository database within the DB2 database system. To create the database, use the scripts that are supplied with the InfoSphere Information Server installation media. Locate the database file system on a disk partition that is shared by all nodes in the cluster. The disk partition must be on highly available media such as a storage area network (SAN).

Operating system	Procedure
AIX	Configuring an existing installation of DB2 database software by using scripts (AIX)
Linux	Configuring an existing installation of DB2 database software by using scripts (Linux)
Windows	"Pre-installing the metadata repository database in a DB2 database system by using scripts (Windows)" on page 115

Running the installation program for a HADR configuration:

Create administrator and component owner accounts and groups. Run the IBM InfoSphere Information Server installation program.

Procedure

1. Create identical administrator and component owner accounts and groups on both high availability disaster recovery (HADR) servers. Make sure that the group IDs and user IDs are identical on both servers. For details on creating administrator users and groups, see "Setting up administrator and owner accounts" on page 107.
2. Run the InfoSphere Information Server installation program in graphical mode. Complete the entire InfoSphere Information Server installation as described in "Scenario D: Clustered metadata repository tier and services tier configuration" on page 61.

WebSphere Application Server Network Deployment: Specifying the alternate address and port for a DB2 HADR database:

For IBM WebSphere Application Server Network Deployment to communicate with more than one server, specify an alternate address and port for a DB2 database.

Procedure

1. Start the WebSphere Application Server administrative console.
2. Log in as the WebSphere Application Server administrator (default: wasadmin), or as another user with administrative privileges.
3. In the navigation pane, expand **Resources**, expand **JDBC**, and click **JDBC Providers**. The JDBC providers page opens.

4. Complete the following steps for each of these JDBC providers that might be listed in the Administrative console:
   - **ASB JDBC Provider**
   - **ASB XA JDBC Provider**
   - **QSSRD Repository JDBC Provider**
   - **QSSRD Repository XA JDBC Provider**

   a. Click the JDBC provider link. The JDBC provider page opens.
   b. Under **Additional Properties**, click **Data sources**
   c. Complete the following steps for each of these data sources that might be listed in the Administrative console:
      - **ASB JDBC DataSource**
      - **ASB JDBC XA DataSource**
      - **ASB JDBC Staging XA DataSource**
      - **JReport JDBC DataSource**
      - **QSSRD DataSource**
      - **QSSRD Global XA DataSource**

      1) Click the data source.
      2) Under Additional Properties, click **Connection Pool Properties**.
      3) In the Purge policy field, select **EntirePool**.
      4) Click **OK**.
      5) Under Additional Properties, click **Custom Properties**.
      6) Click **clientRerouteAlternateServerName**. If this property is not present, click **New** and add the property and value.
      7) In the **Value** field, type the alternate server address:
         - If IBM Tivoli System Automation for Multiplatforms (SA MP) is running in the cluster, type the service IP address that the cluster nodes share.
         - If SA MP is this not running in the cluster, type the IP address of the standby server.
      8) Click **OK**.
      9) Click **clientRerouteAlternatePortNumber**. If this property is not present, click **New** and add the property and value.
      10) In the **Value** field, type the alternate server port number:
         - If SA MP is running in the cluster, type the service port number that the cluster nodes share.
         - If SA MP is not running in the cluster, type the port number of the standby server.
      11) Click **enableSeamlessFailover**. If this property is not present, you can click **New** and add the property and value.
      12) In the **Value** field, type 1.

          This property controls whether the server returns an error that indicates that failover occurred. For more information, see the DB2 database system documentation.
      13) Click **OK**.
      14) Click **New**.
      15) In the **Name** field, type `maxRetriesForClientReroute`

16) In the Value field, type 2. This parameter sets the number of connection retry attempts.
17) Click **Apply**.
18) In the **Name** field, type `retryIntervalForClientReroute` This parameter sets the sleep interval between retry attempts.
19) In the Value file, type 15.

    Note: You might want to set these values differently. The values you choose depend upon the hardware and the topology of your configuration.
20) Click **OK**.
21) Click **Apply**.
22) In the **Messages** box at the top of the page, click **Save** to apply the changes to the master configuration.
23) Click the data source name in the line at the top of the page to return to the data source page.
24) Click **Test Connection** to test the connection.
25) Repeat for all the above data sources that are listed.

d. Repeat for all the above JDBC providers that are listed in the Administrative console.

5. If your metadata repository is configured for active-passive failover, but you did not indicate the virtual host name for your metadata repository database in the IBM InfoSphere Information Server installation program, complete the following steps:
   a. Click **Resources** > **JDBC** > **Data sources**.
   b. In the Data sources panel, click the first data source that is listed.
   c. Under **Common and required data source properties**, update the server name to the virtual host name.
   d. Click **Apply**, then click **OK**.
   e. Repeat each step for every data source that is listed.
   f. Click **Test connection**. Ensure that you see a message that says, `The test connection operation for data source` *`data_source_name`* `on server` *`server_name`* `at node` *`node_name`* `was successful.`
   g. If a message displays that asks you to complete synchronization, follow the instructions.
6. Log out of the console.

WebSphere Application Server Liberty Profile: Specifying the alternate address and port for a DB2 HADR database:

For IBM WebSphere Application Server Liberty Profile to communicate with more than one server, specify an alternate address and port for a DB2 database.

Procedure

1. Edit the following file:

   AIX Linux
   `IS_install_path``/wlp/usr/servers/iis/datasources.xml`

   Windows
   `IS_install_path``\wlp\usr\servers\iis\datasources.xml`

2. For each <dataSource/> element in the file, add the following information in bold as attributes to the <properties.db2.jcc/> child element, substituting the alternate address and port information with your values.

```
<dataSource jndiName="jdbc/ASBDataSource"
    id="DataSource_ASBDataSource" jdbcDriverRef="db2Driver"
     isolationLevel="TRANSACTION_READ_COMMITTED"
     beginTranForResultSetScrollingAPIs="true"
    connectionManagerRef="ConnectionPool_ASBDataSource">
   <properties.db2.jcc databaseName="${iis.db.name}" portNumber="${iis.db.port}"
     driverType="4" serverName="${iis.db.host}"
     password="${iis.db.password}" user="${iis.db.user}"
     clientRerouteAlternateServerName="<alternate_host_name>"
     clientRerouteAlternatePortNumber="<alternate_port>"
     maxRetriesForClientReroute="2"
     retryIntervalForClientReroute="15"
     enableSeamlessFailover="1"
  >
</dataSource>
```

Updating the database properties URL with the alternate host and port:

To update the alternate host and port names for your metadata repository database, you specify the values in the database.properties file located on the services tier computer and run the **iisAdmin -deploy** command to deploy the changes.

Procedure

1. On the services tier computer, open the database.properties file that is used by the application server in a text editor. Ensure you edit the correct file because there are two copies of the file. The file to edit first is located in one of the following directories:
   - **For IBM WebSphere Application Server Network Deployment:**

     AIX Linux
     IS_install_path/ASBServer/apps/lib/iis/classes

     Windows
     IS_install_path\ASBServer\apps\lib\iis\classes
   - **For IBM WebSphere Application Server Liberty Profile:**

     AIX Linux
     IS_install_path/ASBServer/apps/lib/iis/15properties

     Windows
     IS_install_path\ASBServer\apps\lib\iis\15properties
2. In the file, locate the url keyword and change the values of the alternate host name and port to match the values of the alternate metadata repository host.

Table 61. URL keyword and value format for different database systems

Database system	url keyword and value
IBM DB2 cluster or HADR database system	Enter the following on one line. The → symbol indicates a line continuation. url=jdbc\:db2\://*host*\:*port*/*dbname*;→ clientRerouteAlternateServerName\=*alternate_host*;→ clientRerouteAlternatePortNumber\=*alternate_host_port*;→ queryCloseImplicit\=2;

3. On the services tier computer, back up the following file and copy the file that you changed in step 2 on page 128 to this one, overwriting it.

AIX Linux

IS_install_path/ASBServer/conf/database.properties

Windows

IS_install_path\ASBServer\conf\database.properties

4. Propagate the changes through the application server by entering one of the following commands.
   - **For WebSphere Application Server Liberty Profile:**

     AIX Linux

     IS_install_path/ASBServer/bin/iisAdmin.sh -deploy -srcDir *IS_install_path*/ASBServer/apps/lib/iis/15properties -libName iis -libModule 15properties

     Windows

     IS_install_path/ASBServer\bin\iisAdmin.bat -deploy -srcDir *IS_install_path*\ASBServer\apps\lib\iis\15properties -libName iis -libModule 15properties

   - **For WebSphere Application Server Network Deployment stand-alone configuration:**

     AIX Linux

     IS_install_path/ASBServer/bin/iisAdmin.sh -deploy -srcDir *IS_install_path*/ASBServer/apps/lib/iis/classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password*

     Windows

     IS_install_path/ASBServer\bin\iisAdmin.bat -deploy -srcDir *IS_install_path*\ASBServer\apps\lib\iis\classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password*

   - **For WebSphere Application Server Network Deployment clustered configuration:**

     AIX Linux

     IS_install_path/ASBServer/bin/iisAdmin.sh -deploy -srcDir *IS_install_path*/ASBServer/apps/lib/iis/classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password* -cluster *cluster_name*

     Windows

     IS_install_path/ASBServer\bin\iisAdmin.bat -deploy -srcDir *IS_install_path*\ASBServer\apps\lib\iis\classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password* -cluster *cluster_name*

5. Restart the application server.

Application server configuration	Procedure
Stand-alone	Restart WebSphere Application Server
Clustered	Restart the WebSphere Application Server cluster members

Configuring the HADR primary server:

On the primary server, configure the high availability disaster recovery (HADR) variables and verify your settings.

Before you begin

Complete the installation of IBM InfoSphere Information Server. The installation must be operational before you complete the HADR configuration.

Complete HADR configuration as soon as the InfoSphere Information Server installation is operational.

Procedure

1. Log in to the primary server as an administrator.
2. Run the following commands to specify the HADR primary server and standby server:

   ```
   db2 update db cfg for database using hadr_local_host primary_IP
   db2 update db cfg for database using hadr_remote_host standby_IP
   ```

   The *primary_IP* address and the *secondary_IP* address can be an IPv4 address or an IPv6 address. For example:

   ```
   db2 update db cfg for xmeta using hadr_local_host 192.0.2.6  | ::ffff:192.0.2.6
   db2 update db cfg for xmeta using hadr_remote_host 192.0.2.7 | ::ffff:192.0.2.7
   ```

3. Run the following commands to specify the HADR primary and standby service names:

   ```
   db2 update db cfg for database using hadr_local_svc primary_svc_name
   db2 update db cfg for database using hadr_remote_svc standby_svc_name
   ```

   In the commands, the service names must match the HADR port names that are used in the `services` file. For example:

   ```
   db2 update db cfg for xmeta using hadr_local_svc ha_myhadr
   db2 update db cfg for xmeta using hadr_remote_svc ha_myhadr_int
   ```

4. Run the following command to specify the DB2 instance name on the standby server:

   ```
   db2 update db cfg for database using hadr_remote_inst DB2_standby_instance
   ```

   For example:

   ```
   db2 update db cfg for xmeta using hadr_remote_inst db2inst1
   ```

5. Run the following command to specify the time (in seconds) that the HADR process waits before considering a communication attempt to have failed:

   ```
   db2 update db cfg for database using hadr_timeout timeout
   ```

   For example:

   ```
   db2 update db cfg for xmeta using hadr_timeout 120
   ```

6. Run the following command to specify the HADR synchronization mode:

   ```
   db2 update db cfg for database using hadr_syncmode mode
   ```

   For example:

   ```
   db2 update db cfg for xmeta using hadr_syncmode SYNC
   ```

   When configuring the synchronization mode, choose synchronous mode (SYNC) or near synchronous mode (NEARSYNC). Either mode typically provides adequate response time for InfoSphere Information Server. Avoid choosing asynchronous mode (ASYNC) because ASYNC mode has the highest chance of transaction loss if the primary server fails.

7. Run the following command to specify the time (in seconds) that the HADR process waits before considering a communication attempt to have failed:

   ```
   db2 update db cfg for database using hadr_timeout timeout
   ```

For example:

```
db2 update db cfg for xmeta using hadr_timeout 120
```

8. Run the following command to activate the "log index pages created" feature. When logindexbuild is on, index creation, recreation, and reorganization operations are logged. This logging allows indexes to be reconstructed during DB2 roll-forward operations or HADR log replay procedures.

   ```
   db2 update db cfg for xmeta using logindexbuild on
   ```

9. Run the following command to specify that active log files be retained and become online archive log files for use in rollforward recovery.

   ```
   db2 update db cfg for xmeta using logarchmeth1 logretain immediate
   ```

10. Run the following command to set up automatic client reroute:

    ```
    db2 update alternate server for database database using hostname standby_IP
       port port
    ```

    For example:

    ```
    db2 update alternate server for database xmeta using hostname
      192.0.2.7 | ::ffff:192.0.2.7
      port 60000
    ```

11. Run the following command to verify your configuration values:

    ```
    db2 get db cfg for database | grep HADR
    ```

    For example:

    ```
    db2 get db cfg for xmeta | grep HADR
    ```

    The command displays a report:

    ```
    HADR database role                                    = STANDARD
    HADR local host name              (HADR_LOCAL_HOST) = myprimary.example.com
    HADR local service name            (HADR_LOCAL_SVC) = ha_myhadr
    HADR remote host name            (HADR_REMOTE_HOST) = mystandby.example.com
    HADR remote service name          (HADR_REMOTE_SVC) = ha_myhadr_int
    HADR instance name of remote server  (HADR_REMOTE_INST) = db2inst1
    HADR timeout value                       (HADR_TIMEOUT) = 120
    HADR log write synchronization mode     (HADR_SYNCMODE) = SYNC
    ```

Configuring the HADR standby server:

Back up the database from the primary server and restore it on the standby server. Then, set HADR variables on the standby server.

Before you begin

Make sure that neither database has an existing connection with client application. You cannot perform a backup on a database that has an existing connection with its client application.

Procedure

1. Log in to the primary server as an administrator.
2. Run these commands to back up the database on the primary server:

   ```
   mkdir temp_backup_directory
   cd temp_backup_directory
   db2 backup db database
   ```

   For example:

```
mkdir mytempbackup
cd mytempbackup
db2 backup db xmeta
```

Note: You can store more than one database backup in the same temporary directory.
After the backup completes, a confirmation message is displayed:

```
Backup successful. The timestamp for this backup image is :20101101161943
```

3. Transfer the backup files from the primary server to the standby server.
4. Log in to the standby server. Use your DB2 user ID.
5. Run this command to restore the database on the standby server:

   ```
   db2 restore db database from database_backup_directory replace history file
   ```

   For example:

   ```
   db2 restore db xmeta from mytemprestore replace history file
   ```

6. On the standby server, run the following commands to specify the HADR primary server and standby server. In the command, specify the standby server for `hadr_local_host`. Specify the primary server for `hadr_remote_host` (the opposite of the settings on the primary server).

   ```
   db2 update db cfg for database using hadr_local_host standby_IP
   db2 update db cfg for database using hadr_remote_host primary_IP
   ```

   The *primary_IP* address and the *secondary_IP* address can be an IPv4 address or an IPv6 address. For example:

   ```
   db2 update db cfg for xmeta using hadr_local_host 192.0.2.7  | ::ffff:192.0.2.7
   db2 update db cfg for xmeta using hadr_remote_host 192.0.2.6 | ::ffff:192.0.2.6
   ```

7. On the standby server, run the following commands to specify the HADR primary and standby service names. In the command, specify the standby server port name for `hadr_local_svc`. Specify the primary server port name for `hadr_remote_svc` (the opposite of the settings on the primary server).

   ```
   db2 update db cfg for database using hadr_local_svc standby_svc_name
   db2 update db cfg for database using hadr_remote_svc primary_svc_name
   ```

   For example:

   ```
   db2 update db cfg for xmeta using hadr_local_svc ha_myhadr_int
   db2 update db cfg for xmeta using hadr_remote_svc ha_myhadr
   ```

8. On the standby server, run the following command to set up automatic client reroute. Specify the primary server IP address and port number in the command:

   ```
   db2 update alternate server for database database using hostname primary_IP
      port port
   ```

   For example:

   ```
   db2 update alternate server for database xmeta using hostname
     192.0.2.6 | ::ffff:192.0.2.6
     port 60000
   ```

9. On the standby server, run the following command to verify your configuration values:

   ```
   db2 get db cfg for database | grep HADR
   ```

   For example:

   ```
   db2 get db cfg for xmeta | grep HADR
   ```

   The command displays a report:

```
HADR database role                                    = STANDARD
HADR local host name                (HADR_LOCAL_HOST) = mystandby.example.com
HADR local service name              (HADR_LOCAL_SVC) = ha_myhadr_int
HADR remote host name              (HADR_REMOTE_HOST) = myprimary.example.com
HADR remote service name            (HADR_REMOTE_SVC) = ha_myhadr
HADR instance name of remote server  (HADR_REMOTE_INST) = db2inst1
HADR timeout value                       (HADR_TIMEOUT) = 120
HADR log write synchronization mode     (HADR_SYNCMODE) = SYNC
```

Starting HADR:

Start high availability disaster recovery (HADR) on the standby server, and then on the primary server.

Before you begin

Make sure to configure the HADR variables as described in the previous information. If the parameters are not configured correctly, the HADR servers do not start.

Procedure

1. On the standby server, run the following command to deactivate the database:
   ```
   db2 deactivate db database
   ```
   For example:
   ```
   db2 deactivate db xmeta
   ```
2. On the standby server, run the following command to start HADR on the standby server:
   ```
   db2 start hadr on db database as standby
   ```
   For example:
   ```
   db2 start hadr db xmeta as standby
   ```
   The following message is displayed:
   ```
   DB2000I  The START HADR ON DATABASE command completed successfully
   ```
3. On the standby server, run the following command to verify the HADR role of the server:
   ```
   db2 get snapshot for db on database | grep Role
   ```
   For example:
   ```
   db2 get snapshot for db on xmeta  | grep Role
   ```
   The following line is displayed:
   ```
   Role                    = Standby
   ```
4. On the primary server, run the following command to activate the database:
   ```
   db2 activate db database
   ```
   For example:
   ```
   db2 activate db xmeta
   ```
5. On the primary server, run the following command to start HADR on the primary server:
   ```
   db2 start hadr on db database as primary
   ```
   For example:

```
db2 start hadr db xmeta as primary
```

The following message is displayed:

```
DB2000I  The START HADR ON DATABASE command completed successfully
```

6. On the primary server, run the following command to verify the HADR role of the server:

   ```
   db2 get snapshot for db on database | grep Role
   ```

   For example:

   ```
   db2 get snapshot for db on xmeta  | grep Role
   ```

   The following line is displayed:

   ```
   Role                    = Primary
   ```

7. On each server, run the following command to verify that the databases are synchronized:

   ```
   db2 get snapshot for database on database | grep State
   ```

   For example:

   ```
   db2 get snapshot for database on xmeta  | grep State
   ```

   After the standby database makes the connection to the primary database, the state is reported as Peer:

   ```
   State                   = Peer
   ```

Validating the HADR configuration:

After you configure high availability disaster recovery (HADR), validate your configuration and test the takeover function.

Before you begin

Configure HADR on your primary and standby servers, and start HADR on the servers.

About this task

After verifying that you are connected, initiate a takeover and a failback to test the takeover function.

Procedure

1. Verify the connection status, for example, by using the IBM DB2 Control Center.
   a. Locate the database in the left pane of the DB2 Control Center.
   b. Right-click the database name and select **High Availability Disaster Recovery** > **Manage**. The Manage High Availability Disaster Recovery (HADR) window is displayed.
   c. Verify that the **Connection status** field reads `Connected`.
2. Log in to the primary server as an administrator.
3. On the primary server, run the following command to initiate a DB2 HADR takeover:

   ```
   db2 takeover hadr on db database
   ```

   For example:

```
db2 takeover hadr on db xmeta
```

When the takeover is complete, the following result is displayed:

```
DB2000I The TAKEOVER HADR ON DATABASE command completed successfully.
```

4. Log in to the standby server as an administrator.
5. On the standby server, run the following command to determine the status of the database:

   ```
   db2 get snapshot for db on database | grep Role
   ```

   For example:

   ```
   db2 get snapshot for db on xmeta | grep Role
   ```

   When the takeover is successful, the following result is displayed for the standby server:

   ```
   Role                    = Primary
   ```

6. Check the IBM WebSphere Application Server log file `SystemOut.log` to determine whether the application server successfully connected to the standby database. To locate this file, see "Application server system log files" on page 407.

   If automatic client reroute worked correctly, the log file indicates that WebSphere Application Server first tried to connect to the primary server. When that connection failed, automatic client reroute connected successfully to the standby server. The following log file excerpt shows these events:

   ```
   [11/1/11 17:15:39:298 CST] 00000039 ServletWrappe E   SRVE0068E: Uncaught
   exception thrown in one of the service methods of the servlet: /dbview.jsp.
   Exceptionthrown : javax.servlet.ServletException: A connection failed but
   has been re-established. The hostname or IP address is "myprimary.example.com"
   and the service name or port number is 60000 . Special registers may or may
   not be re-attempted (Reason code = 1  DB2ConnectionCorrelator:
   G9056D89.O37F.061101231714
   ```

7. Log in to the standby server as an administrator.
8. On the standby server, run the following command to initiate a DB2 HADR failback:

   ```
   db2 takeover hadr on db database
   ```

   For example:

   ```
   db2 takeover hadr on db xmeta
   ```

   When the failback is complete, the following result is displayed:

   ```
   DB2000I The TAKEOVER HADR ON DATABASE command completed successfully.
   ```

9. On the primary server, run the following command to determine the status of the database:

   ```
   db2 get snapshot for db on database | grep Role
   ```

   For example:

   ```
   db2 get snapshot for db on xmeta | grep Role
   ```

   When the takeover is successful, the following result is displayed for the primary server:

   ```
   Role                    = Primary
   ```

10. On each server, run the following command to verify that the databases are synchronized:

    ```
    db2 get snapshot for database on database | grep State
    ```

For example:

```
db2 get snapshot for database on xmeta  | grep State
```

After the standby database makes the connection to the primary database, the state is displayed as Peer:

```
State                = Peer
```

Creating the IBM InfoSphere QualityStage Match Designer database

The IBM InfoSphere QualityStage Match Designer is a component of InfoSphere QualityStage that is used to design and test match specifications. Match specifications consist of match passes that identify duplicate entities in one or more files.

If you intend to develop match specifications, use a new or existing database that is dedicated to storing the results of match test passes. Do not use the metadata repository or the analysis database to store these results.

Database requirements

The InfoSphere Information Server installation program does not create the Match Designer results database. You create the database before or after the installation, as long as the database is configured and accessible when you use the Match Designer. You can create the database on a computer where the client or engine tier is installed or on any computer that is accessible to both of these tiers. You must configure the database to receive the type of data that is processed in the Match Designer. For example, you must configure the database to receive double-byte data if the Match Designer processes Asian data.

The following table lists the supported databases and describes the configuration requirements:

Table 62. Configuration requirements by database

Database	Configuration requirements
All databases	Configure as an empty database. The Match Designer creates the table structure that it requires when it stores the match pass results in the database.
IBM DB2 database system	Configure the following settings: • Set **Default Buffer Pool and Table Space Page Size** to 32K. • Set **Code Set** to UTF-8.

Database space requirements

After you create the database, use automatic storage management to allocate database space as needed.

To estimate initial space requirements and plan for the appropriate hardware, you need to make some assumptions about the type and volume of matching that you will do. Consider the following factors:

- Estimated number of match specifications
- Average number of match passes per match specification

- Average number of input data records per match specification
- Average number of bytes per input data record
- Average number of input frequency records per match specification

When you calculate the number of bytes in a Match Designer input record, assume that VarChar columns contain maximum width values. Also, keep in mind that most Match Designer input records will be preprocessed through an InfoSphere QualityStage Standardization Stage. This stage will add 500 bytes to a source record or 1000 bytes if the source record is double-byte.

When calculating frequency record volume, consider whether the data values for the source columns that participate in your matches typically have high or low rates of duplication. For data that is highly unique, your frequency record count will be similar to your data record count. For data that has a high rate of duplication, your frequency record count will be significantly less than your data record count.

Review the factors described in Table 2. Then, use the factors in the following formula to estimate how many megabytes of space the Match Designer results database is likely to require.

```
(number of match specifications x (match size + (match size x 10%) +
frequency size) x replication factor) / 1,000,000
```

Table 63. Space requirement factors and calculations

Factor	Description	Calculation
Match specification count	Match specifications define and test criteria for matching data.	Approximate number of match specifications you expect to retain in your database
Match pass count	Match passes define the column-level matching criteria applied during a single pass through match specification input.	Average number of match passes per match specification (used in the Match Size calculation)
Data record count	Data records define the content and format of match specification input.	Average number of input data records per match specification (used in the Match Size calculation)
Data record length	Data records define the content and format of match specification input. When calculating data record length, assume sizing scenarios that require the most space.	Average number of bytes per match specification input data record (used in the Match Size calculation)
Match size (in bytes)	Match size aggregates the match pass and data record count and length factors into the estimated space requirements for the data component of a single match specification.	(Match pass count) x (data record count) x (data record length)

Table 63. Space requirement factors and calculations (continued)

Factor	Description	Calculation
Frequency record count	Frequency records indicate how often a particular value appears in a particular column. For frequency record count, consider high or low rates of duplication for the data values in source columns.	Average number of frequency records per match specification
Frequency size (in bytes)	Frequency size estimates the space requirement for the frequency component of a single match specification.	(Frequency record count per match) x 360
Replication factor	Replication factors account for significant differences in space requirements of different match types: • One-source matches take input from one source, then group and match the data. • Two-source one-to-one matches compare a record on the data source to a record on the reference source. • Two-source many-to-one matches can match any reference source record with many data source records.	Select the replication factor for the match type you expect to run most often: • For one-source matches, use a factor of **1** • For two-source one-to-one matches, use a factor of **2** • For two-source many-to-one matches, use a factor of **5**

Pre-installing the IBM InfoSphere Information Analyzer analysis database

You can install the IBM InfoSphere Information Analyzer analysis database with the IBM InfoSphere Information Server installation program. Alternatively, you can use scripts to manually create the IBM InfoSphere Information Analyzer analysis database in an existing database instance.

Important: IBM certifies and validates InfoSphere Information Server using the provided database creation scripts. If you modify the scripts beyond the changes that are specified here, IBM will not support software failures that might result from the modifications. Modifications to these scripts that do not have a functional impact might require that you make similar changes in the future to scripts that are generated by other administrative tools that depend on the information in the scripts here. You are responsible for making those changes.

Pre-installing the IBM InfoSphere Information Analyzer analysis database in a DB2 database system by using scripts (Windows):

You can use an existing installation of a DB2 database to create an IBM InfoSphere Information Analyzer analysis database for InfoSphere Information Server on Microsoft Windows.

About this task

Use the scripts in the following situations:

- When you cannot run the installation program locally on the target computer.
- When you do not want to use the InfoSphere Information Server installation program to configure the database.

Attention: The scripts run the following commands when configuring the DB2 database system. These changes affect all databases in the DB2 instance:

```
db2set DB2_INLIST_TO_NLJN=YES
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON
db2set DB2_REDUCED_OPTIMIZATION=YES
```

Procedure

1. Navigate to the *install_dir*\DatabaseSupport directory of the InfoSphere Information Server installation media, where *install_dir* is the directory where you extracted the InfoSphere Information Server media.

Database version	Directory
DB2 10.5	*install_dir*\DatabaseSupport\Windows\InformationAnalyzer\DB2_LUW

2. Open the create_ia_db.sql script file in a text editor and make the following changes:

   Important: IBM certifies and validates InfoSphere Information Server using the provided database creation scripts. If you modify the scripts beyond the changes that are specified here, IBM will not support software failures that might result from the modifications. Modifications to these scripts that do not have a functional impact might require that you make similar changes in the future to scripts that are generated by other administrative tools that depend on the information in the scripts here. You are responsible for making those changes.

   a. Replace all occurrences of @DATABASE_NAME@ (including the @ symbols) with the name of the database that you want to create. For example, replace @DATABASE_NAME@ with your database name:

      ```
      iadb
      ```

   b. Replace all occurrences of @INSTALL_ROOT@ (including the @ symbols) with the path to the root directory where your DB2 database is installed. The default root directory is C:\. For example, replace @INSTALL_ROOT@ with the path to the root directory for your DB2 database:

      ```
      C:\
      ```

      When you run the script, DB2 creates a directory under the root directory where your database is installed. This directory contains the data files for your database. The default directory is C:\db2\NODE0000.

   c. Replace all occurrences of @DATABASE_ALIAS@ (including the @ symbols) with the name of the database that you want to create. For example, replace @DATABASE_ALIAS@ with your database name:

      ```
      iadb
      ```

   d. Replace all occurrences of @DATABASE_USERNAME@ (including the @ symbols) with the name of the operating system user that you want to have access to the database. For example, replace @DATABASE_USERNAME@ with the operating system user name:

```
iauser
```

3. Run the operating system script to invoke the SQL script:
   a. Login as instance user.
   a. Open a DB2 Command Window.
   b. From the DB2 Command Window, change directory to the same location as the edited script and run the following command:

      ```
      create_ia_db.cmd [logfile]
      ```

      logfile is the name of the file where you want to save the results of running the command. This argument is optional.
4. Restart DB2.

Pre-installing the IBM InfoSphere Information Analyzer analysis database in a DB2 database system by using scripts (AIX Linux):

You can use an existing installation of a DB2 database to create an IBM InfoSphere Information Analyzer analysis database for InfoSphere Information Server.

About this task

The scripts run the following commands when configuring the DB2 database system. These changes affect all databases in the DB2 instance:

```
db2set DB2_INLIST_TO_NLJN=YES
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON
db2set DB2_REDUCED_OPTIMIZATION=YES
```

Procedure

1. Navigate to the *install_dir*\DatabaseSupport directory of the InfoSphere Information Server installation media, where *install_dir* is the directory where you extracted the InfoSphere Information Server installation media.

Database version	Directory
DB2 10.5	*install_dir*\DatabaseSupport\UNIX_Linux\InformationAnalyzer\DB2_LUW

2. Open the create_ia_db.sql script file in a text editor and make the following changes:

   Important: IBM certifies and validates InfoSphere Information Server using the provided database creation scripts. If you modify the scripts beyond the changes that are specified here, IBM will not support software failures that might result from the modifications. Modifications to these scripts that do not have a functional impact might require that you make similar changes in the future to scripts that are generated by other administrative tools that depend on the information in the scripts here. You are responsible for making those changes.
   a. Replace all occurrences of @DATABASE_NAME@ (including the @ symbols) with the name of the database that you want to create. For example, replace @DATABASE_NAME@ with your database name:

      ```
      iadb
      ```

   b. Replace all occurrences of @INSTALL_ROOT@ (including the @ symbols) with the path to the root directory where your DB2 database is installed. For example, replace @INSTALL_ROOT@ with the path to the root directory for your DB2 database:

```
/home/db2inst1/db/iadb
```

c. Replace all occurrences of `@DATABASE_ALIAS@` (including the @ symbols) with the name of the database that you want to create. For example, replace `@DATABASE_ALIAS@` with your database name:

```
iadb
```

d. Replace all occurrences of `@DATABASE_USERNAME@` (including the @ symbols) with the name of the operating system user that you want to have access to the database. For example, replace `@DATABASE_USERNAME@` with the operating system user name:

```
iauser
```

3. Run the operating system script to invoke the SQL script:
   a. Log on as the DB2 instance owner.
   b. Change to the directory where the edited script is located, and run the following command:

```
create_ia_db.sh [logfile]
```

logfile is the name of the file where you want to save the results of running the command. This argument is optional.

4. Restart DB2.

Setting up a preinstalled application server

For a new installation in a non-clustered environment, you might use the IBM InfoSphere Information Server installation program to install IBM WebSphere Application Server Network Deployment or IBM WebSphere Application Server Liberty Profile. If you choose to use WebSphere Application Server Network Deployment, you can use an existing supported version of the application server or preinstall it yourself.

If you want to implement a cluster for high availability, failover, and scalability, then you must use a preinstalled and configured version of WebSphere Application Server Network Deployment. You cannot opt to use an application server that is installed by the InfoSphere Information Server installation program.

Preinstalling WebSphere Application Server Network Deployment

IBM InfoSphere Information Server supports IBM WebSphere Application Server Network Deployment as a preinstalled application server. If you choose to preinstall an application server, it must be WebSphere Application Server Network Deployment and not IBM WebSphere Application Server Liberty Profile.

Before you begin

The WebSphere Application Server installation must satisfy the following requirements:

- It must be on the same computer where you install the services tier.
- It must be a supported version. See the system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.

InfoSphere Information Server does not support WebSphere Application Server Extended Deployment (XD) or the WebSphere Application Server base version.

About this task

This task describes installing IBM WebSphere Application Server Network Deployment in a non-cluster (stand-alone) installation.

For a cluster installation, see Implementing a WebSphere Application Server cluster for InfoSphere Information Server.

Note: If you preinstall IBM WebSphere Application Server and install a separate front-end HTTP server, ensure that the HTTP server is configured to enable WebSphere Application Server to accept HTTPS calls.

Procedure

1. Run the IBM WebSphere Application Server installation program on the InfoSphere Information Server services tier to install IBM WebSphere Application Server Network Deployment. Within the installation media, the WebSphere Application Server installation program is at the following location:

Operating system	Directory
AIX	*install_image_dir*/payloads/Shared/WAS/payload/WAS/agentInstaller/install
Linux	*install_image_dir*/payloads/Shared/WAS/payload/WAS/agentInstaller/install
Windows	*install_image_dir*\payloads\Shared\WAS\payload\WAS\agentInstaller\install.exe

   See the WebSphere Application Server documentation for more information on installing the application server:
   - For Version 8.5.5.1, go to the WebSphere Application Server Network Deployment information center and read "Installing and configuring your application".
2. Upgrade the JDK if it is not at the correct level. See "Switching to JDK 1.7" on page 162.
3. Create a new application server profile with administrative security disabled. For information about creating a profile, see the WebSphere Application Server documentation:
   - For Version 8.5.5.1, go to the WebSphere Application Server Network Deployment information center and read "Managing profiles on non-z/OS operating systems".
4. If the front-end HTTP server is not already configured for HTTPS, manually configure it for HTTPS. Follow the instructions in the WebSphere Application Server documentation:
   - For Version 8.5.5.1, go to the WebSphere Application Server Network Deployment information center and read "Securing communications".
5. Install any required fix packs and interim fixes. To find the supported fix pack level, Java SDK fix number, and interim fixes that are required for your installed version of IBM WebSphere Application Server Network Deployment, refer to the InfoSphere Information Server system requirements page: www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.

Note: If you use the installation packages for the IBM WebSphere Application Server Network Deployment software that is included in the installation package, the fix packs and interim fixes are automatically installed when you run the installation program.

Setting up an existing non-clustered installation of WebSphere Application Server

You can configure an existing IBM WebSphere Application Server installation to use with InfoSphere Information Server.

Before you begin

The WebSphere Application Server installation must satisfy the following requirements:

- It must be on the same computer where you install the services tier.
- It must be a supported version. See the system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.

InfoSphere Information Server does not support WebSphere Application Server Extended Deployment (XD) or the WebSphere Application Server base version.

Procedure

Configure your WebSphere Application Server installation before you install InfoSphere Information Server.

1. At a command prompt, enter one of the following commands to verify that you have a supported version of WebSphere Application Server and the Java SDK.

Operating system	Command
AIX Linux	`/opt/IBM/WebSphere/AppServer64/bin/ versionInfo.sh` `/opt/IBM/WebSphere/AppServer64/java/bin/ java -fullversion`
Windows	`C:\IBM\WebSphere\AppServer\bin\ versionInfo.bat` `C:\IBM\WebSphere\AppServer\bin\ java -fullversion`

   Compare the values that are returned by these commands to the version numbers in the system requirements to determine if you have a supported version. If you do not have a supported version, do not attempt to configure it for use with InfoSphere Information Server. Upgrade the JDK if it is not at the correct level. See "Switching to JDK 1.7" on page 162.
2. Create a new application server profile with administrative security disabled. For information about creating a profile, see the WebSphere Application Server documentation:
   - For Version 8.5.5.1, go to the WebSphere Application Server Network Deployment information center and read "Managing profiles on non-z/OS operating systems".
3. If the front-end HTTP server is not already configured for HTTPS, manually configure it for HTTPS. Follow the instructions in the WebSphere Application Server documentation:
   - For Version 8.5.5.1, go to the WebSphere Application Server Network Deployment information center and read "Securing communications".

Implementing a WebSphere Application Server cluster for InfoSphere Information Server

You can set up a clustered environment for high availability, failover, and scalability. First, determine an appropriate cluster topology, and then install and configure the appropriate version of IBM WebSphere Application Server Network Deployment, set up clusters, and install a front-end web server.

To deploy a highly available topology, you must have a solid understanding of network technologies such as protocols, layers, and devices. Highly available topologies, especially topologies that involve clustering, are technically complex and require expertise to implement successfully.

You choose an appropriate cluster topology based on your system requirements (level of high availability needed, level of security needed, hardware equipment available, and so on). The cluster topology can include one or more clustered application servers, each running cloned copies of IBM InfoSphere Information Server, one or more front-end web servers (or hardware or software load balancers), one or more firewalls, and so on.

Refer to theWebSphere Application Server documentation for information about cluster configurations and manual tuning:

- For Version 8.5.5.1, go to the WebSphere Application Server information center and read “Introduction: Clusters”.

General considerations for clustering on WebSphere Application Server

IBM WebSphere Application Server Network Deployment provides control over the design of the cluster environment through the deployment manager, managed nodes, and unmanaged nodes (stand-alone InfoSphere Information Server instances):

- You can define multiple Network Deployment cells with their own deployment manager.
- You can install multiple deployment managers on the same computer.
- You can install a deployment manager separately on a dedicated computer, or a deployment manager can share the same computer with another InfoSphere Information Server managed node or stand-alone InfoSphere Information Server instance.

When you create a cluster, be aware of these InfoSphere Information Server restrictions:

- All managed nodes that are participating in a cluster must run on the same operating system. For example, in a cluster, one managed node cannot be running on a Windows system and another managed node running on a UNIX system.
- There is no need to have a WebSphere Application Server cell dedicated to InfoSphere Information Server, however, the following restrictions do apply:
  - The cluster that is dedicated to InfoSphere Information Server cannot contain any other applications in it.
  - The cell where the cluster resides cannot have any prior InfoSphere Information Server applications in it.

Installing IBM WebSphere Application Server Network Deployment for a cluster installation:

IBM InfoSphere Information Server supports IBM WebSphere Application Server Network Deployment, Version 8.5.5.1. You must install a supported version.

Before you begin

1. Determine the cluster topology that you want to use, based on your system requirements.
2. On the computer where you want to install the deployment manager and on all computers that will host a managed node, configure the system to support a large number of file descriptors.

   Refer to your system administrator if you are not familiar with file descriptor allocations.

   For example, if your login shell is `/bin/bash`, complete the following tasks:

Operating system	Procedure
AIX	Configure the following environment variable in `/etc/profile` for system-wide change: `unset LDR_CNTRL`
Linux	1. Add the following line to the `/etc/profile` file: `ulimit -n 10240` 2. Add the following lines to the `/etc/security/limits.conf` file: `* soft nofile 10240` `* hard nofile 10240`

Procedure

1. Install IBM WebSphere Application Server Network Deployment on all the computers that will host the InfoSphere Information Server services tier. This installation includes the deployment manager system and all computers that will host a managed node.

   Run the IBM WebSphere Application Server installation program to install IBM WebSphere Application Server Network Deployment. See the WebSphere Application Server documentation for more information:

   - For Version 8.5.5.1, go to the WebSphere Application Server information center and read “Installing and configuring your application”.

   The InfoSphere Information Server installation package includes installation packages for IBM WebSphere Application Server Network Deployment, Version 8.5.5.1. To preinstall WebSphere Application Server, Version 8.5.5.1 before you run the InfoSphere Information Server installation program, you can run the WebSphere Application Server installation program that is included in the InfoSphere Information Server installation media.

   Within the installation media, the WebSphere Application Server installation program is at the following location:

Operating system	Directory
AIX Linux	payloads/Shared/WAS/payload/WAS/agentInstaller/install
Windows	payloads\Shared\WAS\payload\WAS\agentInstaller\install.exe

When creating a profile during the installation, make sure to disable administrative security.

2. Upgrade the JDK if it is not at the correct level. See Switching to JDK 1.7.
3. Install the required fix pack and interim fixes. To find the supported fix pack level, Java SDK fix number, and interim fixes that are required for your installed version of IBM WebSphere Application Server Network Deployment, refer to the InfoSphere Information Server system requirements page: publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/softwareReqsForProduct.html

Creating a deployment manager profile:

You create a deployment manager profile (for example, *Dmgr01*) that contains the deployment manager process. The deployment manager process is the main administration process that manages all other IBM WebSphere Application Server processes running in the cell, including node agents and application server processes.

Procedure

Create a deployment manager profile by using the Profile Management Tool or the **manageprofiles** command line tool. The Profile Management Tool, if you installed it, and the **manageprofiles** command line tool are located in the following directories of a WebSphere Application Server installation.

Operating system	Directory
AIX Linux	**The Profile Management Tool** *WAS_HOME*/bin/ProfileManagement/pmt.sh **The manageprofiles command line tool** *WAS_HOME*/bin/manageprofiles.sh where *WAS_HOME* is the location of the application server installation, such as /opt/IBM/WebSphere/AppServer.
Windows	**The Profile Management Tool** *WAS_HOME*\bin\ProfileManagement\pmt.bat **The manageprofiles command line tool** *WAS_HOME*\bin\manageprofiles.bat where *WAS_HOME* is the location of the application server installation, such as C:\IBM\WebSphere\AppServer.

See the IBM WebSphere Application Server Network Deployment documentation for more information:

Method	Procedure
Using the Profile Management Tool • On the **Administrative Security** panel, select **Enable administrative security** and provide user credentials. Credentials for the built-in WebSphere Application Server file registry can be used. • On the **Environment Selection** panel, select **Deployment Manager**. • For most cases, choose the **Typical profile creation** option. If you want to modify default attributes of the deployment manager profile (such as the port numbers or whether to run the deployment manager process as a Microsoft Windows service), select the **Advanced profile creation** option.	• For Version 8.5.5.1, go to the WebSphere Application Server Network Deployment information center and read "Creating management profiles with deployment managers"
Using the `manageprofiles` command line tool	• For Version 8.5.5.1, go to the WebSphere Application Server Network Deployment information center and read "manageprofiles command"

Example

The following example illustrates how you use the **manageprofiles** command to create a deployment manager profile in a UNIX environment:

```
manageprofiles.sh -create -templatePath profileTemplates/dmgr/ -isDefault
  -profileName dmgr -adminUserName wasadmin -adminPassword Pa$$w0rd
```

What to do next

Create a custom profile.

Creating a custom profile:

You create a custom profile (for example, *Custom01*) that contains the node agent process and the managed server process that is part of the cluster.

About this task

It is a good practice to install IBM InfoSphere Information Server on a single-node cluster and then, after the installation, to expand your cluster by adding new managed nodes and cluster members. Refer to the post-installation tasks for details about adding new managed nodes and cluster members.

Procedure

Create a custom profile by using the Profile Management Tool or the **manageprofiles** command line tool. The Profile Management Tool is located in the `bin/ProfileManagement` directory of a WebSphere Application Server installation. See the IBM WebSphere Application Server Network Deployment documentation for more information:

Method	Procedure
Using the Profile Management Tool[1]	• For Version 8.5, go to the WebSphere Application Server Network Deployment information center and read "Creating custom profiles"
Using the `manageprofiles` command line tool	• For Version 8.5, go to the WebSphere Application Server Network Deployment information center and read "manageprofiles command"

[1]If you use the Profile Management Tool, adhere to the following guidelines when creating a custom profile:

- On the Federation panel, ensure that the host name and port information match your deployment manager (the default information is localhost and 8879). Make sure that you do not select **Federate this node later** check box. Specify the deployment manager user name and password.
- For most cases, choose the **Typical profile creation** option. If you want to modify default attributes of the deployment manager profile (such as the port numbers), select the **Advanced profile creation** option.

Example

AIX : The following example illustrates how you use the **`manageprofiles`** command to create a custom profile:

```
manageprofiles.sh -create -templatePath profileTemplates/managed/
-profileName <profile_name> -dmgrHost <dmgr_host_name> -dmgrPort
<dmgr_port_number> -adminUserName <admin_user_name> -adminPassword
<admin_password>
```

What to do next

Create a cluster and cluster members.

Creating a cluster and cluster members for the application server:

Use the IBM WebSphere administrative console to create a cluster with one cluster member definition.

Before you begin

1. If you use HTTP session persistence with database-based persistence, you must configure the application server separately before you can create a cluster.
2. Create a managed server (which is not yet part of a cluster) and configure it to use HTTP session database persistence.
3. After you configure the HTTP session with database-based persistence enabled, you can create a cluster and use the newly configured managed server as the initial cluster member.
4. You must choose one of the following two methods available to manage HTTP sessions in a cluster environment: HTTP session memory-to-memory replication or HTTP session database-based persistence. Both methods ensure that HTTP session data survive in the event of a cluster member failure.

Note: HTTP session memory-to-memory replication is easier to set up. If you want to use HTTP session memory-to-memory replication, when creating the cluster in

IBM WebSphere Application Server Network Deployment administrative console, in the Step1: Enter basic cluster information window, select the **Configure HTTP session memory-to-memory replication** check box. By default, the IBM InfoSphere Information Server installation process detects whether HTTP sessions were already configured with database-based session persistence. If it is not yet configured, HTTP session memory-to-memory replication is configured automatically, even if you do not select this check box.

About this task

Create one cluster member to start with. After InfoSphere Information Server is installed, you can create additional cluster members. Refer to the post-installation information for more details on how to add new cluster members.

Remember: The cluster selected for InfoSphere Information Server cannot contain other applications deployed in it, and the cell cannot contain any prior release of InfoSphere Information Server. However, the cell need not be dedicated to InfoSphere Information Server and the cell can contain multiple profiles, nodes, and clusters. For more information, refer to "Implementing a WebSphere Application Server cluster for InfoSphere Information Server" on page 151.

Procedure

Create clusters and cluster members by using the IBM WebSphere Application Server Network Deployment administrative console.

- For Version 8.5.5.1, go to the WebSphere Application Server information center and read "Creating clusters."

Replication domains:

When you create clusters and cluster members, it is useful to understand the various replications domains that are used to create clusters.

Several replication domains are used within an IBM InfoSphere Information Server cluster:

- If you chose a memory-to-memory replication approach for HTTP session management, the replication domain replicates HTTP session information and the stateful EJB state across the cluster.
- Another type of replication domain is used to replicate cached data (data-cached applications that are part of the InfoSphere Information Server suite).

Both types of replication domains replicate cached data to all cluster members (a *full-domain* strategy). Replication domains are created automatically during the InfoSphere Information Server installation. You use the default IBM WebSphere Application Server settings during installation. You can change these settings from the WebSphere Application Server administrative console after the InfoSphere Information Server installation.

Installing a front-end web server:

You install a front-end web server before you run the IBM InfoSphere Information Server installation program.

Before you begin

Determine which web server and web server version to install, depending on your version of IBM WebSphere Application Server Network Deployment.

Review the documentation:

- For Version 8.5.5.1, go to the WebSphere Application Server information center and read "Implementing a web server plug-in."

About this task

The following steps describe how to install the IBM HTTP Server, Versions 8.5.5.1. Refer to the Apache HTTP Server documentation for information about how to install the Apache HTTP Server.

Procedure

1. Determine a web server topology diagram and road map:
   - For Version 8.5.5.1, go to the WebSphere Application Server information center and read "Selecting a web server topology diagram and roadmap."
2. Choose whether to install the web server locally or remotely:
   - For development and testing environments, you might prefer to install the web server locally on a computer that hosts an IBM WebSphere Application Server Network Deployment managed node. When you do so, the web server plug-in file, `plugin-cfg.xml`, is automatically propagated to the web server by IBM WebSphere Application Server.
   - For production environments, install the web server remotely on a dedicated server. Because the front-end web server is typically outside the firewall in the DMZ, the front-end web server must be installed remotely on a separate dedicated computer. You manually copy the `plugin-cfg.xml` file to the computer where the web server is installed. (For more information, see the *IBM InfoSphere Information Server Administration Guide*.)
3. Install IBM HTTP Server:
   - For Version 8.5.5.1, go to the WebSphere Application Server information center and follow the steps for "Installing IBM HTTP Server".

   If you choose to install InfoSphere Information Server by using a cluster member rather than a front-end web server, you can stop here and configure the front-end web server after InfoSphere Information Server installation.
4. Install the web server plug-in module:
   - For Version 8.5.5.1, go to the WebSphere Application Server information center and follow the steps for "Installing web server plug-ins".

   After you install the Workload Manager plug-in binary files, run the `configure<webserver_name>` script to create a web server definition in WebSphere Application Server.

Operating system	Script
AIX Linux	`configure<webserver_name>.sh`
Windows	`configure<webserver_name>.bat`

Note: For remote installations, running the script also creates an unmanaged node.

5. Configure the web server to accept encoded slashes. See “Cannot view report result from InfoSphere Information Server console or InfoSphere Information Server Web console” on page 429.
6. Configure SSL for the IBM HTTP Server, and configure SSL between the front-end IBM HTTP Server and the application server:
   - For Version 8.5.5.1, go to the WebSphere Application Server information center under “Securing IBM HTTP Server” and complete the following tasks:
     – “Securing with SSL communications”
     – “Configure SSL between the front-end IBM HTTP Server and the deployment manager”
7. If you are using a port other than the default 80 or 443 for your front-end web server, add a virtual host entry with the same port value as your front-end web server.
   - For Version 8.5.5.1, see Configuring virtual hosts in the WebSphere Application Server Version 8.5 information center.

What to do next

- Ensure there are no errors seen in the front end web server and the WebSphere Application Server plugin log files. If there are any, they need to fixed before you can install InfoSphere Information Server.
- You can install additional components such as backup load balancers, backup web servers, firewalls, and cache proxies. Refer to the external vendors' documentation.

Installing a front-end load balancer (optional):

You install a load balancer in front of the web servers in an IBM WebSphere Application Server cluster installation if your installation requires very high availability. This configuration is called an IP sprayer topology.

About this task

Steps for installing a front-end load balancer are specific to external vendors and are not provided in this documentation.

Do not install a load balancer directly in front of a WebSphere Application Server cluster without a web server between the load balancer and the cluster. Web servers with the WebSphere Application Server workload management plug-in adjust to changes in the cluster topology such as adding or removing nodes. Load balancers do not have this functionality, and must be manually reconfigured.

Make sure to configure session affinity in the load balancer. There are several methods to achieve session affinity. The method that you choose depends on the product that you use and the forwarding mechanism that you configure. Mechanisms include media access control (MAC) address forwarding, Network Access Translation/Network Address Port Translation (NAT/NAPT) forwarding, and content-based forwarding. Layer 2 load balancers (load balancers that operate at the MAC layer) generally achieve session affinity by using a “stickiness to source IP address” approach. Layer 3 and higher load balancers (load balancers that operate at the IP layer and above) typically achieve session affinity by using passive cookies or a content-based technique. Session affinity with HTTP servers is typically achieved by using these techniques.

What to do next

After you install a load balancer, you can install any additional components that you need, such as backup load balancers and firewalls.

Configuring a cluster to run as a service (Windows):

You can configure an IBM WebSphere Application Server cluster to run as a service. When you run a WebSphere Application Server cluster as a service, the cluster automatically restarts after a failure and keeps processes active after you log out of Microsoft Windows.

About this task

Running a WebSphere Application Server cluster as a service includes the following benefits:

- Enables automatic restarts in the event of a failure
- Enables cluster member processes to remain active after you log out from your Microsoft Windows system account

Procedure

1. Configure your node agent or node agents, to run as a service by using the **WASService** command line tool.
   - For Version 8.5.5.1, go to the WebSphere Application Server information center and read "WASService command".

   The following example code uses the Local System account, SYSTEM, as the default user ID. If you want to use a different account, you can specify your own user ID by using the **-userid** option.

```
C:\IBM\WebSphereND855\AppServer\bin>WASService
-add "myNode01" -servername nodeagent
-profilePath C:\IBM\WebSphereND855\AppServer\profiles\Custom01
-wasHome C:\IBM\WebsphereND855\Appserver
-logFile C:\IBM\WebSphereND855\AppServer\profiles\Custom01\
 logs\nodeagent\startNode.log
-logRoot C:\IBM\WebSphereND855\AppServer\profiles\Custom01\logs\nodeagent
-restart true -startType automatic

Adding Service: ImyNode01
        Config Root: C:\IBM\WebSphereND855\AppServer\profiles\Custom01\config
        Server Name: nodeagent
        Profile Path: C:\IBM\WebSphereND855\AppServer\profiles\Custom01
        Was Home: C:\IBM\WebsphereND855\Appserver\
        Start Args:
        Restart: 1
IBM WebSphere Application Server V8.5.5 - myNode01 service successfully added.
```

2. Configure your cluster members to restart automatically:
   a. Log in to the WebSphere administrative console.
   b. Click **Servers** > **Application Servers** > *server_name*.
   c. Expand **Java and Process Management** and select **Monitoring Policy**.
   d. Change the node restart state to **RUNNING**.
   e. Click **OK and Save**.

   Repeat this step for each server in the cluster.
3. Shut down your node agents by using the **stopNode** command:
   - For Version 8.5.5.1, go to the WebSphere Application Server information center and read "stopNode command".

4. Shut down your cluster by using the WebSphere administrative console.
   - For Version 8.5.5.1, go to the WebSphere Application Server information center and read “Stopping clusters”.
5. Start your node agents by using the Windows Services Console. Your cluster will automatically restart.

   Note: Your node agents and cluster members must now run in the context of the Local System account instead of your own user account, unless you specified a different account when you ran the **WASService** command in step 1 on page 159
   You can now safely log out of your Windows user account and your cluster will continue to run.

Running the installation program (clustered installation):

After you configure your system, run the IBM InfoSphere Information Server installation program.

Before you begin

- **AIX Linux**: Ensure sure that all WebSphere Application Server processes will start with the correct ulimit setting on the computer where you will install the deployment manager and on all computers that will host a managed node:

  Add the following line to the `/etc/profile` file:

  ```
  ulimit -n 10240
  ```
- Ensure that no firewall is deployed within the IBM WebSphere Application Server Network Deployment cell.
- Ensure that the deployment manager process is running.

  Tip: You can check whether the deployment manager is running by starting the WebSphere Application Server administrative console. If the login screen displays correctly, the deployment manager process is running. If you see an error or a blank screen, you can start the deployment manager by using the **startManager** command located under the *<dmgrProfileName>*`/bin` directory.
- Ensure that the managed node and node agent that you created are both running.

  Tip: You can check whether the node agent is running by logging in to the WebSphere Application Server administrative console and clicking **System administration** > **Node agents**. You can start the node agent by using the **startNode** WebSphere Application Server command under the *<customProfileName>*`/bin` directory.
- Verify that you have created one cluster with at least one member.

  Note: Install IBM InfoSphere Information Server on a single-node cluster and then, after the installation, expand your cluster by adding new managed nodes and cluster members. Refer to the post-installation administration tasks for details on how to add new custom profiles.
- If you choose to install against a front-end web server rather than a cluster member, the front-end web server and plug-in module must be installed and configured.

Procedure

1. Disable any firewall that might be running on your computer and ensure that the firewall is disabled during the entire installation. You can reenable the firewall after the installation is finished.
2. Run the InfoSphere Information Server installation program on the computer that hosts the deployment manager process.
   a. Use the steps in the following documentation: Installing IBM InfoSphere Information Server software.
   b. On the WebSphere Application Server Selection panel, select **Configure an existing instance**.
3. Select the WebSphere Application Server deployment manager profile. The installation program detects that you have a clustered installation rather than a stand-alone installation.
4. During the interview process, after you select a profile, you select the cluster, web server and virtual host. You specify to install against a select cluster member or a front-end web server. You have a choice to install using front end web server or install using a cluster member. It is recommended to install against a cluster member unless you are very sure of the following prerequisites:
   - The front-end web server is properly configured, and the certificates at the WebSphere Cell level are imported into the front-end web server and to the web server plug-in's keystore (if applicable).
   - The web server port numbers are specified into WebSphere Virtual Hosts aliases.
   - There are no firewall issues accessing the front-end web server from the machine installing IBM InfoSphere Information Server.

   Further, during IBM InfoSphere Information Server installation, a set of IBM InfoSphere Information Server web applications are deployed into the selected cluster. Some of these web applications are also invoked with REST APIs as part of the installation process. That means that if you select the front-end web server as an installation option, you must make sure that it is pre-configured to properly propagate HTTPS requests to the back end WebSphere Application Server so that these requests are run by the new web application. Or, you must ensure that the front-end web server is fully managed by WebSphere Application Server with the automatic update and propagation of the HTTP plug-in when new web applications are deployed. If you are not sure about the front-end web server configuration and its propagation procedure, select to install against the cluster member.

What to do next

- You configure your cluster with InfoSphere Information Server. For more information, see “WebSphere Application Server cluster configuration” on page 280
- After you configure WebSphere Application Server for your cluster installation of InfoSphere Information Server, you can administer your WebSphere Application Server clusters. For more information, see the *IBM InfoSphere Information Server Administration Guide*.

Note: When the MetadataServer script is run in a stand-alone (non-cluster) environment, it sets the following system settings:

Operating system	Setting
AIX	`uname setting: ulimit -n 10240`
Linux	`uname setting: ulimit -n 10240`
Windows	`uname setting: unset LDR_CNTRL`

These settings are not set when you start a cluster. You can manually produce these settings in a cluster environment, if needed. The MetadataServer script exists in the `ASBServer/bin` directory of the computer hosting the deployment manager.

Switching to JDK 1.7

By default IBM WebSphere Application Server Network Deployment comes with Java Development Kit (JDK) 1.6. IBM InfoSphere Information Server requires JDK 1.7. You can install JDK 1.7 and use the **managesdk** command to switch the version of JDK that is used by WebSphere Application Server Network Deployment.

Procedure

1. On the services tier computer, install JDK 1.7 to the location of your choice.
2. Set JDK 1.7 to be used by the InfoSphere Information Server profile:
   - If you have not yet created a profile for InfoSphere Information Server, you can set JDK 1.7 to be used as the default for all profiles (The ⇒ symbol indicates a line continuation):

     AIX Linux

     ```
     WAS_install_path/bin/managesdk.sh -setNewProfileDefault ⇒
     -sdkname 1.7_64
     WAS_install_path/bin/managesdk.sh -setCommandDefault ⇒
     -sdkname 1.7_64
     ```

     Windows

     ```
     WAS_install_path\bin\managesdk.bat -setNewProfileDefault ⇒
     -sdkname 1.7_64
     WAS_install_path\bin\managesdk.bat -setCommandDefault ⇒
     -sdkname 1.7_64
     ```

   - If you already created a profile for InfoSphere Information Server, run the following commands to enable the profile to use JDK 1.7 (The ⇒ symbol indicates a line continuation):

     AIX Linux

     ```
     WAS_install_path/bin/managesdk.sh -enableProfile -profileName ⇒
     profile_name -sdkname 1.7_64 -enableServers
     WAS_install_path/bin/managesdk.sh -setCommandDefault -sdkname ⇒
     1.7_64
     ```

     Windows

     ```
     WAS_install_path\bin\managesdk.bat -enableProfile -profileName ⇒
     profile_name -sdkname 1.7_64 -enableServers
     WAS_install_path\bin\managesdk.bat -setCommandDefault -sdkname ⇒
     1.7_64
     ```

3. Verify that JDK 1.7 is now used by the profile:
   - You can check by logging into the WebSphere Application Server Administrative Console and selecting the following options: **Application servers** > **server1** > **Java SDKs**
   - You can check by using the **managesdk** command (The ⇒ symbol indicates a line continuation):

     AIX Linux

     ```
     WAS_install_path/bin/managesdk.sh -listEnabledProfile ⇒
     -profileName profile_name
     ```

Windows

```
WAS_install_path\bin\managesdk.bat -listEnabledProfile ⇒
-profileName profile_name
```

Sample output from the command:

```
CWSDK1004I: Profile profile_name:
CWSDK1006I: PROFILE_COMMAND_SDK = 1.7_64
CWSDK1008I: Node ssominenNode01 SDK name: 1.7_64
CWSDK1009I: Server server1 SDK name: 1.7_64
CWSDK1001I: Successfully performed the requested managesdk task.
```

Setting the locale and character encoding for the engine tier target computer

The installation program uses the locale and encoding information of the target computer to determine whether it can provide language support. For languages other than English, you must ensure that the locale and encoding are set correctly.

About this task

When you install IBM InfoSphere Information Server with globalization (NLS) support enabled, the installation program chooses the globalization configuration that is the most suitable for the locale and encoding of the target computer. To ensure that server jobs run and compile correctly, make sure that the target computer uses a supported locale and encoding. Set the encoding to UTF-8 to support a wide variety of languages.

If you do not set the target computer to use a supported locale and encoding, the installation program configures InfoSphere Information Server for the English language, including globalization support.

Complete the following procedure to set the locale. For information about setting the locale in Windows, see the Microsoft documentation.

Procedure

1. Log in to the computer.
2. Enter the following command to determine the current locale:

   ```
   locale
   ```

3. Enter the following command to list all available locales:

   ```
   locale -a
   ```

   The locale names that are returned by **locale** and **locale -a** might be full locale names or alias locale names. A full locale name contains the encoding name in the last component. For example, the full locale name ja_JP.utf8 contains the encoding name, utf8, in the last component. An alias locale name does not include the encoding name. To determine the encoding, you can run the following command:

   ```
   LANG=alias_name locale charmap
   ```

   where *alias_name* is the alias name. For example, the following command returns utf8 on a system where the locale ja_JP uses the utf8 encoding:

   ```
   LANG=ja_JP locale charmap
   ```

4. If the target computer does not use a supported locale and encoding, change the LANG value to a combination that is available on your computer and that is supported by IBM InfoSphere Information Server. To change the locale and encoding, run the **export** command with the following format:

```
export LANG={locale.charmap}
export LC_ALL={locale.charmap}
```

For example, the following commands change the locale to ja_JP and the encoding to UTF-8:

```
export LANG=ja_JP.utf8
export LC_ALL=ja_JP.utf8
```

Character encodings supported by IBM InfoSphere Information Server

To ensure that server jobs run and compile correctly, the operating system of the target computer must use an encoding that is supported by the globalization configuration that is used during installation. Verify that the operating system of the target computer uses a supported character encoding.

The following table lists the languages and encodings that are supported by InfoSphere Information Server on various operating systems. For more information, see the *IBM InfoSphere DataStage and QualityStage Globalization Guide*.

Table 64. Supported encodings on AIX computers for InfoSphere Information Server

Language	Globalization configuration name	Supported encodings
Chinese	DS-CHINESE-SIMPLIFIED	**AIX** EUC-CN, GB18030, GB2312, GBK, IBMEUCCN
Chinese	DS-CHINESE-SIMPLIFIED-UTF8	UTF-8
Chinese, Traditional	DS-CHINESE-TRADITIONAL	BIG5
Chinese, Traditional	DS-CHINESE-TRADITIONAL-UTF8	UTF-8
English	DS-ENGLISH	ISO8859-1
English	DS-ENGLISH-UTF8	UTF-8
French	DS-FRENCH	ISO8859-1
French	DS-FRENCH-UTF8	UTF-8
German	DS-GERMAN	ISO8859-1
German	DS-GERMAN-UTF8	UTF-8
Italian	DS-ITALIAN	ISO8859-1
Italian	DS-ITALIAN-UTF8	UTF-8
Japanese	DS-JAPANESE	
Japanese	DS-JAPANESE-EUC	Not supported
Japanese	DS-JAPANESE-SJIS	IBM932, IBM943, IBM943C, PCK, SHIFT_JIS, SJIS
Japanese	DS-JAPANESE-UTF8	UTF-8
Korean	DS-KOREAN	**AIX** KSC5601
Korean	DS-KOREAN-UTF8	UTF-8
Portuguese	DS-PORTUGUESE-BRAZILIAN	ISO8859-1
Portuguese	DS-PORTUGUESE-BRAZILIAN-UTF8	UTF-8
Spanish	DS-SPANISH	ISO8859-1
Spanish	DS-SPANISH-UTF8	UTF-8
Thai	DS-THAI	TIS620

Table 65. Supported encodings on Linux computers for InfoSphere Information Server

Language	Globalization configuration name	Supported encodings
Chinese	DS-CHINESE-SIMPLIFIED	EUC-CN, GB18030, GB2312, GBK, IBMEUCCN
Chinese	DS-CHINESE-SIMPLIFIED-UTF8	UTF-8
Chinese, Traditional	DS-CHINESE-TRADITIONAL	BIG5
Chinese, Traditional	DS-CHINESE-TRADITIONAL-UTF8	UTF-8
English	DS-ENGLISH	ISO8859-1
English	DS-ENGLISH-UTF8	UTF-8
French	DS-FRENCH	ISO8859-1
French	DS-FRENCH-UTF8	UTF-8
German	DS-GERMAN	ISO8859-1
German	DS-GERMAN-UTF8	UTF-8
Italian	DS-ITALIAN	ISO8859-1
Italian	DS-ITALIAN-UTF8	UTF-8
Japanese	DS-JAPANESE	
Japanese	DS-JAPANESE-EUC	EUC-JP, EUC, IBM33722, IBM33722C, IBMEUCJP, UJIS, XEUCJPOPEN
Japanese	DS-JAPANESE-SJIS	IBM932, IBM943, IBM943C, PCK, SHIFT_JIS, SJIS
Japanese	DS-JAPANESE-UTF8	UTF-8
Korean	DS-KOREAN	EUC, EUC-KR
Korean	DS-KOREAN-UTF8	UTF-8
Portuguese	DS-PORTUGUESE-BRAZILIAN	ISO8859-1
Portuguese	DS-PORTUGUESE-BRAZILIAN-UTF8	UTF-8
Spanish	DS-SPANISH	ISO8859-1
Spanish	DS-SPANISH-UTF8	UTF-8
Thai	DS-THAI	TIS620

Table 66. Supported encodings on Microsoft Windows computers for InfoSphere Information Server

Language	Globalization configuration name	Supported encodings
Chinese	DS-CHINESE-SIMPLIFIED	MS936
Chinese	DS-CHINESE-SIMPLIFIED-UTF8	UTF-8
Chinese, Traditional	DS-CHINESE-TRADITIONAL	MS950
Chinese, Traditional	DS-CHINESE-TRADITIONAL-UTF8	UTF-8
English	DS-ENGLISH	MS1252
English	DS-ENGLISH-UTF8	UTF-8
French	DS-FRENCH	MS1252
French	DS-FRENCH-UTF8	UTF-8
German	DS-GERMAN	MS1252
German	DS-GERMAN-UTF8	UTF-8
Italian	DS-ITALIAN	MS1252

Table 66. Supported encodings on Microsoft Windows computers for InfoSphere Information Server (continued)

Language	Globalization configuration name	Supported encodings
Italian	DS-ITALIAN-UTF8	UTF-8
Japanese	DS-JAPANESE	MS932, MS31J
Japanese	DS-JAPANESE-EUC	Not supported
Japanese	DS-JAPANESE-SJIS	
Japanese	DS-JAPANESE-UTF8	UTF-8
Korean	DS-KOREAN	MS949
Korean	DS-KOREAN-UTF8	UTF-8
Portuguese	DS-PORTUGUESE-BRAZILIAN	MS1252
Portuguese	DS-PORTUGUESE-BRAZILIAN-UTF8	UTF-8
Spanish	DS-SPANISH	MS1252
Spanish	DS-SPANISH-UTF8	UTF-8
Thai	DS-THAI	PC874

Creating a two-server active-passive high availability configuration

You can provide a basic level of high availability by configuring the IBM InfoSphere Information Server software tiers in an active-passive high availability configuration.

In a two-server active-passive high availability configuration, two computers share a storage device. The metadata repository tier, engine tier, and services tier are all installed on the shared storage device. One of the computers (the active server) hosts the tiers. The other computer (the passive server) is started, but is not running any tier software. A high availability software product is installed on both servers, such as IBM Tivoli System Automation for Multiplatforms on Linux or UNIX systems or Microsoft Cluster Service (MSCS) on Microsoft Windows systems. The high availability software maintains a *heartbeat*, which is a periodic signal from the active server to the passive server that indicates that the active server is operational. If the active server fails, the heartbeat also fails. The high availability software restarts all the services on the passive server. This process is called a *failover*.

With this configuration, a floating IP address and virtual host name are defined for the server pair. This information is associated with the active server. All client programs connect to the server by using this address or host name. If the active server fails, the floating IP address and virtual host name are automatically reassociated with the passive server.

If you provide high availability for the services tier or metadata repository tier (or both) by using other methods, you can choose to include only the remaining tiers in the active-passive configuration. For example, you might choose to implement an IBM WebSphere Application Server cluster for your services tier and an IBM DB2 Database for Linux, UNIX, and Windows cluster for your metadata repository tier. In this case, the servers in your active-passive configuration host only the engine tier.

To set up a high availability configuration, you must have a solid understanding of network technologies such as protocols, layers, and devices. Highly available

configurations, especially configurations that involve clustering, are technically complex and require expertise to implement them successfully.

Creating an active-passive configuration (AIX Linux)

Regardless of the high availability software that you use, you must complete the same basic steps to configure a two-server active-passive configuration.

Procedure

1. Set up the hardware and cable the network.

   The servers must be physically separate. Use identical hardware with identical internal storage, memory, and processing power.

   The servers must be able to communicate with each other over high-speed networks. To minimize single points of failure, each server must have redundant connections to at least two physically separate high-speed networks.

   To support failover when a problem occurs, the high availability software creates a heartbeat signal between the servers. Set up at least two more physically separate networks for this signal. These networks can be serial or other networks. The networks that carry the heartbeat signal must be physically separate and isolated from the high-speed networks.
2. Install the high availability software on both the active server and the passive server.
3. Allocate an IP address and associate it with the virtual host name that the client tier will use to connect to the IBM InfoSphere Information Server server-side components.
4. Mount the shared storage device in identical mount points on each node. Create the following mount points:
   - Base installation directory. Typically `/opt/IBM`
   - Base user home directory. Typically `/home`
5. Use the appropriate high availability software administration tools to complete the following steps:
   a. Create a resource group for the InfoSphere Information Server resources.
   b. Add a resource to manage the virtual IP address.
   c. Start the resource group.
   d. Ensure that the resource group is running on the primary server on which InfoSphere Information Server is to run. Ensure that you can connect to the primary server by using the virtual host name. Test failover and resolve any issues before continuing.
6. On both servers, create the directories in the following table within the mount points of these directories. Ensure the directories have identical local user permissions. Create the dsadm, db2inst1, db2fenc1, xmeta, and iauser user IDs with the same names, passwords, group names, group IDs, and home directories on both servers. The names of the directories might be different than indicated in the table, depending on your configuration.

Table 67. Directories for active-passive configuration

Configuration of tiers	Directory
All	Project directory. Default: `/opt/IBM/InformationServer/Server/Projects`

Table 67. Directories for active-passive configuration (continued)

Configuration of tiers	Directory
Engine tier on both servers	• InfoSphere Information Server installation directory. Default: `/opt/IBM/InformationServer` • Home directory for the InfoSphere Information Server engine administrator. Typically `/home/dsadm`
Services tier on both servers	IBM WebSphere Application Server installation directory. Default: `/opt/IBM/WebSphere`
Metadata repository tier on both servers, IBM DB2 database system	• DB2 database system installation directory. Default: `/opt/IBM/db2` • DB2 instance owner home directory. Typically `/home/db2inst1` • DB2 fenced user home directory. Typically `/home/db2fenc1` • Metadata repository owner home directory. Typically `/home/xmeta` • If IBM InfoSphere Information Analyzer is installed, the home directory of the owner of the analysis databases. Typically `/home/iauser`

7. If the active-passive configuration will include only the engine tier (that is, the metadata repository or services tiers are to be located outside of the active-passive configuration), install and configure the metadata repository and services tiers first, before you continue with the engine tier installation and configuration. The engine tier installation requires that the metadata repository tier and services tier are installed first, unless you install them all at the same time in a single pass of the installation program.
8. Run the InfoSphere Information Server installation program on the primary server.

   When prompted by the installation program, specify the virtual host name. Whenever prompted for an installation directory, specify a path within the mount point.
9. Replicate the `/etc/services` entries on both servers.
10. Create scripts to start, stop, and monitor the servers. Store the scripts on the shared storage device.

    Several sample high availability control scripts are included with InfoSphere Information Server. You can use these scripts or modify them to fit your configuration. See "High availability control scripts included with IBM InfoSphere Information Server" on page 169.
11. Use the appropriate high availability software administration tools to complete the following steps:
    a. If the engine tier is included in the high availability configuration, add the `InfoSvrEngine` script or its equivalent resource agent script to the resource group.
    b. If the services tier or metadata repository tier is included in the high availability configuration, add the `InfoSvrServices` script or its equivalent resource agent script to the resource group.
12. Test the setup by simulating different failures.

High availability control scripts included with IBM InfoSphere Information Server:

Several sample high availability control scripts are included with InfoSphere Information Server. To enable your high availability software to start or stop tiers and components, set up these files as application resources within the high availability software.

These scripts are installed by the InfoSphere Information Server installation program in the following directory:

Operating system	Directory
AIX	*installdir*/InformationServer/HAScripts
Linux	*installdir*/InformationServer/HAScripts

Within your high availability configuration, you can use these scripts as they are, or modify them so that they work in your environment. The following table lists and describes the scripts that are available:

Table 68. Sample scripts that are supplied with InfoSphere Information Server

Script	Description
InfoSvrEngine	Starts or stops the InfoSphere Information Server engine tier. Calls ISFAgents as part of its operation.
InfoSvrServices	Starts or stops the InfoSphere Information Server services tier. Calls ISFRepos and ISFServer scripts as part of its operation.
ISFAgents	Calls the NodeAgents.sh script to start or stop InfoSphere Information Server ASB node agents.
ISFServer	Calls the MetadataServer.sh script to start or stop IBM WebSphere Application Server.
ISFRepos	Called by the InfoSvrServices script in an active-passive configuration where an IBM DB2 database system hosts the metadata repository tier. Starts or stops a local DB2 database system instance. The DB2 database system must run on the same server as other InfoSphere Information Server components. Assumes that the DB2 database system is installed on shared storage that is mounted identically on both servers.

The set of scripts you use depends on which tiers you configure in the active-passive high-availability configuration. The following table lists the scripts to use based on the tiers that you include in your configuration:

Table 69. Tiers and resource agent scripts

Tier	Sample resource agent scripts
Engine	• InfoSvrEngine • ISFAgents

Table 69. Tiers and resource agent scripts (continued)

Tier	Sample resource agent scripts
Services	• `InfoSvrServices` • `ISFServer`
Metadata repository	• `InfoSvrServices` • `ISFRepos`

Example: Setting up an active-passive configuration by using IBM Tivoli System Automation for Multiplatforms:

This example describes how to set up a two-server active-passive high availability configuration by using Tivoli System Automation for Multiplatforms.

This example configuration is built on two Red Hat Enterprise Linux Server servers. The servers are connected to each other over a single public network. Each server is also connected to a storage area network (SAN). All IBM InfoSphere Information Server components are installed on the SAN. At any given time, only one of the servers is active; that is, it hosts the InfoSphere Information Server components on the SAN. The other server is passive and does not access the SAN.

Although this example is for Linux, most steps are identical in an AIX environment. Steps that differ are indicated.

The following diagram illustrates the example configuration:

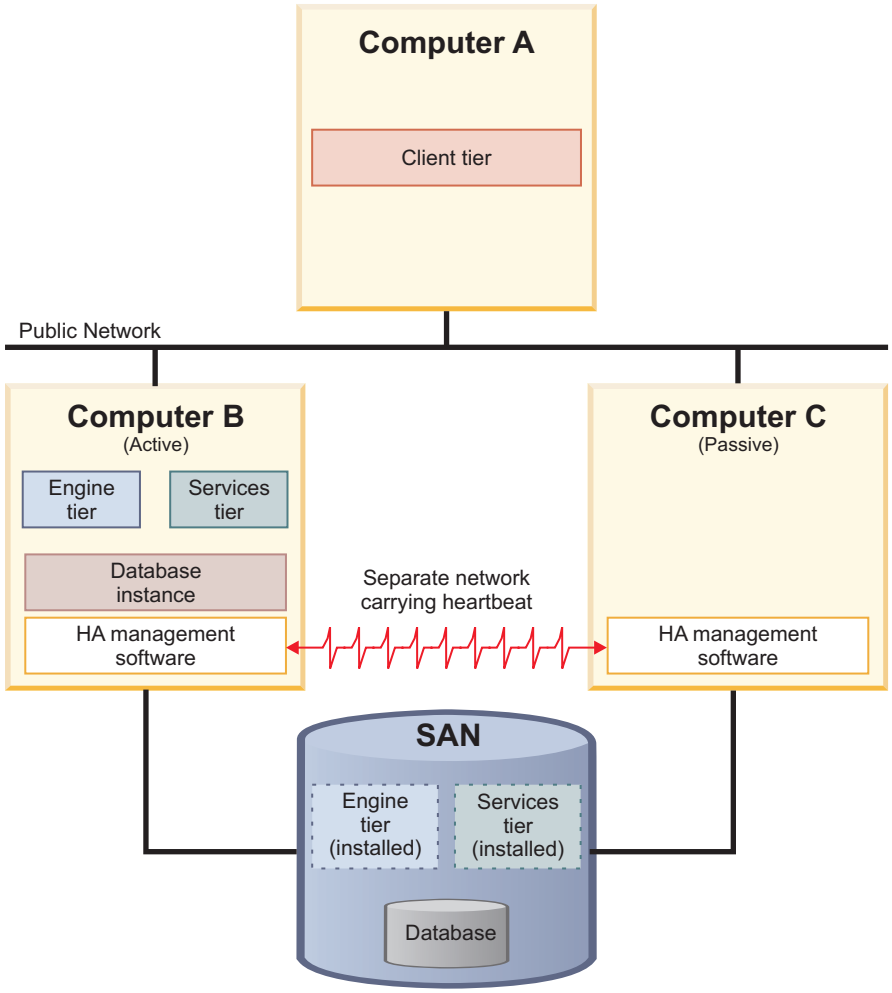


Figure 40. Example of a two-server active-passive high availability configuration

In this example, the following software is used for each server:

- Red Hat Enterprise Linux Server 64-bit
- IBM InfoSphere Information Server
- IBM WebSphere Application Server Network Deployment that is installed with the InfoSphere Information Server installation program
- IBM DB2 Database for Linux, UNIX, and Windows that is installed with the InfoSphere Information Server installation program
- IBM Tivoli System Automation for Multiplatforms that is included in the IBM DB2 installation package

For storage, there is a storage area network (SAN) that is accessible to both servers.

To create the configuration in this example, follow these steps:

1. Set up the hardware. Then configure the network, and configure mount points to the storage area network (SAN). Finally, mount the SAN. See "Setting up the servers, network, and storage (Tivoli System Automation for Multiplatforms example)" on page 172.
2. Install System Automation for Multiplatforms on each server in the cluster, and then create the first level automation domain. System Automation for Multiplatforms is bundled with the IBM DB2 database system installation

package. See "Installing the high availability software (Tivoli System Automation for Multiplatforms example)" on page 173.

3. Define the virtual IP address. Then, configure System Automation for Multiplatforms so that the network adapters on the servers are appropriate and equivalent for the virtual IP address. See "Setting up the virtual IP address (Tivoli System Automation for Multiplatforms example)" on page 174.
4. Define the resource group, and then create the automation policy by using relationship definitions. Create a network tiebreaker. Finally, activate the resource group. See "Defining and starting the resource group (Tivoli System Automation for Multiplatforms example)" on page 175.
5. Create users and user groups on the SAN. Then, run the InfoSphere Information Server installation program. Finally, replicate the `/etc/services` entries on both servers. See "Installing IBM InfoSphere Information Server in the cluster (Tivoli System Automation for Multiplatforms example)" on page 178.
6. Set up scripts to start, stop, and monitor the servers. Then, create the application resources. See "Setting up the high availability control scripts (Tivoli System Automation for Multiplatforms example)" on page 179.
7. To test the configuration, simulate active server failures and network failures and observe the results. See "Testing the configuration (Tivoli System Automation for Multiplatforms example)" on page 180.

Setting up the servers, network, and storage (Tivoli System Automation for Multiplatforms example):

Set up the hardware. Then configure the network, and configure mount points to the storage area network (SAN). Finally, mount the SAN.

About this task

The two servers in this example contain one network interface each (en0). The en0 network interfaces are connected to each other through the external network cloud that forms the public network. The following static IP addresses are assigned to the en0 adapters on the active and passive servers:

- Active server (host name `myprimary`): 192.0.2.0
- Passive server (host name `mysecondary`): 192.0.2.1

In the following examples, you can express either of the IPv4 addresses as IPv6 addresses. For example, you can express `192.0.2.0` as `::ffff:192.0.2.0`.

Procedure

1. Set up the servers, network, and storage hardware.
2. Adjust kernel parameters and user limits identically on both servers. For information about these adjustments, see "Modifying kernel parameters and user limits (AIX Linux)" on page 101.
3. Open the `/etc/hosts` file on each server. In each file, verify that the active and passive server names are mapped to their corresponding public IP addresses. For example:

```
192.0.2.0    myprimary.example.com       myprimary
192.0.2.1    mysecondary.example.com     mysecondary
```

If the host names are defined in the /etc/hosts file, the servers can identify each other if a DNS server failure occurs. All servers in the cluster must have the same entries in this file. In addition, all servers must have static IP addresses.

4. Validate that the active and passive servers can ping each other successfully over the public network: For example:

```
ping myprimary
ping mysecondary
```

5. Create identical directories on each server to be used as the mount points for the shared SAN. Give the directories permissions of 775:

```
mkdir /opt/IBM
chmod 775 /opt/IBM
mkdir /home
chmod 775 /home
```

6. Open the /etc/fstab file on each server.

   AIX: The file name is /etc/filesystems.

7. In each file, define the mount points for the shared SAN. For each mount point, specify the noauto option to prevent the mount point from being automatically mounted on more than one server in the cluster. For example:

```
192.0.2.2:/newfs   /opt/IBM   nfs   defaults   0 0
192.0.2.2:/newfs   /home      nfs   defaults   0 0
```

8. Mount the SAN on each server.
9. Create the following directories, with identical local user permissions:
   - /opt/IBM/InformationServer/Server/Projects
   - /opt/IBM/InformationServer
   - /opt/IBM/WebSphere
   - /opt/IBM/db2

Installing the high availability software (Tivoli System Automation for Multiplatforms example):

Install IBM Tivoli System Automation for Multiplatforms on each server in the cluster, and then create the first level automation domain. System Automation for Multiplatforms is bundled with the IBM DB2 database system installation package.

Procedure

1. Install the System Automation for Multiplatforms base component on each server. The installation directory must be on the local storage on each server, not on the storage area network (SAN).
2. Edit the /etc/profile file and ensure that the *CT_MANAGEMENT_SCOPE* environment variable is set to 2 for all users.
3. On each server that is to be part of the automation domain, run the **preprpnode** command to prepare security on the server so that it can be defined in a peer domain.

```
/usr/sbin/rsct/bin/preprpnode myprimary mysecondary
```

4. Create the automation domain. In the example, a single automation domain named ha_domain is created that contains the myprimary and mysecondary servers. To create the domain, run the **mkrpdomain** command on one of the servers (you can choose either server):

```
mkrpdomain ha_domain myprimary mysecondary
```

After this command is run once on a server, any Tivoli System Automation for Multiplatforms or Reliable Scalable Cluster Technology (RSCT) command can be issued from any server in the automation domain.

5. The newly created automation domain has an operational state (OpState) of offline. Run the **startrpdomain** command to bring the domain online:

```
startrpdomain ha_domain
```

Setting up the virtual IP address (Tivoli System Automation for Multiplatforms example):

Define the virtual IP address. Then, configure IBM Tivoli System Automation for Multiplatforms so that the network adapters on the servers are appropriate and equivalent for the virtual IP address.

About this task

In the following examples, a file named `myserviceip.cfg` is created that defines a virtual IP address named `haip`. The `IPAddress` value is the network alias IP address, which can be an IPv4 address or an IPv6 address. In the following examples, you can express the IPv4 address `192.0.2.10` as the IPv6 address `::ffff:192.0.2.10`.

Procedure

1. Create a configuration file in the `/opt/IBM/hascripts` directory that defines the virtual IP address and specifies the names of the servers to associate with the address. The `NetMask` value is the network mask of this IP address:

```
PersistentResourceAttributes::
     NodeNameList={"myprimary", "mysecondary"}
     Name="haip"
     NetMask=255.255.255.128
     IPAddress=192.0.2.10
     ResourceType=1
```

2. Run the **mkrsrc** command to create the virtual IP resource:

```
mkrsrc -f /opt/IBM/hascripts/myserviceip.cfg IBM.ServiceIP
```

3. Run the **lsrsrc** command to verify that the resource was created correctly:

```
lsrsrc -l IBM.ServiceIP
```

The command displays the virtual IP resource configuration. For example:

```
Resource Persistent Attributes for IBM.ServiceIP
resource 1:
        Name             = "haip"
        ResourceType     = 0
        AggregateResource = "0x2029 0xffff 0x8657b47c 0x9cbcbccb 0x918205be 0xc296f798"
        IPAddress        = "192.0.2.10"
        NetMask          = "255.255.255.128"
        ProtectionMode   = 1
        NetPrefix        = 0
        ActivePeerDomain  = "ha_domain"
        NodeNameList     = {"mysecondary"}
resource 2:
        Name             = "haip"
        ResourceType     = 0
        AggregateResource = "0x2029 0xffff 0x8657b47c 0x9cbcbccb 0x918205be 0xc296f798"
        IPAddress        = "192.0.2.10"
        NetMask          = "255.255.255.128"
        ProtectionMode   = 1
        NetPrefix        = 0
        ActivePeerDomain  = "ha_domain"
        NodeNameList     = {"myprimary.example.com"}
resource 3:
        Name             = "haip"
        ResourceType     = 1
```

```
        AggregateResource = "0x3fff 0xffff 0x00000000 0x00000000 0x00000000 0x00000000"
        IPAddress         = "192.0.2.10"
        NetMask           = "255.255.255.128"
        ProtectionMode    = 1
        NetPrefix         = 0
        ActivePeerDomain  = "ha_domain"
        NodeNameList      = {"myprimary.example.com","mysecondary"}
```

4. Run the **mkequ** command to create an equivalency definition that specifies the network adapters that are to carry the virtual IP address. For example, the following command creates a static equivalency that is named ha_nieq:

   ```
   mkequ ha_nieq IBM.NetworkInterface:eth0:myprimary,eth0:mysecondary
   ```

5. Run the **lsequ** command to verify the equivalency definition:

   ```
   lsequ -e ha_nieq
   ```

   The command displays the equivalency definition configuration. For example:

   ```
   Displaying Equivalency information:
   For Equivalency "ha_nieq".

   Equivalency 1:
           Name                      = ha_nieq
           MemberClass               = IBM.NetworkInterface
           Resource:Node[Membership] = {eth0:myprimary.example.com,eth0:mysecondary}
           SelectString              = ""
           SelectFromPolicy          = ANY
           MinimumNecessary          = 1
           Subscription              = {}
           Color                     = 0
           ActivePeerDomain          = ha_domain
           ConfigValidity            =
   ```

6. On both servers, in the /etc/hosts file, add the mapping for the virtual host name to the virtual IP address. For example:

   ```
   192.0.2.10    harg  harg.example.com
   ```

Defining and starting the resource group (Tivoli System Automation for Multiplatforms example):

Define the resource group, and then create the automation policy by using relationship definitions. Create a network tiebreaker. Finally, activate the resource group.

Procedure

1. Run the **mkrg** command to create a resource group. In this configuration example, the resource group that is created is named harg:

   ```
   mkrg harg
   ```

2. Run the **addrgmbr** command to add the virtual IP address and application resources to the resource group.

   ```
   addrgmbr -g harg IBM.ServiceIP:haip
   ```

3. Run the **lsrg** command to verify the resource group:

   ```
   lsrg -m
   ```

   The command displays the resource group definition.

   ```
   Displaying Member Resource information:
   Class:Resource:Node[ManagedResource] Mandatory MemberOf OpState WinSource Location
   IBM.ServiceIP:haip                   True      harg     offline  Nominal   myprimary
   ```

4. Run the **mkrel** command to define the relationship between the virtual IP resource (haip) and the network interface equivalency (ha_nieq). The virtual IP depends on the network interface equivalency to be active.

```
mkrel -p DependsOn -S IBM.ServiceIP:haip -G IBM.Equivalency:ha_nieq
   haip_dependson_ha_nieq
```

5. Run the **lsrel** command to verify the relationships:

   ```
   lsrel -l
   ```

   The command displays managed relations.

   ```
   Displaying Managed Relations :

   Managed Relationship 1:
           Name                          = haip_dependson_ha_nieq
           Class:Resource:Node[Source] = IBM.ServiceIP:haip
           ResourceGroup[Source]       = harg
   ```

6. Create a network quorum (tiebreaker). The tiebreaker ensures that Tivoli System Automation for Multiplatforms can decide which node still owns the cluster if communication between the nodes is lost. This safeguard prevents processes from being started on both nodes simultaneously.

   a. Edit the `/usr/sbin/cluster/netmon.cf` file on both servers. In each file, add the IP address of each server in either IPv4 format or IPv6 format. For example, you can use the address `192.0.2.0` or its equivalent IPv6 address `::ffff:192.0.2.0`.

      ```
      cat /usr/sbin/cluster/netmon.cf
      192.0.2.0 | ::ffff:192.0.2.0
      192.0.2.1 | ::ffff:192.0.2.1
      ```

   b. Run the **mkrsrc** command to create the network tie breaker:

      ```
      mkrsrc IBM.TieBreaker Type="EXEC" Name="ha_nettiebkr" DeviceInfo='PATHNAME=
         /usr/sbin/rsct/bin/samtb_net Address=192.0.2.24 Log=1 Count=2'
         PostReserveWaitTime=30;
      ```

      The network address used in the command (192.0.2.24 in the example) can be any valid IP address that responds to the **ping** command from both nodes.

   c. Run the **chrsrc** command to set the persistent attributes for the tie breaker:

      ```
      chrsrc -c IBM.PeerNode OpQuorumTieBreaker="ha_nettiebkr"
      ```

   d. Run the **lsrsrc** command to verify that the tie breaker is created:

      ```
      lsrsrc -Ab IBM.TieBreaker
      ```

      The command displays the attributes for the tie breaker.

      ```
      Resource Persistent and Dynamic Attributes for IBM.TieBreaker
      resource 1:
              Name                = "ha_nettiebkr"
              Type                = "EXEC"
              DeviceInfo          = "PATHNAME=/usr/sbin/rsct/bin/samtb_net
                                     Address=192.0.2.24 Log=1 Count=2"
              ReprobeData         = ""
              ReleaseRetryPeriod  = 0
              HeartbeatPeriod     = 0
              PreReserveWaitTime  = 0
              PostReserveWaitTime = 30
              NodeInfo            = {}
              ActivePeerDomain    = "ha_domain"
              ConfigChanged       = 0
      resource 2:
              Name                = "Fail"
              Type                = "Fail"
              DeviceInfo          = ""
              ReprobeData         = ""
              ReleaseRetryPeriod  = 0
              HeartbeatPeriod     = 0
              PreReserveWaitTime  = 0
      ```

```
        PostReserveWaitTime = 0
        NodeInfo            = {}
        ActivePeerDomain    = "ha_domain"
        ConfigChanged       = 0
resource 3:
        Name                = "Operator"
        Type                = "Operator"
        DeviceInfo          = ""
        ReprobeData         = ""
        ReleaseRetryPeriod  = 0
        HeartbeatPeriod     = 0
        PreReserveWaitTime  = 0
        PostReserveWaitTime = 0
        NodeInfo            = {}
        ActivePeerDomain    = "ha_domain"
        ConfigChanged       = 0
```

7. Run the **chrg** command to activate the resource group and all of the resources that are defined as its members:

   ```
   chrg -o online harg
   ```

8. Run the **lsrg** command to verify the status of the resource group:

   ```
   lsrg -g harg
   ```

   The command displays information for the specified resource group.

   ```
   Displaying Resource Group information:
   For Resource Group "harg".
   Resource Group 1:
           Name                = harg
           MemberLocation      = Collocated
           Priority            = 0
           AllowedNode         = ALL
           NominalState        = Online
           ExcludedList        = {}
           Subscription        = {}
           Owner               =
           Description         =
           InfoLink            =
           Requests            = {}
           ActivePeerDomain    = ha_domain
           OpState             = Online
           TopGroup            = harg
           ConfigValidity      =
           TopGroupNominalState = Online
   ```

9. Run the **lssam** command to check the status of the cluster:

   ```
   lssam
   ```

   The command displays the status of each component of the cluster.

   ```
   Online IBM.ResourceGroup:harg Nominal=Online
           '- Online IBM.ServiceIP:haip
                   |- Online IBM.ServiceIP:haip:myprimary
                   '- Offline IBM.ServiceIP:haip:mysecondary
   ```

10. Test the cluster to verify that it is working correctly.
    a. Verify that the virtual IP address connects to the primary node.
    b. Run the **rgreq** command to force a failover to the secondary node. In the command, specify the name of the secondary node and the resource group name.

       ```
       rgreq -o move -n mysecondary harg
       ```

    c. Run the **lssam** command again to check the status of the cluster:

       ```
       lssam
       ```

If the failover was successful, the command output indicates that the primary node is now offline and the secondary node is online.

```
Online IBM.ResourceGroup:harg Nominal=Online
        '- Online IBM.ServiceIP:haip
                |- Offline IBM.ServiceIP:haip:myprimary
                '- Online IBM.ServiceIP:haip:mysecondary
```

d. Verify that the virtual IP address now connects to the secondary node.

e. Run the **rgreq** command to force a failover back to the primary node.

```
rgreq -o move -n myprimary harg
```

f. Run the **lssam** command again to check the status of the cluster:

```
lssam
```

If the failover was successful, the command output indicates that the secondary node is now offline and the primary node is online.

```
Online IBM.ResourceGroup:harg Nominal=Online
        '- Online IBM.ServiceIP:haip
                |- Online IBM.ServiceIP:haip:myprimary
                '- Offline IBM.ServiceIP:haip:mysecondary
```

Installing IBM InfoSphere Information Server in the cluster (Tivoli System Automation for Multiplatforms example):

Create users and user groups on the storage area network (SAN). Then, run the InfoSphere Information Server installation program. Finally, replicate the `/etc/services` entries on both servers.

About this task

The configuration for this task matches Scenario C, which depicts a two-server active-passive high availability configuration.

Procedure

1. On the storage area network (SAN), create all users and user groups.

   The following sample Linux commands create some of the administrator user accounts and groups. In the example, the SAN is mounted at `/home`:

```
/usr/sbin/groupadd db2iadm1
/usr/sbin/groupadd db2fadm1
/usr/sbin/groupadd xmeta
/usr/sbin/groupadd iauser
/usr/sbin/groupadd dstage

/usr/sbin/useradd -g db2iadm1 -d /home/db2inst1 db2inst1
/usr/sbin/useradd -g db2fadm1 -d /home/db2fenc1 db2fenc1
/usr/sbin/useradd -g xmeta    -d /home/xmeta xmeta
/usr/sbin/useradd -g iauser   -d /home/iauser iauser
/usr/sbin/useradd -g dstage   -d /home/dsadm dsadm
```

   The following sample AIX commands create some of the administrator user accounts and groups:

```
/usr/bin/mkgroup id='202' db2iadm1
/usr/bin/mkgroup id='203' db2fadm1
/usr/bin/mkgroup id='205' xmeta
/usr/bin/mkgroup id='206' iauser
/usr/bin/mkgroup id='207' dstage

/usr/bin/mkuser pgrp=db2iadm1 groups=db2iadm1 home=/home/db2inst1 db2inst1
/usr/bin/mkuser pgrp=db2fadm1 groups=db2fadm1 home=/home/db2fenc1 db2fenc1
```

```
/usr/bin/mkuser pgrp=xmeta    groups=xmeta    home=/home/xmeta    xmeta
/usr/bin/mkuser prgp=iauser   groups=iauser   home=/home/iauser   iauser
/usr/bin/mkuser pgrp=dstage   groups=dstage   home=/home/dsadm    dsadm
```

For details about creating administrator users and groups, see "Setting up administrator and owner accounts" on page 107.

2. Run the InfoSphere Information Server installation program in graphical mode. See Installing InfoSphere Information Server by using the wizard.

   Whenever you are prompted for a directory, choose a directory from the base directory mount point (/opt/IBM).

   On the High Availability Server Cluster Configuration page, click **This is an active-passive server cluster configuration**. In the **Virtual host name** field, enter the virtual host name (harg.example.com in this example).
3. Let the installation program complete.
4. Replicate the /etc/services entries on both servers.

Setting up the high availability control scripts (Tivoli System Automation for Multiplatforms example):

Set up scripts to start, stop, and monitor the servers. Then, create the application resources.

About this task

Several sample high availability control scripts are supplied with IBM InfoSphere Information Server. These scripts are used in this example. See "High availability control scripts included with IBM InfoSphere Information Server" on page 169.

Procedure

1. Copy the script files to a directory on the storage area network (SAN). For example: /opt/IBM/hascripts
2. Create an application resource configuration file for the services tier. Put the file in the same directory as the script files. For example: /opt/IBM/hascripts. The following sample file defines an application resource for the services tier. The file is named /opt/IBM/hascripts/InfoSvrServices.IBM.Application:

   ```
   PersistentResourceAttributes::
       Name="ServicesApp"
       StartCommand="/opt/IBM/hascripts/InfoSvrServices start"
       StopCommand="/opt/IBM/hascripts/InfoSvrServices stop"
       MonitorCommand="/opt/IBM/hascripts/InfoSvrServices status"
       MonitorCommandPeriod=60
       MonitorCommandTimeout=60
       NodeNameList={"myprimary","mysecondary"}
       StartCommandTimeout=240
       StopCommandTimeout=240
       UserName="root"
       ResourceType=1
   ```
3. Run the **mkrsrc** command to create the application resource for the services tier:

   ```
   mkrsrc -f /opt/IBM/hascripts/InfoSvrServices.IBM.Application IBM.Application
   ```
4. Create an application resource configuration file for the engine tier. The following sample file defines an application resource for the engine tier. The file is named /opt/IBM/hascripts/InfoSvrEngine.IBM.Application:

   ```
   PersistentResourceAttributes::
       Name="EngineApp"
       StartCommand="/opt/IBM/hascripts/InfoSvrEngine start"
       StopCommand="/opr/hascripts/InfoSvrEngine stop"
   ```

```
MonitorCommand="/opt/IBM/hascripts/InfoSvrEngine status"
MonitorCommandPeriod=60
MonitorCommandTimeout=60
NodeNameList={"myprimary","mysecondary"}
StartCommandTimeout=240
StopCommandTimeout=240
UserName="root"
ResourceType=1
```

5. Run the **mkrsrc** command to create the application resource for the engine tier:
   ```
   mkrsrc -f /opt/IBM/hascripts/InfoSvrEngine.IBM.Application IBM.Application
   ```
6. Run the **rgreq** command to lock the harg resource group:
   ```
   rgreq -o lock harg
   ```
7. Run the **addrgmbr** command to add the services tier application resource to the resource group harg that was created:
   ```
   addrgmbr -g harg IBM.Application:ServicesApp
   ```
8. Run the **addrgmbr** command to add the engine tier application resource to the resource group harg that was created:
   ```
   addrgmbr -g harg IBM.Application:EngineApp
   ```
9. Run the **mkrel** command to define the relationship between the InfoSphere Information Server services tier application resource (ServicesApp) and the virtual IP resource (haip):
   ```
   mkrel -p StartAfter -S IBM.Application:ServicesApp -G IBM.ServiceIP:haip
      ServicesApp_startafter_haip
   ```
10. Run the **mkrel** command to define the relationship between the InfoSphere Information Server engine tier application resource (EngineApp) and the virtual IP resource (haip):
    ```
    mkrel -p StartAfter -S IBM.Application:EngineApp -G IBM.ServiceIP:haip
       EngineApp_startafter_haip
    ```
11. Run the **lsrel** command to verify the relationships:
    ```
    lsrel -l
    ```

    The command displays all the managed relations:
    ```
    Displaying Managed Relations :


    Managed Relationship 1:
            Name                          = haip_dependson_ha_nieq
            Class:Resource:Node[Source] = IBM.ServiceIP:haip
            ResourceGroup[Source]       = harg

    Managed Relationship 2:
            Name                          = ServicesApp_startafter_haip
            Class:Resource:Node[Source] = IBM.Application:ServicesApp
            ResourceGroup[Source]       = harg
    Managed Relationship 3:
            Name                          = EngineApp_startafter_haip
            Class:Resource:Node[Source] = IBM.Application:EngineApp
            ResourceGroup[Source]       = harg
    ```
12. Run the **rgreq** command again to unlock the resource group:
    ```
    rgreq -o unlock harg
    ```

Testing the configuration (Tivoli System Automation for Multiplatforms example):

To test the configuration, simulate active server failures and network failures and observe the results.

About this task

During the test, run the **lssam** command periodically to view the status of the servers in the cluster. When the active server in the cluster is operational and the passive server is on standby, the **lssam** command displays output like this example:

```
Online IBM.ResourceGroup:harg Nominal=Online
        |- Online IBM.Application:EngineApp
                |- Online IBM.Application:EngineApp:myprimary
                '- Offline IBM.Application:EngineApp:mysecondary
        |- Online IBM.Application:ServicesApp
                |- Online IBM.Application:ServicesApp:myprimary
                '- Offline IBM.Application:ServicesApp:mysecondary
        '- Online IBM.ServiceIP:haip
                |- Online IBM.ServiceIP:haip:myprimary
                '- Offline IBM.ServiceIP:haip:mysecondary
```

The first line indicates the resource group. In the preceding example, the resource group name is `harg`.

The following lines indicate the status of the member resources: EngineApp (the IBM InfoSphere Information Server engine tier control application), ServicesApp (the InfoSphere Information Server services tier and metadata repository tier control application), and `haip` (the virtual IP resource). Resources and constituent servers can be in any of the following states:

- Online: The resource was started and is functioning normally:
- Offline: The resource was stopped successfully.
- Failed Offline: The resource malfunctioned.

Procedure

1. Run the **lssam** command to verify that the active server in the cluster is operational and functioning normally.
2. Simulate a power failure on the active server. For example, turn off the power or unplug the power cable. Verify that all resources fail over to the other server and remain online there.
3. Run the **lssam** command to verify that the failover was successful. The command output is similar to the following text:

   ```
   Online IBM.ResourceGroup:harg Nominal=Online
           |- Online IBM.Application:EngineApp
              |- Failed offline IBM.Application:EngineApp:myprimary Node=Offline
              '- Online IBM.Application:EngineApp:mysecondary
           |- Online IBM.Application:ServicesApp
              |- Failed offline IBM.Application:ServicesApp:myprimary Node=Offline
              '- Online IBM.Application:ServicesApp:mysecondary
           '- Online IBM.ServiceIP:haip
              |- Failed offline IBM.ServiceIP:haip:myprimary Node=Offline
              '- Online IBM.ServiceIP:haip:mysecondary
   ```

   These lines indicate that the application resource failed over successfully to the passive server (`mysecondary`). Also, the virtual IP resource failed over successfully. Verify that the formerly active server (`myprimary`) is now in the failed offline state.
4. Use an InfoSphere Information Server client to access the system and verify that it is functioning normally.
5. Restore power to the formerly active server (`myprimary`).
6. Run the **lssam** command. The command output is similar to the following text:

```
Online IBM.ResourceGroup:harg Nominal=Online
        |- Online IBM.Application:EngineApp
             |- Offline IBM.Application:EngineApp:myprimary
             '- Online IBM.Application:EngineApp:mysecondary
        |- Online IBM.Application:ServicesApp
             |- Offline IBM.Application:ServicesApp:myprimary
             '- Online IBM.Application:ServicesApp:mysecondary
        '- Online IBM.ServiceIP:haip
             |- Offline IBM.ServiceIP:haip:myprimary
             '- Online IBM.ServiceIP:haip:mysecondary
```

Verify that the formerly active server (myprimary) is now in the offline state. You might need to run the command several times before the status is displayed as shown in the preceding text.

7. Simulate a power failure on the formerly passive (mysecondary) server. Verify that all resources fail over to the myprimary server and remain online there.
8. Run the **lssam** command to verify that the failover was successful. The command output is similar to the following text:

```
Online IBM.ResourceGroup:harg Nominal=Online
        |- Online IBM.Application:EngineApp
           |- Online IBM.Application:EngineApp:myprimary
           '- Failed offline IBM.Application:EngineApp:mysecondary Node=Offline
        |- Online IBM.Application:ServicesApp
           |- Online IBM.Application:EngineApp:myprimary
           '- Failed offline IBM.Application:EngineApp:mysecondary Node=Offline
        '- Online IBM.ServiceIP:haip
           |- Online IBM.ServiceIP:haip:myprimary
           '- Failed offline IBM.ServiceIP:haip:mysecondary Node=Offline
```

These lines indicate that the application resource failed over successfully to the myprimary server. Also, the virtual IP resource failed over successfully. Verify that the mysecondary server is now in the failed offline state. You might need to run the command several times before the status is displayed as shown in the preceding text.

9. Use an InfoSphere Information Server client to access the system and verify that it is functioning normally.
10. Restore power to the mysecondary server.
11. Run the **lssam** command. The command output is similar to the following text:

```
Online IBM.ResourceGroup:harg Nominal=Online
        |- Online IBM.Application:EngineApp
           |- Online IBM.Application:EngineApp:myprimary
           '- Offline IBM.Application:EngineApp:mysecondary
        |- Online IBM.Application:ServicesApp
           |- Online IBM.Application:EngineApp:myprimary
           '- Offline IBM.Application:EngineApp:mysecondary
        '- Online IBM.ServiceIP:haip
           |- Online IBM.ServiceIP:haip:myprimary
           '- Offline IBM.ServiceIP:haip:mysecondary
```

Verify that the mysecondary server is now in the offline state. You might need to run the command several times before the status is displayed as shown in the preceding text.

12. Repeat steps 1 through 11. During this test cycle, restart each server in turn instead of disabling power. For example, to restart a server, you can log in to the server as root and run the **reboot** command.
13. Repeat steps 1 through 11. During this test cycle, simulate a network failure by disconnecting the public network cable from each server in turn.

Creating an active-passive configuration (Microsoft Windows)

Regardless of the Microsoft Windows operating system version that you use, you must complete the same basic steps to configure a two-server active-passive configuration.

About this task

This procedure describes an active-passive configuration where both the services and engine tiers have a standby server available. Refer to Tech Note 1591193 (http://www.ibm.com/support/docview.wss?uid=swg21591193), for a procedure on an active-passive engine configuration. The tech note contains detailed instructions and screen captures.

Procedure

1. Set up the servers and storage, and cable the network.

   The servers must be able to communicate with each other over high-speed networks. To minimize single points of failure, each server must have redundant connections to at least two physically separate high-speed networks.

   To support failover when a problem occurs, the high availability software creates a heartbeat signal between the servers. Set up at least two more physically separate networks for this signal. The networks that carry the heartbeat signal must be physically separate and isolated from the high-speed networks.

   Each server must connect to a shared storage device. The storage device can be connected to each server by using a Fibre Channel Switch or SCSI interface. The storage device is listed as a local disk on both servers. At any given time, only one node has exclusive access to the storage device.

   Configure a separate server as the Windows domain controller.
2. Install the operating system on both servers. The operating system includes high availability functions.
3. Set up the cluster. Verify that all cluster resources are running correctly, and that all resources failover properly from server to server.
4. Set up the virtual IP address.
5. Create users and user groups on the shared storage device.
6. On both servers, make sure that the directories that are listed in the following table are shared. The directories might differ depending on your configuration.

Table 70. Directories for active-passive configuration

Configuration of tiers	Directory	Default directory
All	Project directory	`C:\IBM\InformationServer\Server\Projects`
Engine tier on both servers	IBM InfoSphere Information Server installation directory	`C:\IBM\InformationServer`
Services tier on both servers	IBM WebSphere Application Server installation directory	`C:\IBM\WebSphere`
Metadata repository tier on both servers, IBM DB2 database system	• DB2 database system installation directory • DB2 administrator home directory	`C:\IBM\db2`

7. If the active-passive configuration will include only the engine tier (that is, the metadata repository or services tiers are to be located outside of the active-passive configuration), install and configure the metadata repository and services tiers first, before you continue with the engine tier installation and configuration. The engine tier installation requires that the metadata repository tier and services tier are installed first, unless you install them all at the same time in a single pass of the installation program.
8. Run the InfoSphere Information Server installation program on the primary server.

   When prompted by the installation program, specify the virtual host name. Whenever prompted for an installation directory, specify a path within the shared storage device.
9. Cause a failover to the backup server.
10. Delete the InfoSphere Information Server installation directory on the shared storage device. Repeat the installation on the backup server with the same user names, passwords, and directories that are used on the primary server.
11. Use the appropriate high availability software administration tools to complete the following steps:
    a. If the engine tier is included in the high availability configuration, add Generic Service resources for each of the engine tier services.
    b. If the services tier or metadata repository tier is included in the high availability configuration, add Generic Service resources for WebSphere and DB2 services.
    c. Create a Generic Script resource and add the supplied script in *installdir*\InformationServer\HAScripts\DSHARestart.vbs.
    d. For each of the resources you just added, set the appropriate dependency in the cluster software.
12. Test the setup by simulating different failures.

Example: Setting up an active-passive configuration by using the Microsoft failover clustering service:

This example describes how to set up a two-server active-passive high availability configuration by using the Microsoft failover clustering service in Microsoft Windows Server 2008.

This example configuration is built on two Microsoft Windows Server 2008 servers. The servers are connected to each other over a single public network. Each server is also connected to a storage area network (SAN). All IBM InfoSphere Information Server components are installed on the SAN. At any given time, only one of the servers is active; that is, it hosts the InfoSphere Information Server components on the SAN. The other server is passive and does not access the SAN.

The following diagram illustrates the example configuration:

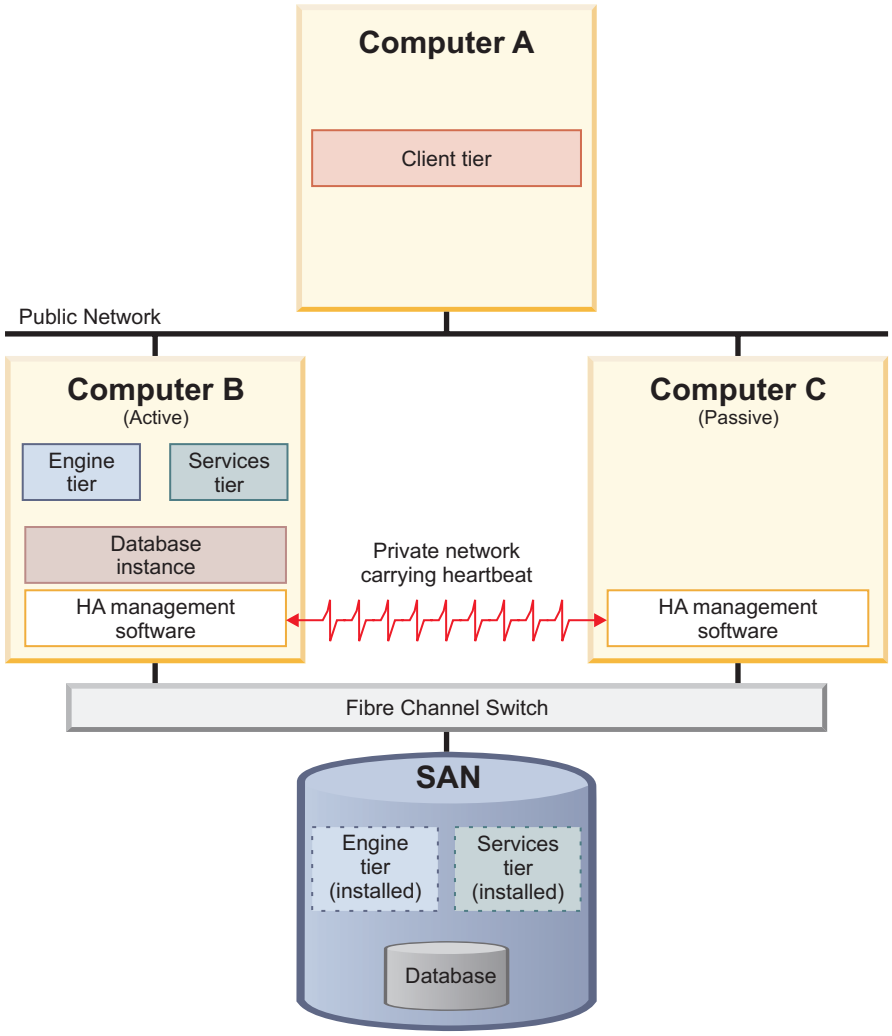


Figure 41. Example of a two-server active-passive high availability configuration

In this example, the following software is used for each server:

- Microsoft Windows Server 2008 64-bit
- IBM InfoSphere Information Server
- IBM WebSphere Application Server Network Deployment (installed with the InfoSphere Information Server installation program)
- IBM DB2 Database for Linux, UNIX, and Windows (installed with the InfoSphere Information Server installation program)

For storage, there is a shared storage area network (SAN) that is accessible to both servers. The SAN is connected to the servers through a Fibre Channel switch.

To create the configuration in this example, follow these steps:

1. Set up the hardware. Then configure the network, and configure the connection to the storage area network (SAN). See "Setting up the servers, network, and storage (Microsoft example)" on page 186.
2. Use the Failover Cluster Manager to set up the cluster. See "Setting up the cluster (Microsoft example)" on page 186.
3. Create resources for the virtual IP address and virtual host name. Then, map the virtual host name to the virtual IP address in the `etc\hosts` file on each

server. Finally, add environment variables that specify the virtual host name. See “Setting up the virtual IP address (Microsoft example)” on page 187.

4. Create users and user groups on the SAN. Then, run the InfoSphere Information Server installation program on the primary server. Finally, cause a failover to the backup server, and run the installation program on the backup server. See “Installing IBM InfoSphere Information Server in the cluster (Microsoft example)” on page 187.
5. Set up the *installdir*\InformationServer\HAScripts\DSHARestart.vbs script to start and stop the servers. Then, create the application resources.
6. Use the Failover Cluster Manager to test the configuration, simulate active server failures and network failures, and observe the results. See “Testing the configuration (Microsoft example)” on page 189.

Setting up the servers, network, and storage (Microsoft example):

Set up the hardware. Then, configure the network and configure the connection to the storage area network (SAN).

About this task

For detailed information about each step in the procedure, see the Microsoft failover clustering documentation.

Procedure

1. Set up the servers, network, and storage hardware.
2. Install the Windows Server 2008 operating system on both servers.
3. Within Windows on each server, rename the local area network icons so that they reflect the purpose of each network. For example, name the public network icon “Public”. Name the private network icon “Private”.
4. Configure the binding order for the networks on each server. Set up the following order: Public, then Private, then Remote Access Connections.
5. Configure the private network adapter on each server. Set the protocol to TCP/IP, and specify a static IP address. Specify a network speed, and duplex mode.
6. Configure the public network adapter on each server. Specify a static IP address. Verify connectivity by pinging all IP addresses from each server.
7. Add the servers to the same domain in the domain controller. Verify that each server can access the domain controller.
8. Set up a domain user account that is dedicated to the Cluster service.
9. Set up a quorum disk. The quorum disk is a separate disk that stores checkpoint data and other information that is required for cluster operation.
10. Verify that the SAN can be accessed from both servers, and that it is functioning properly.

Setting up the cluster (Microsoft example):

Use the Cluster Configuration wizard to set up the cluster.

About this task

For detailed information about each step in the procedure, see the Microsoft failover clustering documentation.

Procedure

1. Turn off the backup server. Turn on the storage area network (SAN).
2. Log in to the primary server. Use a domain account that has local administrator privileges.
3. Using the Failover Cluster Manager, complete the following configuration tasks to configure the primary server:
   a. Specify a unique NetBIOS cluster name.
   b. Make the server the first node in the cluster.
   c. Analyze the node to make sure that there are no hardware or software problems that might interfere with the cluster configuration.
   d. Associate the IP address of the cluster with the public network.
   e. Specify the name of the domain account that you created that is dedicated to the cluster service.
   f. Create the cluster.
4. Use the Failover Cluster Manager to verify that all cluster resources are running correctly.
5. Leave the primary server and SAN on. Turn on the backup server.
6. Use the Failover Cluster Manager to add the backup server to the cluster.
7. Use the Failover Cluster Manager to configure the heartbeat and prioritize the network adapters for the private network.
8. Use the Failover Cluster Manager to configure the quorum disk.
9. Set up a boot delay so that simultaneous restoration of power to all nodes does not cause the cluster service to fail to start.
10. Test that all resources failover properly from server to server.

Setting up the virtual IP address (Microsoft example):

Create resources for the virtual IP address and virtual host name. Then, map the virtual host name to the virtual IP address in the `etc\hosts` file on each server. Finally, add environment variables that specify the virtual host name.

About this task

For detailed information about each step in the procedure, see the Microsoft failover clustering documentation.

Procedure

1. Use the Failover Cluster Manager to create a resource for the virtual IP address. Use the IP address resource type.
2. On each server, open the `C:\Windows\System32\drivers\etc\hosts` file. In each file, map the virtual host name to the virtual IP address.

Installing IBM InfoSphere Information Server in the cluster (Microsoft example):

Create users and user groups on the storage area network (SAN). Then, run the InfoSphere Information Server installation program on the primary server. Finally, cause a failover to the backup server, and run the installation program on the backup server.

Procedure

1. On the SAN, create all users and user groups. For details about creating administrator users and groups, see Setting up administrator and owner accounts.
2. Run the InfoSphere Information Server installation program on the primary server in graphical mode. See Installing IBM InfoSphere Information Server by using the wizard.

   Whenever you are prompted for a directory, specify a directory that is located on the SAN.

   In the High Availability Server Cluster Configuration page, click **This is an active-passive server cluster configuration**. In the **Virtual host name** field, enter the virtual host name.

   Important: If your active-passive configuration will include only the engine tier, do not select the high availability option when installing the engine tier on the backup server.
3. After the installation program completes the installation process, cause a failover to the backup server. See the Microsoft failover clustering documentation for more information.
4. Delete the InfoSphere Information Server installation folder on the SAN.
5. Run the InfoSphere Information Server installation program on the backup server in graphical mode. Be sure to specify the same user names, passwords, and path names as you did when you ran the installation program on the primary server.

Setting up the engine and services tier resources for high availability (Microsoft example):

You set up the engine and services tier Generic Service resources and a DSHARestart Generic Script resource.

About this task

For information on how to create Generic service resources and set dependencies as indicated in the following procedure, refer to Tech Note 1591193.

Procedure

In the Failover Cluster Manager:

1. Select to add a Generic Service resource for each of the engine tier services: ASBAgent, DSTelnet, DSEngine and DSRPC.
2. Select to add a Generic Service resource for DB2 service, DB2 - DB2COPY - DB2-0 and Websphere service, IBMWAS70Service - hostnameNode01
3. Select Generic Script resource and add the supplied *installdir*`\InformationServer\HAScripts\DSHARestart.vbs` script.
4. For each of the resources you just created, set up dependencies as indicated in the following table.

Resource	Depends on
ASBAgent	Clustered Disk, IP Address, and WebSphere service
DSTelnet	Clustered Disk, IP Address, DSEngine

Resource	Depends on
DSRPC	Clustered Disk, IP Address, DSHARestart Generic Script
DSEngine	Clustered Disk, IP Address, ASBAgent
DSHARestart	Clustered Disk, IP Address, DSEngine
DB2 service	Clustered Disk and IP address
WebSphere service	Clustered Disk, IP address and DB2 generic service

Testing the configuration (Microsoft example):

Use the Failover Cluster Manager to test the configuration, simulate active server failures and network failures, and observe the results.

About this task

For detailed information about each step in the procedure, see the Microsoft failover clustering documentation.

Procedure

1. Using the Failover Cluster Manager, verify that all cluster resources are running correctly.
2. Simulate a power failure on the active server. For example, turn off the power or unplug the power cable.
3. Verify that all resources fail over to the other server and remain online there.
4. Use an IBM InfoSphere Information Server client to access the system and verify that it is functioning normally.
5. Restore power to the formerly active server.
6. Verify that the formerly active server is now in the offline state.
7. Simulate a power failure on the formerly passive server.
8. Verify that all resources fail over to the active server and remain online there. Verify that the passive server is now in the failed offline state.
9. Use an InfoSphere Information Server client to access the system and verify that it is functioning normally.
10. Restore power to the passive server.
11. Verify that the backup server is now in the offline state
12. Repeat the previous steps. During this test cycle, restart each server in turn instead of disabling power. For example, to restart a server, you can log in to the server and run the **reboot** command.
13. Repeat the previous steps. During this test cycle, simulate a network failure by disconnecting the public network cable from each server in turn.

Part 2. Installing IBM InfoSphere Information Server software

You can create an installation of IBM InfoSphere Information Server or update an existing installation with a patch or fix pack.

© Copyright IBM Corp. 2005, 2014

Chapter 3. Preparing to install software

Before you install InfoSphere Information Server, ensure that you have the correct installation media and that your computers are ready for installation.

Selecting the installation order in a multi-computer topology

If you are installing IBM InfoSphere Information Server on more than one computer, install the tiers in the following order. If you are installing multiple tiers on the same computer, install the tiers in the same installation session. For example, you might choose to install both the repository and services tiers on the same computer. In the installation program, you would select to install the two tiers in the same installation session.

1. **Repository tier:** Each of the following tiers requires connection information for the metadata repository so that their components can be registered.
2. **Services tier:** The services tier registers with the metadata repository.
3. **Engine tier:** Each engine is registered with both the metadata repository and services tier.
4. **Client tier:** Although you can install the client tier at any time, if you are installing IBM InfoSphere Metadata Asset Manager, the metadata interchange server is automatically registered with the metadata repository if you install the client tier after installing the metadata repository.

   The client tier has no dependencies and can be installed by itself on a computer. However, if you decide to install any of the other tiers on the same computer where a client has already been installed, the existing client must first be uninstalled and installed again along with the additional tiers.

Obtaining IBM InfoSphere Information Server software

Acquire the InfoSphere Information Server software and applicable fix packs. Copy the files to each tier computer.

Before you begin

To ensure uninterrupted connectivity to the installation media, copy the contents of the installation media to a local file system.

Procedure

1. Acquire the installation media.

Option	Description
If you have the installation media in physical form	Check that you have all of the installation disks.
If you do not have the installation media in physical form	Download the installation image files from Passport Advantage. Use the Knowledge Collection support document to determine the image files that are appropriate for your operating systems and configurations.

2. Download any applicable fix packs from Fix Central.

© Copyright IBM Corp. 2005, 2014

3. If you downloaded installation image files, uncompress the installation image files onto each tier computer.

   Windows

   To avoid corrupting the files in the image, do not use the **Extract All** feature of the MicrosoftWindows uncompression utility to uncompress the files. Use another standard uncompression utility such as WinZip. Also, extract the files into a short directory path, such as `C:\temp` rather than into a long directory path.

Configuring your browser

To run the installation program successfully, you must configure your browser to support the requirements of the installation program.

All browsers

- Verify that you have a supported browser. For details, go to the system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.
- Make sure that JavaScript is enabled.

Mozilla Firefox

Newer versions of Firefox enable JavaScript and load images automatically by default. If you have older versions of Mozilla Firefox and have disabled these options, enable them before running the installation program:

1. In the **Options**, menu select **Options**. The Options window displays.
2. Select **Content**.
3. Enable **Load images automatically**.
4. Enable **Enable JavaScript**.
5. Click **OK**.

Microsoft Internet Explorer

Restriction: Graphics in the installation wizard might not display correctly in Microsoft Internet Explorer. To see the graphics in the wizard, use Mozilla Firefox if you have it installed.

Make sure that the security zone for the computer allows the installation program to run:

1. From Microsoft Internet Explorer, in the **Tools** menu, select **Internet Options**.
2. Click the **Security** tab.
3. Click the **Trusted Sites** icon.
4. Click **Sites**.
5. In the **Add this Website to the zone** field, type `https://`*hostname* where *hostname* is the host name of the computer on which the installation program will run.
6. Click **Add**.
7. Click **OK**.
8. Click **OK**.

Preparing to install InfoSphere Metadata Asset Manager and bridges to import and export metadata

Components of InfoSphere Metadata Asset Manager are installed automatically with the suite on the services and repository tiers. To use bridges and connectors you must install additional software. When you install the metadata interchange agent with bridges, you can use bridges to import and export metadata with InfoSphere Metadata Asset Manager.

About InfoSphere Metadata Asset Manager

InfoSphere Metadata Asset Manager provides robust functionality for importing metadata. You use bridges or connectors in InfoSphere Metadata Asset Manager to import metadata assets from applications, files, and databases and to export database schemas.

InfoSphere Metadata Asset Manager has replaced Import Export Manager as the tool for importing and exporting metadata.

You import the metadata to a staging area where you can analyze the source metadata and preview the results before you share it to the metadata repository. You can browse and search the metadata repository and manage common metadata assets. The staging area is a schema within the metadata repository database that stores staged imports for InfoSphere Metadata Asset Manager.

For more information about InfoSphere Metadata Asset Manager, see the *IBM InfoSphere Information Server Guide to Managing Common Metadata*.

Preparing to use bridges and connectors

Table 71. Installing bridges and connectors to use with InfoSphere Metadata Asset Manager

Bridge or connector	Tier	Requirements
Bridges: CA ERwin 7 Data Modeler (File) CA ERwin 8 Data Modeler (File) CA ERwin 8Data Modeler (File) Embarcadero ER/Studio Business Architect (via DM1) Embarcadero ER/Studio Data Architect IBM Cognos BI Reporting - Content Manager IBM DB2 Cube Views IBM DB2 Warehouse Manager (via CWM XMI) IBM InfoSphere Data Architect MetaBroker IBM InfoSphere Discovery MetaBroker Microsoft SQL Server Analysis and Reporting Services (Repository) MicroStrategy OMG CWM 1 XMI 1 Oracle Business Intelligence Enterprise Edition Oracle Hyperion Application Builder (via CWM XMI) Oracle Warehouse Builder (via CWM XMI)SAP BusinessObjects Repository SAS Data Integration Studio (via CWM XMI) Sybase PowerDesigner CDM Sybase PowerDesigner PDM	Client	Install **Metadata interchange agent and bridges**. All of the listed bridges are installed. The metadata interchange agent enables interaction between these bridges and InfoSphere Metadata Asset Manager. The computer on which you install the metadata interchange agent and bridges must meet the prerequisites and system requirements for installing bridges. For a list of installation prerequisites, see the technote *System requirements for InfoSphere Metadata Integration Bridges*: http://www.ibm.com/support/docview.wss?&uid=swg27037866.

Table 71. Installing bridges and connectors to use with InfoSphere Metadata Asset Manager (continued)

Bridge or connector	Tier	Requirements
IBM InfoSphere Streams bridge IBM InfoSphere Master Data Management bridge	Services	These bridges are installed automatically when you install the services tier.
HDFS bridge	Engine tier	The HDFS bridge is installed automatically when you install the engine tier on a Linux or AIX computer.
Connectors: Amazon S3 Greenplum IBMInfoSphereDB2 IBMNetezza® JDBC ODBC Oracle Teradata	Engine tier	Install one of the following products on the engine tier: • InfoSphere DataStage and QualityStage • InfoSphere Information Analyzer All of the listed connectors are installed.

When you install the metadata interchange agent with bridges, the following components are automatically installed on the client tier computer:

- Metadata interchange agent
- InfoSphere Metadata Integration Bridges
- Websphere Application Server Liberty Profile. If the client tier and the services tier are on the same computer and Liberty Profile is already installed, an additional profile is installed for the metadata interchange agent.

Designating metadata interchange servers

During installation, a *metadata interchange server* is usually automatically designated for each client tier computer where the bridges are installed and each engine tier computer where the connectors are installed.

When you import or export by using InfoSphere Metadata Asset Manager, you select a metadata interchange server on which bridges, connectors, or both are installed. Then you specify the bridge or connector that you want to use for the import or export.

In a distributed computer topology, when you install the bridges on a computer, the computer's metadata interchange server is automatically designated unless you are installing a client computer and IBM InfoSphere Information Server has not yet been installed on the services tier computer. If you install bridges on the client computer before you install InfoSphere Information Server on the services tier computer, you can manually designate the client computer's metadata interchange server on the Administration tab of InfoSphere Metadata Asset Manager or on the command line after the installation of InfoSphere Information Server. When you manually designate a metadata interchange server, you must specify the HTTPS port for the computer. The default HTTPS port is 19443.

The IBM InfoSphere Streams bridge and the IBM InfoSphere Master Data Management bridge. which are installed on the services tier automatically, are available to use with any metadata interchange server.

Preparing to run the installation program

Before you run the installation program to create an installation or add product modules or tiers, make sure that the computers are ready for installation. Also back up critical data.

Before you begin

AIX Linux

DB2 requires that the 32-bit Pluggable Authentication Modules (PAM) be installed on the metadata repository tier computer for access to the libpam library.

Windows

Verify that you have a valid installation of the Windows Script Engine. See the Microsoft documentation for more information. Additionally, ensure that the following directories are in your system path, and exist on the target computer. Remove any paths to Java executables from the system PATH.

- `%SystemRoot%\`
- `%SystemRoot%\system32`
- `%SystemRoot%\system32\wbem`

Procedure

Follow this procedure for each target computer:

1. Log in to the target computer.

Operating system	Description
AIX	Log in as the root user.
Linux	Log in as the root user.
Windows	Log in as a local administrator or as a local or domain user who is directly assigned to the local Administrators group.

2. Back up your systems.
   a. Back up all critical data.
   b. Back up existing instances of IBM InfoSphere Information Server if your computer already hosts an instance.
   c. Back up the following directories.

Operating system	Directories
AIX	`/etc/services`, `/etc/inittab`, `/etc/passwd`, and `/etc/group`
Linux	`/etc/services`, `/etc/inittab`, `/etc/passwd`, and `/etc/group`
Windows	Microsoft Windows registry and the `C:\Windows\System32\drivers\etc`

3. If you removed a previous installation of InfoSphere Information Server on the computer, reboot the computer before continuing.
4. The installation program writes required temporary files to a location that is defined by the following environment variable.

Operating system	Environment variable
AIX	*TMP*
Linux	*TMP*
Windows	*%temp%*

There must be roughly 2.5 GB of temporary space available.

You can redirect most of the temporary files to a different directory if desired. If you do this, about 1.5 GB will be redirected to the specified directory, but 1 GB will still be required in the directory specified above.If you want to redirect 1.5 GB of temporary files to a different temporary directory, create a new file called `.jvm_args` in the root installation directory and add the following line with the new temporary directory:

`-Djava.io.tmpdir=`*`installation_program_tempdir`*

After the installation completes, do not delete the installation log files that are in the temporary directory. Preserve these files for troubleshooting and verification purposes.

5. **Linux**: Because WebSphere Application Server is not compatible with Security Enhanced Linux (SELinux), disable SELinux or set it to permissive mode before installing InfoSphere Information Server. For more information about how to disable SELinux, see the Linux documentation.
6. **Windows**: For a client tier or engine tier installation, if you installed the MKS Toolkit software for UNIX and Linux compatability, uninstall it and reboot the computer before you run the InfoSphere Information Server installation program.
7. **Windows**: Ensure that the length of your system PATH environment variable will not exceed 2047 characters after installation.

   Note: The installation of IBM InfoSphere Information Server will add approximately 600 characters to the PATH system environment variable if you accept the default directory paths. Therefore, to avoid potential problems with installation, ensure that the length of your PATH prior to installation does not exceed 1447 characters. For more information, refer to the Microsoft documentation, http://support.microsoft.com/kb/2685893.
8. **Windows**: To import metadata by using bridges and connectors with InfoSphere Metadata Asset Manager, take the steps in the following table. For more information, see "Preparing to install InfoSphere Metadata Asset Manager and bridges to import and export metadata" on page 195.

Option	Description
To use bridges with InfoSphere Metadata Asset Manager and Import Export Manager	1. Remove any existing installations of bridges from the client tier Microsoft Windows computer before you run the InfoSphere Information Server installation program. For detailed instructions, see the system requirements for bridges: http://www.ibm.com/support/docview.wss?&uid=swg27037866. 2. When you run the installation program, for the client tier select **Metadata interchange agent and bridges**.
To use connectors with InfoSphere Metadata Asset Manager	When you run the installation program, install InfoSphere DataStage and QualityStage on the engine tier.

9. If you are installing an instance of InfoSphere Information Server on a computer with an existing instance, disable all scheduled jobs, scheduled tasks, scheduled reports, and enabled IBM InfoSphere Information Services Director applications.
10. Disable firewall software and antivirus software before starting the installation.
11. For an installation that involves WebSphere Application Server clustering:
    - All management processes including the node agents must be started before the installation. Also, sufficient file descriptor resources must be allocated to the processes by using the **ulimit** command. See "Running the installation program (clustered installation)" on page 160.
    - Verify that you can start and stop the cluster, and restart node agents. See the IBM InfoSphere Information Server Administration Guide.
    - Verify that you can stop and start the web server.
    - For optimal results, install into a cluster that has one Deployment Manager and one WebSphere Application Server node. Add additional nodes as business needs dictate post-installation.
12. If the WebSphere Application Server processes or node agents are currently running under a non-root user, there is no longer a need to change them to run as root before running the installation or uninstallation program. However, after running the program, the application server will have been switched to run as root. Therefore, if you want to run it as non-root again, you will need to switch it to run as the non-root user afterward.

Chapter 4. Installing IBM InfoSphere Information Server software

Use the InfoSphere Information Server installation program to install the product. You can also use the installation program to add product modules or software tiers to an existing installation.

The installation program can also install IBM DB2 and IBM WebSphere Application Server Network Deployment if your installation requires stand-alone configurations of these components. Alternatively, the installation program can set up InfoSphere Information Server to use preinstalled instances of these components, or other supported database systems and application servers.

The installation program can run in either of two modes:

Graphical mode
: In this mode, the installation program runs within a standard Web browser, either locally on the target computer or remotely. The installation program presents wizard pages for you to select options. When you have selected your options, the program saves them to a response file on the target computer. Then the program performs a prerequisites check. If the check succeeds, the program installs the software.

Silent mode
: In this mode, the installation program does not present a user interface. Instead, it reads settings from a response file. It performs a prerequisites check. If the check succeeds, the program installs the software. As it performs the installation, it can save the report to a text file or display results on the console.

Important: During the initial installation session, install all of the product modules that you plan to use either now or in the future. If you install only some of the product modules now, and then install fixes and service packs, you might need to reinstall the fixes and service packs if you add product modules later.

Installing IBM InfoSphere Information Server by using the wizard (graphical mode)

You can use a web browser to run the installation program locally in graphical mode on the target computer. Alternatively, you can use a web browser on any system that has network access to the computer where the installation program is running.

Before you begin

Choose your installation topology, review the installation checklists, and prepare the tier computers. Obtain the product media to include the installation and entitlement packages. Configure your browser correctly so that the installation program runs smoothly.

Microsoft Windows: Start an administrator session by using one of the following methods:

- Run the following command:

© Copyright IBM Corp. 2005, 2014

```
runas /user:Administrator cmd.exe
```

- Open an elevated command window.
  1. Click **Start** > **Search**.
  2. In the Search Results box, enter `cmd` and press Ctrl+Shift+Enter.
  3. Click **Continue** to accept the prompt asking for your permission to continue.

AIX Linux: Log in as the root user.

About this task

When you first start the installation program, it presents a non-graphical interface. The program starts a self-contained web server on the computer. You can then start a web browser on the computer or on a remote computer and access the graphical interface.

After you enter your settings in the wizard pages, the program stores your settings in a response file. If you run the installation program again, you can load your settings from this file into the program instead of entering them again.

Procedure

1. In the newly created session, run the following command to clear any Java options that might be set.

Operating system	Command
AIX	In the sh or bash shell: `export IBM_JAVA_OPTIONS=` In the csh or tcsh shell: `unsetenv IBM_JAVA_OPTIONS`
Linux	In the sh or bash shell: `export IBM_JAVA_OPTIONS=` In the csh or tcsh shell: `unsetenv IBM_JAVA_OPTIONS`
Windows	`set IBM_JAVA_OPTIONS=`

2. **AIX Linux**: Configure the following settings for the session to help ensure that the installation is successful:
   - Set the file creation mask to 022:
     ```
     umask 022
     ```
   - Ensure that you have no less than 10240 file descriptor resources allocated.
   - Check that the current setting for the NOFILES kernel parameter is set to a value that is at least equal to the value that you specify for **ulimit**.
3. Extract the installation package. For example:

Operating system	Command
AIX	`gunzip -c is-suite.*.tar.gz | tar xvf -`
Linux	`tar zxvf is-suite.*.tar.gz`

Operating system	Command
Windows	The installation image contains some deep directory paths and long file names. Therefore, to avoid running into the Windows 255 character path name limit during later installation steps, extract the contents into a short initial directory path, such as `C:\temp`. Also, instead of using the Windows File Explorer extraction utility, use a utility such as PKZip, WinZip, or 7-zip. For example: **`7z.exe x is-suite.*.zip -oC:\temp`**

The package is named `is-suite.*` for the full product or `is-client.*` for the client only.

4. The is-suite images require the use of a Bundle Entitlement package; the is-client images do not. If you are installing an is-suite image, extract the contents of the entitlement package into the `is-suite` directory that is decompressed in the previous step. Ensure that you have downloaded the correct entitlement package for your operating system platform, licensed products, and suite installation package. A suite entitlement package will only work for a suite installation package. For example:

   ```
   unzip -d is-suite Bundle.suite.*.zip
   ```

   On AIX, if you don't already have the GNU zip utility, it is available as `gzip-`*`program_version.aix_version`*`.ppc.rpm` in the IBM AIX Toolbox.
5. Change to the `is-suite` or `is-client` directory.
6. Run the installation program:

Operating system	Command
AIX	**`./setup [-verbose]`**
Linux	**`./setup [-verbose]`**
Windows	**`setup.exe [-verbose]`**

The `-verbose` option is optional. The option causes the installation program to display log messages to the console and also write them to the log file. If the `-verbose` option is not specified, log messages are written to the log file only.

Several messages are displayed. The program then displays a message about how to access the installation wizard:

```
======> Enter the following URL to your web browser to begin the
installation process:
https://localhost:8443/ISInstall
```

7. Open a web browser on the computer or on a remote computer, and navigate to the URL that is listed in the message from starting the installation program. (Use the full host name if using a remote browser.)



Figure 42. Installation program opens in a web browser

Microsoft Windows: The installation program automatically launches the default web browser on the local computer to display the graphical interface.

A message might appear in the browser that reads "The publisher cannot be verified..." You can safely ignore the message and continue the installation. This message does not indicate a security risk when it appears during InfoSphere Information Server installation.

a. Select a language from the menu for the installation program to use, and click **Login**.

   This selection applies only to this installation session. Later in the installation process, you select the language to apply to the InfoSphere Information Server installation.

b. Follow the prompts in the installation program to continue installing the product.

The installation program runs an early prerequisites check to ensure that your installation is configured correctly, as shown in the following figure. The wizard page displays each check. If the check fails, it is marked FAILED in the page. If the check found issues that might cause problems with the installation,

it is marked WARNING in the page.

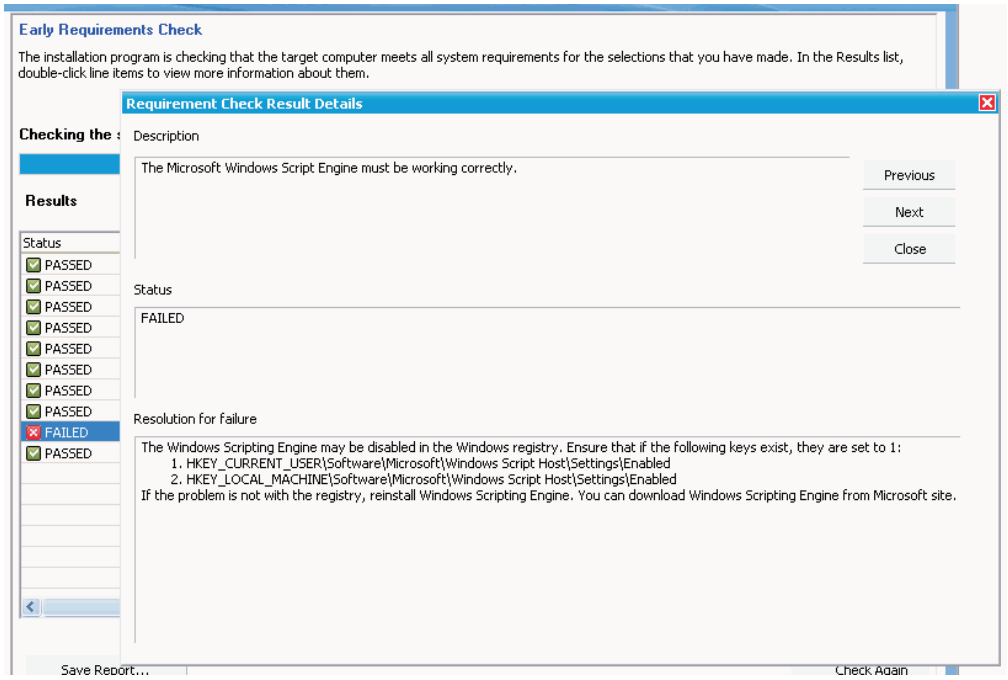


Figure 43. Early requirements check showing a failed message for the Windows Script Engine

8. When all prerequisites checks pass, continue to follow the prompts in the installation program.

Taking action if the installation prerequisites check fails

The installation program runs an early prerequisites check before installing the product to ensure that the target computer satisfies the remainder of the installation sequence.

Symptoms

During the early prerequisites check, you received FAILED or WARNING messages

User response

Complete these steps by following the prompts in the wizard:

1. Double-click each message to learn more about resolving it.
2. Try to resolve each issue without exiting the installation program. Then, click **Check Again** in the Prerequisites Check page.
3. If you must exit the installation program, click **Cancel**, close the web browser window. Resolve the issues and then restart the installation program.

The following figure shows a portion of the Early Requirements Check page. When you select a FAILED message and click **Details**, resolution information is displayed.

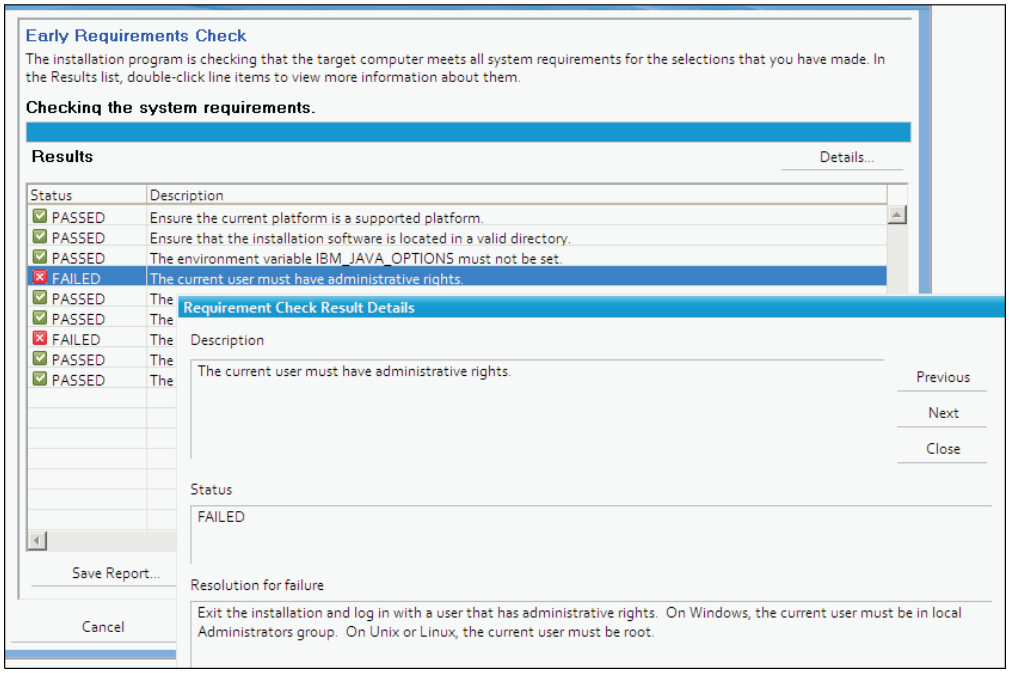


Figure 44. Requirement Check Result Details window

Specifying tiers and product modules to install

After you start the installation program, specify the software tiers and product modules that you want to install for each of the tier computers where you are installing InfoSphere Information Server. The following topics detail the options available for each tier. However, when you run the installation program on particular computer, select every tier that you want to install on that computer and install them at the same time. Reference the following topics for information on some of the options available for each tier.

Specifying installation options (repository tier)

Use the InfoSphere Information Server installation program to specify options for the computer that will host the metadata repository, such as the property information for your database system.

Procedure

1. Select the directories where you want to install the product that you determined when planning your installation.
2. Select **New installation**.
3. Select the tiers to install on the target computer.
4. Enter the InfoSphere Information Server administrator user name and password. Use the information that you determined when planning your installation.
5. Select the product modules that you want to install. The product modules that are displayed are only those that are part of the entitlement package that you purchased. Select the same set of product modules for each computer.
6. If you are configuring your topology for high availability, enter the preconfigured floating host name to share between the servers.
7. Specify the information for the database to use for the metadata repository.

If you plan to install a new DB2 database system, specify the directory where you want to install it. If you are configuring an existing DB2 instance, specify the location of the existing DB2 database system installation and the port number for connecting to the instance.

a. **AIX Linux**: Specify the DB2 instance owner name. The instance owner controls all the DB2 processes and owns all the file systems and devices that are used by the databases within the DB2 instance.
   - If you created the instance owner account before starting the installation program, click **Existing user** and specify the DB2 instance owner name. Also specify a port number for the new instance.
   - If you did not create the instance owner account, click **Create a user as an instance owner** and specify the user information. This option is not available if the user registry is Lightweight Directory Access Protocol (LDAP) or NIS.

b. **AIX Linux**: Specify the DB2 fenced user name. Any fenced user-defined functions (UDFs) and stored procedures run under this user and group.
   - If you created the fenced user name before starting the installation program, click **Existing user** and specify the DB2 fenced user name.
   - If you did not create the fenced user name, click **Create a new user as a fenced user** and specify the user information. This option is not available if the user registry is LDAP or NIS.

c. Enter the metadata repository owner information and database name.

8. Specify the following additional information for the metadata repository database:
   a. User name and password for the staging area user (xmetasr). The staging area is a schema within the metadata repository database that stores staged imports for IBM InfoSphere Metadata Asset Manager.
9. Depending on the product modules you selected, the installation wizard will ask for information for the repository tier data stores for the product modules (for example, the operations database for IBM InfoSphere DataStage and the standardization rules designer database for IBM InfoSphere QualityStage). The installation program will install them for you if you chose to install the DB2 database system.

   If you decide to install them on another computer or if you are using a different database system, then when you install the services tier, the databases are registered and SQL scripts are provided for you to create the supporting data stores after installation.

Specifying installation options (services tier)

Use the InfoSphere Information Server installation program to specify options for the service tier computer, such as the location where you want to install IBM WebSphere Application Server, the property information for your database system, and configuration information for other product modules.

Procedure

1. Select the directories where you want to install the product that you determined when planning your installation.
2. Select **New installation**.
3. Select the tiers to install on the target computer.
4. Enter the InfoSphere Information Server administrator user name and password. Use the information that you determined when planning your installation.

5. Select the product modules that you want to install. The product modules that are displayed are only those that are part of the entitlement package that you purchased. Select the same set of product modules for each computer.
6. If you are installing InfoSphere DataStage, choose the product edition and optional features that you want to install.
7. If you are configuring your topology for high availability, enter the preconfigured floating host name to share between the servers.
8. Select the application server:
   - You can choose to install IBM WebSphere Application Server Network Deployment:
     a. Specify the installation location for the application server.
     b. Select **Customize WebSphere Application Server profile ports** if you want to define custom port assignments for WebSphere Application Server within the profile.
     c. Enter the server administrator user name and password.
   - You can choose to use an existing installation of IBM WebSphere Application Server Network Deployment:
     a. Specify the application server instance to use.
     b. Specify an application server profile to use. You must have at least one profile available for installation. If you are installing with a stand-alone application server, then the profile that you selected for installation must not have security configured and the profile must be empty. If the installation program detects that the type of the selected profile is *dmgr*, it installs InfoSphere Information Server on a WebSphere Application Server cluster. For a clustered installation, the deployment manager must have security already configured.
     c. In this case of a cluster, you select the cluster on which to install InfoSphere Information Server and specify information about the virtual host and web server to use. If you have configured a front-end web server before installation, you can choose that; otherwise, you can choose to install against a cluster member and configure the front-end later. For more information, see "Implementing a WebSphere Application Server cluster for InfoSphere Information Server" on page 151.
     d. Enter the server administrator user name and password.
   - You can choose to install IBM WebSphere Application Server Liberty Profile:

     Because of convention, such as in support documents, it is recommended to accept the default port numbers unless they conflict with existing ports on your computer. The SSL information that you provide is used to create an SSL certificate for secure connections to the application server. The installation program creates the certificate for you. For the **SSL subject distinguished name**, you can provide information that is meaningful to your organization. The certificate information is used to provide secure communication with the application and application server. The distinguished name has the format:

     `CN=Common_Name,OU=Org_Unit,O=Org_Name,C=Country_name`

     Example:

     `CN=analytics.example.com,OU=Sales Division,O=Example Co.,C=US`

     Set the Common Name (CN) field to the value of the host name of the services tier computer, to be used by remote clients to access the server. As

part of the SSL handshake, clients verify that the host name that is used to access the services tier matches the certificate CN value (or one of the values if there are more than one).

Specify any length of days for the certificate to be valid from the time that it is created (for example, 365). When the certificate expires, clients might no longer be able to connect until you update the certificate.

9. Enter the InfoSphere Information Server administrator user name and password.
10. If the metadata repository is installed on a separate computer, specify the connection information for the metadata repository.
11. If the metadata repository is installed on a separate computer, then depending on the product modules you selected, the installation wizard will ask you for connection information for the modules' supporting data stores (for example, the operations database for IBM InfoSphere DataStage and the standardization rules designer database for IBM InfoSphere QualityStage). The databases are registered and SQL scripts are provided for you to create the supporting data stores after installation.

Specifying installation options (engine tier)

Use the IBM InfoSphere Information Server installation program to specify options for the engine tier computer, such as the ports to use for the InfoSphere Information Server ASB agent (communication agent), as well as connection details for other product modules.

Procedure

1. Select the directories where you want to install the product that you determined when planning your installation.
2. Select **New installation**.
3. Select the tiers to install on the target computer.
   - If you are installing an engine on this computer but are not installing the services tier, specify the connection information for the services tier.
   - If your services tier is configured in a clustered implementation of IBM WebSphere Application Server, specify the host name and port number of the front-end HTTP dispatcher. Otherwise, specify the host name of the services tier computer and the port number for WebSphere Application Server (typically 9443).
4. Enter the InfoSphere Information Server administrator user name and password. Use the information that you determined when planning your installation.
5. Select the product modules that you want to install. The product modules that are displayed are only those that are part of the entitlement package that you purchased. Select the same set of product modules for each computer.
6. If you are configuring your topology for high availability, enter the preconfigured floating host name to share between the servers.
7. Enter the InfoSphere Information Server administrator user name and password.
8. Specify the ports to use for the InfoSphere Information Server ASB agent (communication agent). The ASB agent facilitates communication between tiers.
9. Specify connection details for additional InfoSphere Information Server components.
   a. Specify the InfoSphere DataStage administrator user name.

- If you created the administrator user before starting the installation program, click **Existing user** and specify the user name.
- If you did not create the administrator user, click **Create a user as the IBM InfoSphere DataStage administrator** and specify the user information.

b. Specify two TCP/IP ports to use for the IBM InfoSphere DataStage and QualityStage job monitor, which takes snapshots of job performance. The job monitor uses one port to request job information and the second port to report job information. The default port settings are acceptable in most configurations.

c. If an installed instance of the InfoSphere Information Server engine exists on the target computer, and you want to install another engine instance on the same computer, specify the tag information for the engine.

d. If you want this installation to handle data in languages other than English, install globalization support. Some server jobs might run more slowly after you install globalization, but the performance of parallel jobs is not affected. The code page that is currently configured on the target computer must be supported by InfoSphere Information Server. If the code page is not supported, change to a supported code page and restart the installation program.

   Important: You cannot change the globalization support after installation.

e. If you have existing jobs that depend on the IBMWebSphere MQ Plug-in, choose to install the plug-in. For new jobs, you can use the IBMWebSphere MQ Connector.

f. If you have existing jobs that use SAS, enable a legacy SAS configuration and select the SAS that you want to configure.

g. Specify the InfoSphere DataStage projects to create. You must create at least one project during installation. The installation wizard defaults to creating a project called “dstage1.”

10. Specify the ports for the metadata interchange agent to use. You must specify the HTTPS port number for the engine tier computer. The default HTTPS port is **19443**.
11. If the metadata repository is installed on a separate computer, specify the connection information for the metadata repository.
12. If the metadata repository is installed on a separate computer, then depending on the product modules you selected, the installation wizard will ask you for connection information for the modules' supporting data stores (for example, the operations database for IBM InfoSphere DataStage and the standardization rules designer database for IBM InfoSphere QualityStage). The databases are registered and SQL scripts are provided for you to create the supporting data stores after installation.

Specifying installation options (client tier)

Use the InfoSphere Information Server installation program to specify options for the client tier computer, such as the property information for your database system.

Procedure

1. Select the directories where you want to install the product that you determined when planning your installation.
2. Select **New installation**.
3. Select the tiers to install on the target computer.

- If you are installing a client on this computer but are not installing the services tier, specify the connection information for the services tier.
- If your services tier is configured in a clustered implementation of IBM WebSphere Application Server, specify the host name and port number of the front-end HTTP dispatcher. Otherwise, specify the host name of the services tier computer and the port number for WebSphere Application Server (typically 9443).

4. Enter the InfoSphere Information Server administrator user name and password. Use the information that you determined when planning your installation.
5. Select the product modules that you want to install. The product modules that are displayed are only those that are part of the entitlement package that you purchased. Select the same set of product modules for each computer.

   Select **Metadata Interchange Agent** if you want to use bridges with IBM InfoSphere Metadata Asset Manager. Installing the metadata interchange agent automatically installs bridges on the same computer.

Running the installation step

After the installation program completes the prerequisites check and you specify installation options for each tier, you specify a response file and proceed with the installation process. The response file contains configuration information that tells the installation program how to perform an installation.

About this task

Saving a response file is important if you want to cancel the current installation session and use the response file in a later session. Use the installation wizard to specify a name and location for your response file.

Procedure

1. In the installation wizard, specify a name and location for the response file in the **Response file location** field (at this point, you can choose to cancel the installation and use the saved response file in a later session).
2. Choose to install the product and click **Next** to begin the product installation. The program displays a summary of the choices that you made in the wizard. As the installation proceeds, the installation program reports the progress in the window.

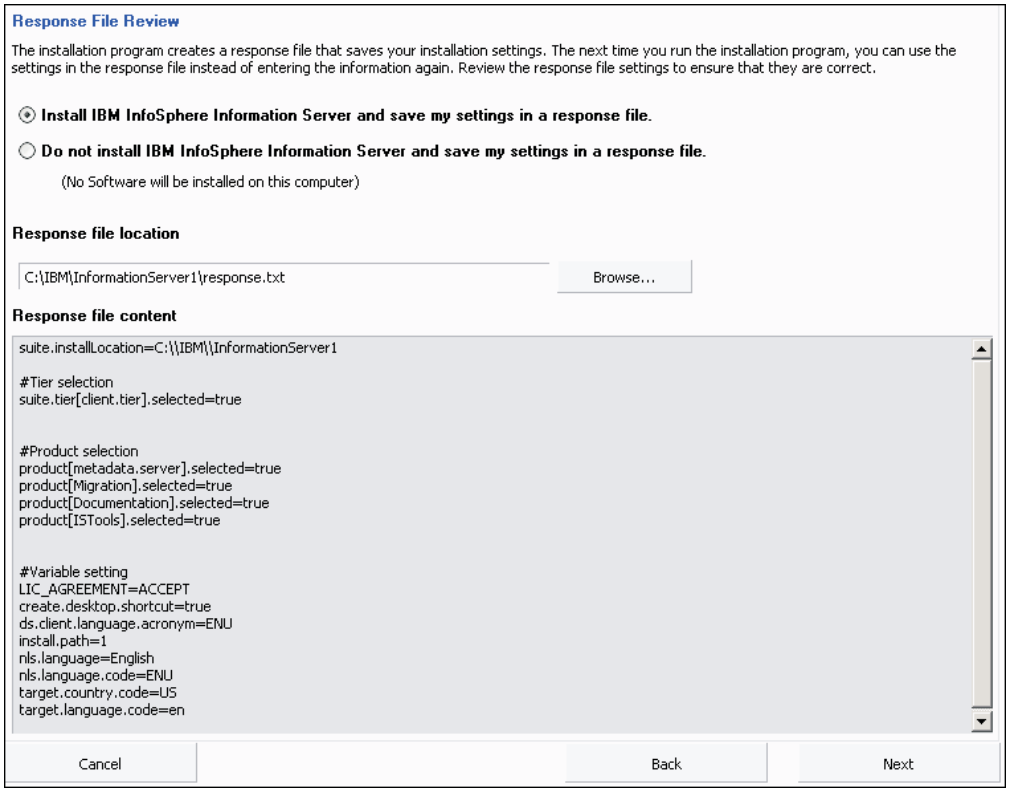


Figure 45. Response File Review page

Be mindful of the following conditions while the installation program is running:

- Leave the command prompt open until the installation is complete, as shown in the following figure.

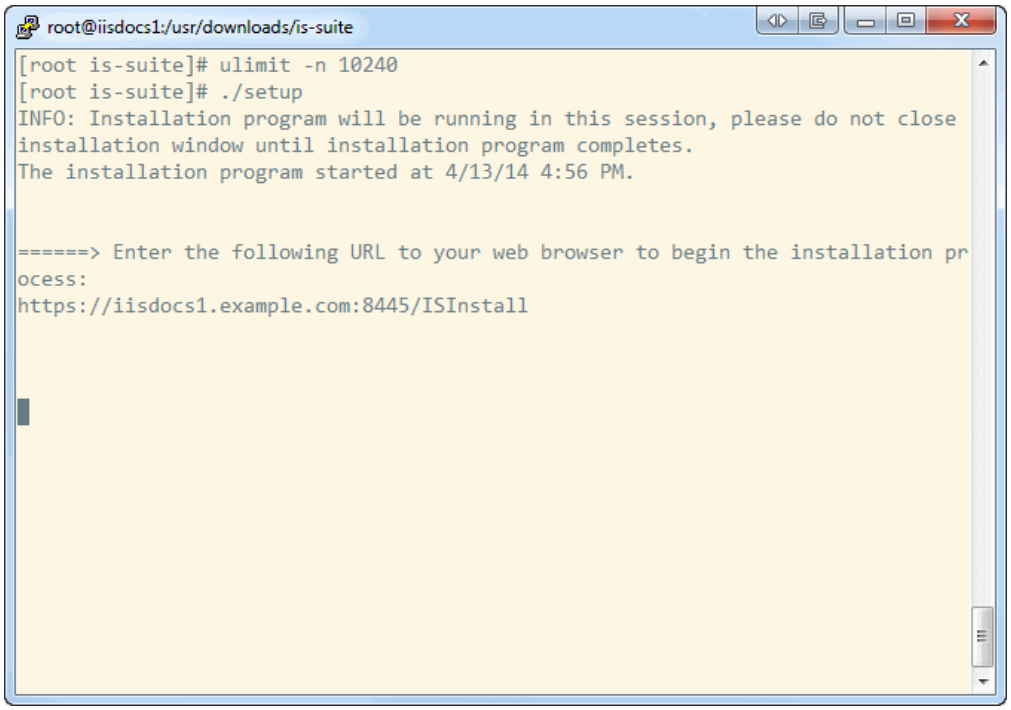


Figure 46. Command prompt open during installation

- If the system encounters an error, it might occasionally request a response from you. Check periodically to make sure that the system is not waiting for you to respond.
- Pop-up windows might be displayed, that you must respond to for the installation to proceed. Occasionally a pop-up window might be displayed behind the web browser. If the installation halts for a long interval, check for pop-up windows displayed behind the web browser. You can click on the navigation line on the left side to bring the exception window to the front.
- After you have started the installation process, keep the web browser open to monitor progress and to respond if there is an error. If you close the web browser, the installation continues to the end. You can reconnect to the session by opening a new web browser to the same URL. You can also monitor the installation process by reading the log files.
- Depending on your security settings, warning windows might display during the installation of the IBM DB2 database system or other components. Ignore the warnings, and click **OK** to continue the installation.
- Do not install InfoSphere Information Server in a directory path that contains space characters.

Taking action if the installation program encounters an error

After you start the installation process, the installation program might encounter an error that requires your action to resolve.

Symptoms

The installation program encounters an error and an Installation Exception window is displayed.

User response

Complete these steps in the Installation Exception window:

1. If the error window is still displayed with the option to Abort, Retry, or Continue the installation leave the window up. Likewise, if you are running the installation program as a wizard and you have not closed it, do not close it.
2. Review the last statements in the log file to determine what caused the error. Investigate and fix the error, if possible. Refer to the troubleshooting topics or the support tech notes online and take the specified user action.
3. Retry the installation:
   - If you did not close the installation program, click **Retry failed action**. If the issue is resolved, then it will proceed. Otherwise, if it runs into the same exception, you will see the exception window again.
   - If you did close the installation program, run it again. If you were able to save the results in a response file, run the installation program with the response file so that you do not need to re-enter the parameters. Use the following command if you want to run the installation program with the wizard and response file:

     ```
     setup -rsp response_file_name -uiMode graphical
     ```
4. If the subsequent attempt fails, manually uninstall the product files and try to install again.
5. If failure still occurs, contact IBM support for assistance.

Completing the installation process

After the installation program has collected your selections, it runs another, more detailed prerequisites check based on your selections. During this check, it analyzes your computer resources and file system to determine whether the installation is likely to succeed.

About this task

The wizard page displays each check. If the check fails, it is marked FAILED in the page. If the check found issues that might cause problems with the installation, it is marked WARNING in the page.

When the installation process is complete, an Installation Summary for Current Computer page is displayed that contains the product, tiers, and product components that you installed, as well as information on how long the installation program ran and whether the installation was successful.

If the installation process was successful, the following message is written to the log file:

```
Installation Engine completed all steps successfully.
```

Procedure

1. Click **Finish** or close the browser window.
2. Restart the computer.

Installing IBM InfoSphere Information Server in silent mode by using the response file

You can use silent mode to run an unattended installation of InfoSphere Information Server. In silent mode, the installation program does not display a user interface. Instead, it reads settings from a response file that you create, performs a prerequisites check, and installs the software if the check succeeds.

About this task

You might run a silent installation with a response file if you want to reinstall the software after cleaning up a failed installation. Or, you might have multiple engine computers and don't want to run through the wizard interview process on each computer.

To run a silent installation, you first run through the installation wizard and save your responses in the response file at the end of the interview process. Before running the `setup` program in silent mode with the response file, you edit the response file providing password values and changing any values that are specific for the computer that you are installing the software on. Then you run `setup -rsp` *`response_file_name`* to run in silent mode with the response file.

Creating a response file

A response file is required to run the installation program in silent mode. You use the installation program to create a response file.

Procedure

1. Create a response file.

Option	Description
Using the installation program	1. Run the installation program or software removal program in graphical mode. 2. Specify your parameters in the wizard. 3. When the Response File Review page displays, select either **Install IBM InfoSphere Information Server and save my settings in a response file** or **Do not install IBM InfoSphere Information Server and save my settings in a response file**. Either selection causes a response file to be created. 4. Specify a file name and location and click **Next**. The program saves your parameters in a response file in the specified location. The program then either performs the operation or exits, depending upon your selections in the window.

2. Optional: Run the **`genkey`** command to generate encrypted information. The **`genkey`** command is located in the tools directory on your installation media:

Operating system	Command
AIX Linux	*install_media*/tools/genkey.sh -generate *info_to_encrypt* Example: /tmp/isinstall/is-suite/tools/genkey.sh -generate Mypa55word
Windows	*install_media*\tools\genkey.cmd -generate *info_to_encrypt* Example: C:\tmp\dcinstall\is-suite\tools\genkey.cmd -generate Mypa55word

Messages are displayed in the terminal window. Then the command displays the encrypted version of the specified information. For example:

```
{xor}8Y3lImYJo0sxGOl4igw/pA==
```

3. Locations in the response file where passwords must be included are indicated with markers. Change all password markers to your plain text or encrypted passwords. If you use an encrypted password, append `.encrypted` to the name. For example:

   The name-value pair generated by the installation program:

   ```
   xmeta.db.password=<CHANGE_TO_YOUR_PASSWORD>
   ```

   Replacing the value with a plain text password:

   ```
   xmeta.db.password=mYPa$$w0rd
   ```

   Replacing the value with an encrypted password:

   ```
   xmeta.db.password.encrypted={xor}8Y3lImYJo0sxGOl4igw/pA==
   ```

4. Save and store the response file for use when installing InfoSphere Information Server.

Encrypting information in a response file

You can encrypt passwords and other information in response files.

About this task

By default, information is stored in plain text in response files. If you want to encrypt passwords or other information in your response files, follow this procedure to encrypt information and include it in the files.

You can encrypt any values in a response file. However, this procedure is typically used to encrypt passwords.

Procedure

1. Create the response file.
2. Run the **genkey** command to generate encrypted information. The **genkey** command is located in the tools directory on your installation media:

Operating system	Directory
AIX	*install_media*/tools/genkey.sh *info_to_encrypt*
Linux	*install_media*/tools/genkey.sh *info_to_encrypt*
Windows	*install_media*\tools\genkey.cmd *info_to_encrypt*

For example:

```
C:\is-suite\tools\genkey.cmd -generate mypassword
```

Messages are displayed in the terminal window. Then, the command displays the encrypted version of the specified information. For example:

```
{xor}8Y3lImYJo0sxGOl4igw/pA==
```

3. In the response file, append the extension `.encrypted` to the parameter. Replace the value with the encrypted version. For example:
   - Original parameter specification:

     ```
     app.server.password=<CHANGE_TO_YOUR_PASSWORD>
     ```
   - Encrypted parameter specification:

     ```
     app.server.password.encrypted={xor}8Y3lImYJo0sxGOl4igw/pA==
     ```

About the response files

A response file contains parameters and values that tell the IBM InfoSphere Information Server installation program how to perform an installation. You can also use a response file during software removal. Create a response file by using the InfoSphere Information Server installation or software removal program.

Different sets of parameters and values are used depending on the type of software installation or removal. For example, a certain set of parameters and values is used for an installation that includes application server clustering. Another set is used for an installation where the application server is set up in a stand-alone configuration.

Response file syntax

A response file is a text file. In the file, parameters and values are specified one per line. In each line, the parameter is separated from the value with an equals sign (=). A line that begins with a pound sign (#) indicates a comment. Parameters and values are case sensitive. Blank lines are ignored.

To specify a backslash in the response file, include two backslashes. For example:

```
suite.installLocation=C:\\IBM\\InformationServer
```

To set a parameter to a null value, omit the parameter value or specify the parameter without a value. For example:

```
xmeta.db.tablespace=
```

Passwords in a response file

To use a response file, you must specify passwords in the file. If the file is created by the installation or software removal program, the program does not store the passwords in the file. Instead, locations in the file where passwords must be included are indicated with markers. For example:

```
app.server.password=<CHANGE_TO_YOUR_PASSWORD>
```

Replace each marker with the plain text or encrypted password. For example:

```
app.server.password=mypassword
```

Response file parameters (Windows)

Response files vary depending on the conditions of your particular installation. You can use this reference to view details about the response file parameters that apply to your installation. Response parameters and their values are case-sensitive.

allow.customize.was.profile.ports=[true|false]
Set to *true* if you are specifying the settings for all IBM WebSphere Application Server ports instead of installing with the default ports.

app.server.password=*change_to_application_server_admin_password*
Set the password for the application server administrator.

app.server.type=[was]
Set the type of application server that is being configured.

An example value is *was*.

app.server.user=*admin user ID*
Set the user ID for the application server administrator. This ID is created in the local repository for authentication.

An example value is *wasadmin*.

DataStage=[true|false]
Set to *true* if installing IBM InfoSphere DataStage.

DataStage.BalOp=[true|false]
Set to *true* if installing IBM InfoSphere DataStage Balanced Optimization. This option can be *true* only if DataStage is *true*.

DataStage.SAS=[true|false]
Set to *true* if installing IBM InfoSphere DataStage Pack for SAS. This option can be *true* only if DataStage is *true*.

DataStage.Server=[true|false]
Set to *true* if installing IBM InfoSphere DataStage.

db.driver=
[com.ibm.db2.jcc.DB2Driver

]
Set the driver that the metadata repository uses to connect to the database server. This option controls which database to use, and is available if you are using a preinstalled database.

An example value is *com.ibm.db2.jcc.DB2Driver*.

db.url=*database_url*
Set the database URL for the metadata repository or given database. These sample values apply only if you created the database using the DatabaseSupport scripts or if the database is on a remote system. The following database URLs are allowed:

`jdbc:db2://hostname:port/database`

db2.admin.user.name=*DB2_admin_username*
Set to the operating system ID to be used as the IBM DB2 administrator ID. This value applies only if you are installing the repository tier.

An example value is *db2admin*.

db2.admin.user.password=*DB2_admin_password*
Set to the password for the DB2 administrator ID. This value applies only if you are installing the repository tier.

db2.copy.name=*DB2_copy_name*
Set to the name of this DB2 Copy. This value applies only if you are installing the repository tier.

db2.install.location=*DB2_absolute_installation_path*
Set to the absolute path where DB2 is to be installed.

An example path is: `C:\IBM\SQLLIB`.

db2.port.number=*DB2_instance_port_number*
Set to the TCPIP port for this DB2 instance. This value applies only if you are installing the repository tier.

An example value is *50000*.

ds.client.language.acronym=*[n|n1|n2...nn]*
Set to the language code for InfoSphere DataStage clients, where *n* can be any of the codes in the following table.

Table 72. Supported language codes for InfoSphere DataStage

Language	Code
English	ENU
French	FRA
German	DEU
Italian	ITA
Spanish	ESP
Portuguese-Brazil	PTB
Chinese-Simplified	CHS
Chinese-Traditional	CHT
Japanese	JPN
Korean	KOR

ds.project.1.name=*DataStage_project_name*
Set the value to the InfoSphere DataStage project name that you want. Replace the number, *1*, with the project number that increments, starting with 1.

An example value is *dstage1*.

ds.project.1.path=*path_to_DataStage_project_files*
Set the value for the path to the InfoSphere DataStage project files. Replace the number, *1*, with the project number that increments, starting with 1.

An example value is `C:\\IBM\\InformationServer\\Server\\Projects\\dstage1`.

ds.project.1.version=**New**
Replace the number with the incrementing number, starting with *1*, of the project to be created for InfoSphere DataStage.

Set the value to *New*.

Note: To make a project protected, add `(Protected)` in English, as demonstrated in the last line in the example.
For example:

```
ds.project.1.name=dstage1
ds.project.1.path=C:\\IBM\\InformationServer\\Server\\Projects\\dstage1
ds.project.1.version=New
ds.project.2.name=addedproj
ds.project.2.path=C:\\IBM\\InformationServer\\Server\\Projects\\addedproj
```

```
ds.project.2.version=New
ds.project.3.name=protect
ds.project.3.path=C:\\IBM\\InformationServer\\Server\\Projects\\protect
ds.project.3.version=New (Protected)
```

https.truststore.path=*path*
In a secure HTTP (HTTPS) installation, set to the path to the truststore file that contains the certificate to accept.

An example location is `C:\\tempMyFile.jks`.

ia.db.hostname=*Information_Analyzer_db_host_name*
Set to the IBM InfoSphere Information Analyzer database server name.

The host name is the output of the **hostname** command on the system where the response file was generated.

ia.db.location=*Information_Analyzer_db_location*
Set to the path where the InfoSphere Information Analyzer database will be stored.

An example location is `C:\\`.

ia.db.name=*Information_Analyzer_db_name*
Set to the name of the InfoSphere Information Analyzer database to be created.

An example database name is *iadb*.

ia.db.password=*Information_Analyzer_db_owner_password*
Set to the password for the InfoSphere Information Analyzer database owner.

ia.db.port=*port_number_for_DB2_instance*
Set to the port for the DB2 instance where the database is to be created. If this is a repository installation and the InfoSphere Information Analyzer database is to be created in the local DB2 instance, **ia.db.port** must use the same port as the port used for **db2.port.number**.

An example port is *50000*.

ia.db.type=**[db2]**
Set to database type for the InfoSphere Information Analyzer database.

ia.db.username=*Information_Analyzer_db_owner_ID*
Set to the operating system ID to be used as the InfoSphere Information Analyzer database owner ID.

An example value is *iauser*.

im.shared.directory=*installation_manager_shared_path*
Set to the IBM Installation Manager shared installation path. IBM Installation Manager is used to install WebSphere Application Server. Set this attribute on the services tier only.

An example value is `C:\IBM\IMShared`.

imamMbbAgent.ajp.port=**port_number**
Set to the port number of the Apache JServ Protocol connector for bridges. Set this attribute on the client tier only.

An example value is 19979.

imamMbbAgent.https.port=**port_number**
Set to the HTTPS port number for bridges. Set this attribute on the client tier only.

An example value is 19443.

imamMbbAgent.shutdown.port=port_number
Set to the port number on which the metadata interchange agent waits for requests to shut down. Set this attribute on the client tier only.

An example value is 19978.

imam.register.server=[true|false]
This attribute is used to register the metadata interchange server with InfoSphere Metadata Asset Manager. It is used only on the client tier. Set this to true only if you have already installed the services tier. If you set this attribute to true, you must also set `services.tier.host` and `services.tier.port`.

install.path=[1|2|3|4|5|6]
Set to the type of installation:

1 = New Installation

2 = Add Product

3 = Add Tier

4 = Remove Product

5 = Remove Tier

6 = Uninstall

An example value is 1.

is.admin.password=*Information_Server_Admin_password*
Set to the password for the InfoSphere Information Server Administrator ID.

is.admin.user=*Information_Server_Admin_ID*
Set to the user name that you want for the InfoSphere Information Server Administrator ID. This ID will be created in the local repository.

An example value is *admin*.

is.console.port=*port*
Set to the HTTP or HTTPS port where InfoSphere Information Server is installed.

In HTTP installation, the value is typically *9080*. In an HTTPS installation, the value is typically *9443* or *443*.

is.ha.selected=[true|false]
Set to *true* if this is an active-passive server cluster configuration installation.

is.https=[true|false]
Set to *true* if this is an HTTPS installation.

is.was.cluster=[true|false]
Set to *true* if you are installing into a WebSphere Application Server cluster configuration.

isf.agent.port=[true|false]
Set to the port number for the InfoSphere Information Server Agent.

An example value is *31531*.

isf.http.host=*host_name*
Set the host name of the front-end dispatcher for a cluster installation.

This value is the short form of *hostname*, excluding the domain. (For example, only the *HTTP_host_name* portion of *HTTP_host_name.domain*.com is used.)

isf.http.port=*port_number*
Set the HTTP or HTTPS port number of the front-end dispatcher for a cluster installation.

An example value is *80*.

isf.server.host=*ISF_server_host_name*
Set the server host name for a cluster environment.

This value is the complete form of *hostname*, including the complete domain.

jobmon.port.1=*port_number_for_first_job_monitor_port*
Set to the port number for the first job monitor port.

An example value is *13400*.

jobmon.port.2=*port_number_for_second_job_monitor_port*
Set to the port number for the second job monitor port.

An example value is *13401*.

LIC_AGREEMENT=ACCEPT
Set to *ACCEPT* to accept the license agreement.

mq.install.choice=[NONE|SERVER|CLIENT]
Set to *SERVER* if you want to install the WebSphere MQ Server plug-in.

Set to *CLIENT* if you want to install the WebSphere MQ Client plug-in.

nls.enabled=[true|false]
Set to true if you want globalization support for InfoSphere DataStage and IBM InfoSphere QualityStage.

The value *false* is supported only for systems that have the **nls.language** parameter set to *English* and a supported system code page.

nls.language=[*English-UTF8*|*English*|*Chinese-Simplified-UTF8*|*Chinese-Simplified*|*Chinese-Traditional-UTF8*|*Chinese-Traditional*|*French-UTF8*|*French*|*German-UTF8*|*German*|*Italian-UTF8*|*Italian*|*Japanese-EUC*|*Japanese-SJIS*|*Japanese-UTF8*|*Japanese*|*Korean-UTF8*|*Portuguese-Brazilian-UTF8*|*Portuguese-Brazilian*|*Spanish-UTF8*|*Spanish*|*Thai*]
Set to the language to configure for InfoSphere DataStage and InfoSphere QualityStage. The installation computer code page must be compatible and supported by InfoSphere DataStage.

An example value is English.

nls.language.code=*[n|n1|n2...nn]*
Set to the three-letter language code for InfoSphere DataStage and InfoSphere QualityStage support, where *n* can be any of the codes in the following table.

Table 73. Supported language codes

Language	Code
English	ENU
French	FRA
German	DEU
Italian	ITA
Spanish	ESP
Portuguese-Brazil	PTB
Chinese-Simplified	CHS

Table 73. Supported language codes (continued)

Language	Code
Chinese-Traditional	CHT
Japanese	JPN
Korean	KOR

product[AVI].selected=[true|false]
Set to *true* if installing IBM InfoSphere QualityStage Address Verification Interface.

product[Blueprint].selected=[true|false]
Set to *true* if installing IBM InfoSphere Blueprint Director.

product[BusinessGlossary].selected=[true|false]
Set to *true* if installing IBM InfoSphere Information Governance Catalog.

product[BusinessGlossaryAnywhere].selected=[true|false]
Set to *true* if installing IBM InfoSphere Glossary Anywhere.

product[DataQualityConsole].selected=[true|false]
Set to *true* if installing IBM InfoSphere Data Quality Console.

product[datastage].selected=[true|false]
Set to *true* if installing InfoSphere DataStage.

product[DPID].selected=[true|false]
Set to *true* if installing IBM InfoSphere QualityStage Module for DPID.

product[FastTrack].selected=[true|false]
Set to *true* if installing IBM InfoSphere FastTrack.

product[IMAM].selected=true
Set to *true* to install IBM InfoSphere Metadata Asset Manager.

product[IMAMEngine].selected=true
Set to *true* to install IBM InfoSphere Metadata Asset Manager. Set this attribute on the client tier only.

product[InformationAnalyzer].selected=[true|false]
Set to *true* if installing InfoSphere Information Analyzer.

product[ISALite].selected=[true|false]
Set to *true* if installing IBM Support Assistant Lite. It is recommended to leave this set to `true`.

product[ISTools].selected=[true|false]
Set to *true* if installing InfoSphere Information Server Manager.

product[metadata.server].selected=[true]
Set to *true* for installing IBM InfoSphere Metadata Server.

product[MetadataWorkbench].selected=[true|false]
Set to *true* if installing IBM InfoSphere Information Governance Catalog.

product[qualitystage].selected=[true|false]
Set to *true* if installing InfoSphere QualityStage.

product[Recovery].selected=[true|false]
Set to *true* if installing backup and recovery software. It is recommended to leave this set to `true`.

product[SERP].selected=[true|false]
Set to *true* if installing IBM InfoSphere QualityStage Module for SERP.

product[WISD].selected=[true|false]
Set to *true* if installing IBM InfoSphere Information Services Director.

srd.db.location=*location_path*
Set to the path for the IBM InfoSphere QualityStage Standardization Rules Designer database. This value applies only if you are installing the repository and services tiers.

srd.db.name=*database_name*
Set to the name of the database that the IBM InfoSphere QualityStage Standardization Rules Designer data store is in. This value applies only if you are installing the repository and services tiers.

srd.db.schemaName=*schema_name*
Set to the name of the IBM InfoSphere QualityStage Standardization Rules Designer schema. This value applies only if you are installing the repository and services tiers.

srd.db.tablespace=*tablespace_name*
Set to the tablespace name of the IBM InfoSphere QualityStage Standardization Rules Designer data store. This value applies only if you are installing the repository and services tiers.

srd.db.user.name=*owner_ID*
Set to the owner ID for the IBM InfoSphere QualityStage Standardization Rules Designer data store. This value applies only if you are installing the repository and services tiers.

srd.db.user.password=*password*
Set to the password of the IBM InfoSphere QualityStage Standardization Rules Designer data store owner ID. This value applies only if you are installing the repository and services tiers.

suite.installLocation=*IS_installation_path*
Set to the absolute directory path where InfoSphere Information Server is to be installed.

An example value is `C:\\IBM\\InformationServer`.

Do not use spaces in the path name.

suite.tier[app.server.tier].selected=[true|false]
Set to *true* if installing the Services tier.

suite.tier[client.tier].selected=[true|false]
Set to *true* if installing the Client tier.

suite.tier[database.tier].selected=[true|false]
Set to *true* if installing the Repository tier.

suite.tier[engine.tier].selected=[true|false]
Set to *true* if installing the Engine tier.

target.country.code=*[n|n1|n2...nn]*
Set the two-letter country or region code for the InfoSphere Information Server installation, where *n* can be any of the codes in the following table.

Table 74. Supported country or region codes

Language	Code
Brazil	BR
China	CN
France	FR

Table 74. Supported country or region codes (continued)

Language	Code
Germany	DE
Italy	IT
Japan	JP
Korea	KR
Spain	ES
Taiwan	TW
United States	US

target.language.code=*[n|n1|n2...nn]*
Set the language code for the InfoSphere Information Server installation, where *n* can be any of the codes in the following table.

Table 75. Supported language codes

Language	Code
English	en
French	fr
German	de
Italian	it
Spanish	es
Portuguese-Brazil	pt_BR
Chinese-Simplified	zh
Chinese-Traditional	zh_TW
Japanese	ja
Korean	ko

use.bundled.db2=[true|false]
Set to *true* if installing the DB2 database that is bundled with InfoSphere Information Server.

use.bundled.was=[true|false]
Install the IBM WebSphere Application Server Network Deployment that is bundled with InfoSphere Information Server.

was.admin.host.port=*was_console_port_number*
Set to the port for the WebSphere Application Server administrative console.

An example value is *9060*.

was.bootstrap.port=*websphere_bootstrap_port_number*
Set to the WebSphere bootstrap port.

An example value is *2809*.

was.dcs.unicast.port=*was_ha_manager_port_number*
Set to WebSphere high availability manager communication port.

An example value is *9353*.

was.default.host.port=*was_http_transport_port_number*
Set to the port number for the HTTP transport port.

An example value is *9080*.

was.install.location=*was_installation_location*
Set to the directory where installing WebSphere Application Server. The directory must be empty.

An example value is `C:\\IBM\\WebSphere\\AppServer`.

was.ipc.connector.port=*port_number*
Set to the port number for the WebSphere Application Server Inter-Process Communications connector. The directory must be empty.

An example value is 9633.

was.mutual.auth.listener.port=*was_mutual_bootstrap_port_number*
Set to the port number for the WebSphere mutual bootstrap port.

An example value is *9402*.

was.orb.listener.port=*was_orb_listener_port_number*
Set to the port number for the WebSphere ORB listening port.

An example value is *9100*.

was.profile.dir=*was_profile_directory*
Set the WebSphere profile directory.

An example value is `C:\\IBM\\WebSphere\\AppServer\\profiles\\InfoSphere`.

was.profile.name=*was_profile_name*
Set the WebSphere profile name.

An example value is *InfoSphere*.

was.sas.ssl.port=*SAS_SSL_port*
Set to the port number for the SAS SSL Server Authentication listener.

An example value is *9401*.

was.secure.admin.host.port=*secure_was_port*
Set to the port number for the secure WebSphere Application Server administrative console.

An example value is *9043*.

was.secure.default.host.port=*default_secure_HTTP_transport_port_number*
Set to the port number for the default secure HTTP transport port.

An example value is *9443*.

was.secure.sib.mq.port=*secure_service_integration_MQ_port_number*
Set to the port number for the secure service integration WebSphere MQ interoperability port.

An example value is *5578*.

was.secure.sib.port=*secure_service_integration_port_number*
Set to the port number for the secure service integration port.

An example value is *7286*.

was.server.auth.listener.port=*CSIV2_server_auth_port_number*
Set to the port number for the CSIV2 server authentication port.

An example value is *9403*.

was.server.name=*was_server_name*
Set the WebSphere Application Server server name.

An example value is *server1*.

was.sib.mq.port=*service_integration_MQ_port_number*
Set to the port number for the service integration WebSphere MQ interoperability port.

An example value is *5558*.

was.sib.port=*service_integration_port_number*
Set to the port number for the service integration port.

An example value is *7276*.

was.sip.port=*SIP_port_number*
Set to the port number for the SIP port.

An example value is *5060*.

was.sip.secured.port=*SIP_secured_port_number*
Set to the port number for the secured SIP port.

An example value is *5061*.

was.soap.connector.port=*SOAP_connector_port_number*
Set to the port number for the SOAP connector port.

An example value is *8880*.

xmeta.db.config.type=[*db2_10_5_nls*]
To set the database configuration type for metadata repository.

xmeta.db.hostname=*metadata_repository_server_host_name*
Set to the host name for the metadata repository database server.

An example value is *myhostname.mydomain.com*.

xmeta.db.location=*metadata_repository_db_location*
Set to the drive on Windows where you want the metadata repository database to be created.

An example value is `C:\\`.

xmeta.db.name=*metadata_repository_db_name*
Set to the database name for the metadata repository.

An example value is *xmeta*.

xmeta.db.password=*metadata_repository_db_owner_password*
Set to the password for the metadata repository database owner.

xmeta.db.port=*metadata_repository_db_port*
Set to the port for the DB2 instance where the database is to be created. The port for **xmeta.db.port** must be the same as **db2.port.number** if this is a repository installation and the InfoSphere Information Analyzer database is to be created in the local DB2 instance.

An example value is *50000*.

xmeta.db.type=**[db2]**
Set the database type for the metadata repository.

xmeta.db.username=*metadata_repository_db_owner_ID*
Set to the operating system ID to be used as the metadata repository database owner ID.

An example value is *xmeta*.

xmeta.driver.type=[*DB2_10*]
Set the driver type for the metadata repository.

xmeta.odb.db.hostname=*host_name*
Set to the host name of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.location=*location_path*
Set to the path of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.name=*database_name*
Set to the name of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.port=*port_number*
Set to the port number of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.registration=*[true|false]*
Set to true if the IBM InfoSphere DataStage and QualityStage operations database repository is to be registered in the metadata repository. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.schemaName=*schema_name*
Set to the schema name of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.tablespace=*tablespace_name*
Set to the tablespace name of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.user.name=*owner_id*
Set to the owner ID of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.user.password=*password*
Set to the password of the IBM InfoSphere DataStage and QualityStage operations database owner ID. This value applies only if you are installing the repository and services tiers.

xmeta.staging.db.password=*password*
Set to the password of the IBM InfoSphere Metadata Asset Manager staging area owner ID. This value applies only if you are installing the repository and services tiers.

xmeta.staging.db.username=*owner_ID*
Set to the owner ID of the IBM InfoSphere Metadata Asset Manager staging area. This value applies only if you are installing the repository and services tiers.

Response file parameters (AIX Linux)

Response files vary depending on the conditions of your particular installation. You can use this reference to view details about the file response parameters that apply to your installation. Response parameters and their values are case sensitive.

allow.customize.was.profile.ports=[true|false]
Set to true if you are specifying the settings for all IBM WebSphere Application Server ports instead of installing with the default ports.

app.server.password=*change_to_application_server_admin_password*
Set the password for the application server administrator.

app.server.type=[was]
Set this property to the type of application server that is being configured.

An example value is was.

app.server.user=*admin user ID*
Set to the user ID for the application server administrator. This ID will be created in the local repository for authentication.

An example value is wasadmin.

DataStage=[true|false]
Set to true if installing IBM InfoSphere DataStage.

DataStage.BalOp=[true|false]
Set to true if installing IBM InfoSphere DataStage Balanced Optimization. This option can be true only if DataStage is true.

DataStage.SAS=[true|false]
Set to true if installing IBM InfoSphere DataStage Pack for SAS. This option can be true only if DataStage is true.

DataStage.Server=[true|false]
Set to true if installing IBM InfoSphere DataStage.

db.driver=
[com.ibm.db2.jcc.DB2Driver

]
Set the driver to be used by the metadata repository for connection. This option controls which database to use, and is available if you are using a preinstalled database.

db.url=*database_url*
Set the database URL for the metadata repository or given database. These sample values apply only if you created the database using the DatabaseSupport scripts or if the database is on a remote system. The following database URLs are allowed:

jdbc:db2://*hostname*:*port*/*database*

db2.admin.group.name=*DB2_admin_group_name*
Set to the IBM DB2 administrator group name.

An example value is dasadm.

db2.admin.home.directory=*DB2_admin_home_directory*
Set to the DB2 administrator home directory.

An example value is /home/dasusr.

db2.admin.user.name=*DB2_admin_username*
Set to the operating system ID to be used as the DB2 administrator ID. This value applies only if you are installing the repository tier.

An example value is dasuser.

db2.fenced.group.name=*DB2_fenced_group_name*
Set to the DB2 fenced group name.

An example value is db2fadm.

db2.fenced.home.directory=*DB2_fenced_home_directory*
Set to the DB2 fenced home directory.

An example path is /home/db2fenc1.

db2.fenced.user.name=*DB2_fenced_user_name*
Set to the DB2 fenced user name.

An example value is db2fenc1.

db2.fenced.user.password=*change_to_DB2_fenced_password*
Set to the DB2 fenced user password

db2.install.location=*DB2_absolute_installation_path*
Set to the absolute path where DB2 is to be installed.

An example value is /opt/IBM/db2.

db2.instance.group=*DB2_instance_group_name*
Set to the group name of the installed DB2 instance.

An example value is db2iadm.

db2.instance.home.directory=*DB2_instance_home_directory*
Set to the directory of the installed instance of DB2.

An example value is /home/db2inst1.

db2.instance.name=*DB2_instance_name*
Set to a string that will identify the DB2 instance. The DB2INSTANCE environment variable will be set to this DB2 instance name by default.

An example value is db2inst1.

db2.instance.user.name=*DB2_instance_user_name*
Set to the user name of the instance of DB2.

An example value is db2inst1.

db2.instance.user.password=*change_to_DB2_instance_user_password*
Set to the password of the DB2 instance user.

db2.port.number=*DB2_instance_port_number*
Set to the TCPIP port for this DB2 instance. This value applies only if you are installing the repository tier.

An example value is 50000.

ds.admin.gid=*DataStage_admin_group_name*
Set to the InfoSphere DataStage administrator group name.

An example value is dstage.

ds.admin.uid=*DataStage_admin_username*
Set to the InfoSphere DataStage administrator user name.

An example value is dsadm.

ds.client.language.acronym=[ENU|FRA|DEU|ITA|ESP|PTB|CHS|CHT|JPN| KOR]
Set to the language code for InfoSphere DataStage clients.

Table 76. Supported language codes for InfoSphere DataStage

Language	Code
English	ENU
French	FRA
German	DEU
Italian	ITA
Spanish	ESP
Portuguese-Brazil	PTB
Chinese-Simplified	CHS
Chinese-Traditional	CHT
Japanese	JPN
Korean	KOR

An example value is ENU.

ds.project.1.name=*DataStage_project_name*
Set the value to the InfoSphere DataStage project name that you want. Replace the number, *1*, with the project number that increments, starting with 1.

An example value is dstage1.

ds.project.1.path=*path_to_DataStage_project_files*
Set the value for the path to the InfoSphere DataStage project files. Replace the number, *1*, with the project number that increments, starting with 1.

An example value is `/opt/IBM/InformationServer/Server/Projects/dstage1`.

ds.project.1.version=New
Replace the number with the incrementing number, starting with *1*, of the project to be created for InfoSphere DataStage.

Set the value to New for this release.

Note: To make a project protected, add `(Protected)` in English, as demonstrated in the last line in the following example.
For example:

```
ds.project.1.name=dstage1
ds.project.1.path=/opt/IBM/InformationServer/Server/Projects/dstage1
ds.project.1.version=New
ds.project.2.name=addedproj
ds.project.2.path=/opt/IBM/InformationServer/Server/Projects/addedproj
ds.project.2.version=New
ds.project.3.name=protect
ds.project.3.path=/opt/IBM/InformationServer/Server/Projects/protect
ds.project.3.version=New (Protected)
```

ds.rpc.instance.tag=*ITAG_value*
Set to the InfoSphere DataStage rpc instance tag (ITAG). The ITAG value is a three-digit hexadecimal. ACE, ADE, and DCD are reserved values. The ITAG value must be unique on the system.

ADE is the default value if you do not make this an ITAG installation. If you specify a value for an ITAG installation, use a value other than ACE, ADE, or DCD.

dsrpc.port=*DataStage_dsrpc_port_number*
Set to the InfoSphere DataStage dsrpc port.

An example value is 31538.

https.truststore.path=*path*
In a secure HTTP (HTTPS) installation, set to the path to the truststore file that contains the certificate to accept.

An example location is `tmp/MyFile.jks`.

ia.db.hostname=*Information_Analyzer_db_host_name*
Set to the IBM InfoSphere Information Analyzer database server name.

The host name is the output of the **hostname** command on the system were the response file was generated.

ia.db.location=*Information_Analyzer_db_location*
Set to the path where the InfoSphere Information Analyzer database will be stored.

An example location is `/opt/IBM/InformationServer/Repos/iadb`.

ia.db.name=*Information_Analyzer_db_location*
Set to the name of the InfoSphere Information Analyzer database to be created.

An example database name is iadb.

ia.db.password=*Information_Analyzer_db_owner_password*
Set to the password for the InfoSphere Information Analyzer database owner.

ia.db.port=*port_number_for_DB2_instance*
Set to the port for the DB2 instance where the database is to be created. If this is a repository installation and the InfoSphere Information Analyzer database is to be created in the local DB2 instance, **ia.db.port** must use the same port as the port used for **db2.port.number**.

An example value is 50000.

ia.db.type=**[db2]**
Set to database type for the InfoSphere Information Analyzer database.

ia.db.username=*Information_Analyzer_db_owner_ID*
Set to the operating system ID to be used as the InfoSphere Information Analyzer database owner ID.

An example value is iauser.

im.shared.directory=*installation_manager_shared_path*
Set to the IBM Installation Manager shared installation path. IBM Installation Manager is used to install WebSphere Application Server. Set this attribute on the services tier only.

An example value is `/opt/IBM/IMShared`.

install.path=**[1|2|3|4|5|6]**
Set to the type of installation:

1 = New Installation

2 = Add Product

3 = Add Tier

4 = Remove Product

5 = Remove Tier

6 = Uninstall

An example value is 1.

is.admin.password=*Information_Server_Admin_password*
Set to the password for the InfoSphere Information Server Administrator ID.

is.admin.user=*Information_Server_Admin_ID*
Set to the user name that you want for the InfoSphere Information Server Administrator ID. This ID will be created in the local repository.

An example value is admin.

is.console.port=*port*
Set to the HTTP or HTTPS port where InfoSphere Information Server is installed.

In HTTP installation, the value is typically 9080. In an HTTPS installation, the value is typically 9443 or 443.

is.ha.selected=**[true|false]**
Set to true if this is an active-passive server cluster configuration installation.

is.https=**[true|false]**
Set to true if this is an HTTPS installation.

is.sas.selected=**[true|false]**
Set to true if you want to enable legacy IBM InfoSphere DataStage SAS jobs. This option can be true only if sas.install.choice is set to a valid value, and DataStage.SAS is true.

is.was.cluster=**[true|false]**
Set to true if you are installing into a WebSphere Application Server cluster configuration.

isf.agent.port=**[true|false]**
Set to the port number for the InfoSphere Information Server Agent.

An example value is 31531.

isf.http.host=*host_name*
Set the host name of the front-end dispatcher for a cluster installation.

This value is the short form of *hostname*, excluding the domain. (For example, only the *HTTP_host_name* portion of *HTTP_host_name.domain*.com is used.)

isf.http.port=*port_number*
Set the HTTP or HTTPS port number of the front-end dispatcher for a cluster installation.

An example value is 80.

isf.server.host=*ISF_server_host_name*
Set the server host name for a cluster environment.

This value is the complete form of *hostname*, including the complete domain.

jobmon.port.1=*port_number_for_first_job_monitor_port*
Set to the port number for the first job monitor port.

An example value is 13400.

jobmon.port.2=*port_number_for_second_job_monitor_port*
Set to the port number for the second job monitor port.

An example value is 13401.

LIC_AGREEMENT=**ACCEPT**
Set to ACCEPT to accept the license agreement.

mq.install.choice=[NONE|SERVER|CLIENT]
Set to SERVER if you want to install the MQ Series Server plug-in.

Set to CLIENT if you want to install the MQ Series Client plug-in.

nls.enabled=[true|false]
Set to true if you want globalization support for InfoSphere DataStage and IBM InfoSphere QualityStage.

The value, false, is supported only for systems that have the **nls.language** parameter set to English and a supported system code page.

nls.language=[*English-UTF8*|*English*|*Chinese-Simplified-UTF8*|*Chinese-Simplified*|*Chinese-Traditional-UTF8*|*Chinese-Traditional*|*French-UTF8*|*French*|*German-UTF8*|*German*|*Italian-UTF8*|*Italian*|*Japanese-EUC*|*Japanese-SJIS*|*Japanese-UTF8*|*Japanese*|*Korean-UTF8*|*Portuguese-Brazilian-UTF8*|*Portuguese-Brazilian*|*Spanish-UTF8*|*Spanish*|*Thai*]
Set to the language to configure for InfoSphere DataStage and InfoSphere QualityStage. The installation computer code page must be compatible and supported by InfoSphere DataStage.

An example value is English.

nls.language.code=[ENU|FRA|PTB|JPN|ITA|KOR|DEU|CHT|CHS|ESP]
Set to the three-letter language code for InfoSphere DataStage and InfoSphere QualityStage support.

Table 77. Supported language codes

Language	Code
English	ENU
French	FRA
German	DEU
Italian	ITA
Spanish	ESP
Portuguese-Brazil	PTB
Chinese-Simplified	CHS
Chinese-Traditional	CHT
Japanese	JPN
Korean	KOR

An example value is ENU.

product[AVI].selected=[true|false]
Set to true if installing IBM InfoSphere QualityStage Address Verification Interface.

product[Blueprint].selected=[true|false]
Set to *true* if installing IBM InfoSphere Blueprint Director.

product[BusinessGlossary].selected=[true|false]
Set to true if installing IBM InfoSphere Information Governance Catalog.

product[BusinessGlossaryAnywhere].selected=[true|false]
Set to true if installing IBM InfoSphere Glossary Anywhere.

product[DataQualityConsole].selected=[true|false]
Set to *true* if installing IBM InfoSphere Data Quality Console.

product[datastage].selected=[true|false]
Set to true if installing InfoSphere DataStage.

product[DPID].selected=[true|false]
Set to true if installing IBM InfoSphere QualityStage Module for DPID.

product[FastTrack].selected=[true|false]
Set to true if installing IBM InfoSphere FastTrack.

product[IMAM].selected=true
Set to *true* to install IBM InfoSphere Metadata Asset Manager.

product[InformationAnalyzer].selected=[true|false]
Set to true if installing InfoSphere Information Analyzer.

product[ISALite].selected=[true|false]
Set to *true* if installing IBM Support Assistant Lite. It is recommended to leave this set to `true`.

product[ISTools].selected=[true|false]
Set to true if installing InfoSphere Information Server Manager.

product[metadata.server].selected=[true]
Set to true for installing IBM InfoSphere Metadata Server.

product[MetadataWorkbench].selected=[true|false]
Set to true if installing IBM InfoSphere Information Governance Catalog.

product[qualitystage].selected=[true|false]
Set to true if installing InfoSphere QualityStage.

product[Recovery].selected=[true|false]
Set to *true* if installing backup and recovery software. It is recommended to leave this set to `true`.

product[SERP].selected=[true|false]
Set to true if installing IBM InfoSphere QualityStage Module for SERP.

product[WISD].selected=[true|false]
Set to true if installing IBM InfoSphere Information Services Director.

sas.install.choice=[SASVersion6|SASVersion8|SASVersion9]
Select the version of InfoSphere DataStage SAS that you are installing. If you are setting is.sas.selected to true to enable legacy InfoSphere DataStage SAS jobs, you must provide a value for sas.install.choice.

An example value is SASVersion6.

srd.db.location=*location_path*
Set to the path for the IBM InfoSphere QualityStage Standardization Rules Designer database. This value applies only if you are installing the repository and services tiers.

srd.db.name=*database_name*
Set to the name of the database that the IBM InfoSphere QualityStage Standardization Rules Designer data store is in. This value applies only if you are installing the repository and services tiers.

srd.db.schemaName=*schema_name*
Set to the name of the IBM InfoSphere QualityStage Standardization Rules Designer schema. This value applies only if you are installing the repository and services tiers.

srd.db.tablespace=*tablespace_name*
Set to the tablespace name of the IBM InfoSphere QualityStage Standardization Rules Designer data store. This value applies only if you are installing the repository and services tiers.

srd.db.user.name=*owner_ID*
Set to the owner ID for the IBM InfoSphere QualityStage Standardization Rules Designer data store. This value applies only if you are installing the repository and services tiers.

srd.db.user.password=*password*
Set to the password of the IBM InfoSphere QualityStage Standardization Rules Designer data store owner ID. This value applies only if you are installing the repository and services tiers.

suite.installLocation=*IS_installation_path*
Set to the absolute directory path where InfoSphere Information Server is to be installed.

An example value is `/opt/IBM/InformationServer`.

suite.tier[app.server.tier].selected=[true|false]
Set to true if installing the Services tier.

suite.tier[client.tier].selected=[true|false]
Set to true if installing the Client tier.

suite.tier[database.tier].selected=[true|false]
Set to true if installing the Repository tier.

suite.tier[engine.tier].selected=[true|false]
Set to true if installing the Engine tier.

target.country.code=[US|FR|BR|JP|IT|KR|DE|TW|CN|ES]
Set the two-letter country or region code for the InfoSphere Information Server installation.

Table 78. Supported country or region codes

Language	Code
Brazil	BR
China	CN
France	FR
Germany	DE
Italy	IT
Japan	JP
Korea	KR
Spain	ES
Taiwan	TW
United States	US

An example value is US.

target.language.code=[en|fr|de|it|es|pt_BR|zh|zh_TW|ja|ko]
Set the language code for the InfoSphere Information Server installation.

Table 79. Supported language codes

Language	Code
English	en
French	fr
German	de
Italian	it
Spanish	es
Portuguese-Brazil	pt_BR
Chinese-Simplified	zh
Chinese-Traditional	zh_TW
Japanese	ja
Korean	ko

An example value is en.

use.bundled.db2=[true|false]
Set to true if installing the DB2 database that is bundled with InfoSphere Information Server.

use.bundled.was=[true|false]
Set to true if installing the IBM WebSphere Application Server Network Deployment that is bundled with InfoSphere Information Server.

was.admin.host.port=*was_console_port_number*
Set to the port for the WebSphere Application Server administrative console.

An example value is 9060.

was.bootstrap.port=*websphere_bootstrap_port_number*
Set to the WebSphere bootstrap port.

An example value is 2809.

was.dcs.unicast.port=*was_ha_manager_port_number*
Set to WebSphere high availability manager communication port.

An example value is 9353.

was.default.host.port=*was_http_transport_port_number*
Set to the port number for the HTTP transport port.

An example value is 9080.

was.install.location=*empty_was_installation_location*
The directory must be empty.

An example value is `/opt/IBM/WebSphere/AppServer`.

was.mutual.auth.listener.port=*was_mutual_bootstrap_port_number*
Set to the port number for the WebSphere mutual bootstrap port.

An example value is `9402`.

was.orb.listener.port=*was_orb_listener_port_number*
Set to the port number for the WebSphere ORB listening port.

An example value is `9100`.

was.profile.dir=*was_profile_directory*
Set the WebSphere profile directory.

An example value is /opt/IBM/WebSphere/AppServer/profiles/InfoSphere.

was.profile.name=*was_profile_name*
Set the WebSphere profile name.

An example value is InfoSphere.

was.sas.ssl.port=*SAS_SSL_port*
Set to the port number for the SAS SSL Server Authentication listener.

An example value is 9401.

was.secure.admin.host.port=*secure_was_port*
Set to the port number for the secure WebSphere Application Server administrative console.

An example value is 9043.

was.secure.default.host.port=*default_secure_HTTP_transport_port*
Set to the port number for the default secure HTTP transport port.

An example value is 9443.

was.secure.sib.mq.port=*secure_service_integration_MQ_port_number*
Set to the port number for the secure service integration MQ interoperability port.

An example value is 5578.

was.secure.sib.port=*secure_service_integration_port_number*
Set to the port number for the secure service integration port.

An example value is 7286.

was.server.auth.listener.port=*CSIV2_server_auth_port_number*
Set to the port number for the CSIV2 server authentication port.

An example value is 9403.

was.server.name=*was_server_name*
Set the WebSphere Application Server server name.

An example value is server1.

was.sib.mq.port=*service_integration_MQ_port_number*
Set to the port number for the service integration MQ interoperability port.

An example value is 5558.

was.sib.port=*service_integration_port*
Set to the port number for the service integration port.

An example value is 7276.

was.sip.port=*SIP_port_number*
Set to the port number for the SIP port.

An example value is 5060.

was.sip.secured.port=*SIP_secured_port_number*
Set to the port number for the secured SIP port.

An example value is 5061.

was.soap.connector.port=*SOAP_connector_port_number*
Set to the port number for the SOAP connector port.

An example value is 8880.

xmeta.db.config.type=[*db2_10_5_nls*]
To set the database configuration type for metadata repository.

xmeta.db.hostname=*metadata_repository_server_host_name*
Set to the host name for the metadata repository database server.

An example value is *myhostname*.*mydomain*.com.

xmeta.db.location=*metadata_repository_db_location*
Set to the path for Linux or UNIX where you want the metadata repository database to be created.

An example value is /opt/IBM/InformationServer/Repos/xmeta.

xmeta.db.name=*metadata_repository_db_name*
Set to the database name for the metadata repository.

An example value is xmeta.

xmeta.db.password=*metadata_repository_db_owner_password*
Set to the password for the metadata repository database owner.

xmeta.db.port=*metadata_repository_db_port*
Set to the port for the DB2 instance where the database is to be created. The port for **xmeta.db.port** must be the same as **db2.port.number** if this is a repository installation and the InfoSphere Information Analyzer database is to be created in the local DB2 instance.

An example value is 50000.

xmeta.db.type=**[db2]**
Set the database type for the metadata repository.

xmeta.db.username=*metadata_repository_db_owner_ID*
Set to the operating system ID to be used as the metadata repository database owner ID.

An example value is xmeta.

xmeta.driver.type=[*DB2_10*]
Set the driver type for the metadata repository.

xmeta.odb.db.hostname=*host_name*
Set to the host name of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.location=*location_path*
Set to the path of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.name=*database_name*
Set to the name of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.port=*port_number*
Set to the port number of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.registration=**[true|false]**
Set to true if the IBM InfoSphere DataStage and QualityStage operations

database repository is to be registered in the metadata repository. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.schemaName=*schema_name*
Set to the schema name of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.tablespace=*tablespace_name*
Set to the tablespace name of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.user.name=*owner_id*
Set to the owner ID of the IBM InfoSphere DataStage and QualityStage operations database. This value applies only if you are installing the repository and services tiers.

xmeta.odb.db.user.password=*password*
Set to the password of the IBM InfoSphere DataStage and QualityStage operations database owner ID. This value applies only if you are installing the repository and services tiers.

xmeta.staging.db.password=*password*
Set to the password of the IBM InfoSphere Metadata Asset Manager staging area owner ID. This value applies only if you are installing the repository and services tiers.

xmeta.staging.db.username=*owner_ID*
Set to the owner ID of the IBM InfoSphere Metadata Asset Manager staging area. This value applies only if you are installing the repository and services tiers.

Starting the installation program with a response file

You can run a silent installation of InfoSphere Information Server by using a response file rather than by using the installation wizard. Follow the progress of the installation by monitoring the installation log file.

Procedure

1. Optional: Modify the EXCEPTION_POLICY environment variable. This variable determines how errors are processed that might occur when running in silent mode.

Operating System	Command
AIX	In the sh or bash shell: `export EXCEPTION_POLICY=n` In the csh or tcsh shell: `setenv EXCEPTION_POLICY=n`
Linux	In the sh or bash shell: `export EXCEPTION_POLICY=n` In the csh or tcsh shell: `setenv EXCEPTION_POLICY=n`
Windows	`set EXCEPTION_POLICY=n`

n can be any of the following values:

1 Prompt user to try again on error (default behavior)

2 Continue on error

Important: Continuing with an installation after an error occurs could result in a nonoperational installation, or might cause existing installed software to stop functioning. Use this option only if directed by IBM Support.

3 Exit on error

2. On Microsoft Windows computers, verify that the user is a member of the local Administrators group. You must be a local administrator or a local user who is directly assigned to the local Administrators group to run the installation program.
   a. Open a command prompt and enter `compmgmt.msc`.
   b. In the navigation pane, click **Local Users and Groups** > **Groups**.
   c. In the right pane, right-click **Administrators** and click **Properties**.
   d. In the Administrators Properties window, locate the user name in the **Members** list. If the user name is in the list, the user is part of the Administrators group.

      If the user name is not in the list, the user is not part of the group. You must use a different user account to run the installation program, or ask a system administrator to add the current account into the local Administrators group.
3. Log in to the target computer.

Operating system	Description
AIX	Log in as the root user.
Linux	Log in as the root user.
Windows	Log in as a local administrator or as a local or domain user who is directly assigned to the local Administrators group.

4. Open a command prompt or shell window.

Operating system	Command
AIX	Open a shell window on the target computer, or use a graphical desktop sharing system. Keep the shell window open during the entire installation process. If you use telnet or ssh, make sure that the session does not shut down or timeout for the duration of the installation.
Linux	Open a shell window on the target computer, or use a graphical desktop sharing system. Keep the shell window open during the entire installation process. If you use telnet or ssh, make sure that the session does not shut down or timeout for the duration of the installation.

Operating system	Command
Windows	Start an administrator session by using one of the following methods: • Run the following command: `runas /user:Administrator cmd.exe` • Open an elevated command window. 1. Click **Start** > **Search**. 2. In the Search Results box, enter `cmd` and press Ctrl+Shift+Enter. 3. Click **Continue** to accept the prompt asking for your permission to continue.

5. In the newly created session, run the following command to clear any Java options that might be set.

Operating system	Command
AIX	In the sh or bash shell: `export IBM_JAVA_OPTIONS=` In the csh or tcsh shell: `unsetenv IBM_JAVA_OPTIONS`
Linux	In the sh or bash shell: `export IBM_JAVA_OPTIONS=` In the csh or tcsh shell: `unsetenv IBM_JAVA_OPTIONS`
Windows	`set IBM_JAVA_OPTIONS=`

6. **AIX Linux**: Configure the following settings for the session to help ensure that the installation is successful:
   - Set the file creation mask to 022:
     `umask 022`
   - Ensure that you have no less than 10240 file descriptor resources allocated.
   - Check that the current setting for the NOFILES kernel parameter is set to a value that is at least equal to the value that you specify for **ulimit**.
7. Extract the installation package. For example:

Operating system	Command
AIX	**`gunzip -c is-suite.*.tar.gz | tar xvf -`**
Linux	**`tar zxvf is-suite.*.tar.gz`**
Windows	The installation image contains some deep directory paths and long file names. Therefore, to avoid running into the Windows 255 character path name limit during later installation steps, extract the contents into a short initial directory path, such as `C:\temp`. Also, instead of using the Windows File Explorer extraction utility, use a utility such as PKZip, WinZip, or 7-zip. For example: **`7z.exe x is-suite.*.zip -oC:\temp`**

The package is named `is-suite.*` for the full product or `is-client.*` for the client only.

8. The is-suite images require the use of a Bundle Entitlement package; the is-client images do not. If you are installing an is-suite image, extract the contents of the entitlement package into the `is-suite` directory that is decompressed in the previous step. Ensure that you have downloaded the correct entitlement package for your operating system platform, licensed products, and suite installation package. A suite entitlement package will only work for a suite installation package. For example:

   ```
   unzip -d is-suite Bundle.suite.*.zip
   ```

   On AIX, if you don't already have the GNU zip utility, it is available as `gzip-`*`program_version.aix_version`*`.ppc.rpm` in the IBM AIX Toolbox.
9. Change to the `is-suite` or `is-client` directory.
10. Run the setup command with the **-reportOnly** parameter to perform the prerequisite check. For example,

    ```
    setup -reportOnly
    ```

    The installation program runs through a prerequisites check phase. If all prerequisite checks pass, installation continues until the installation process finishes. If one or more prerequisite checks fail, the installation displays errors and stops. The prerequisites check report contains messages like the following line:

    ```
    FAIL - The DB2 installation directory C:\IBM\SQLLIB must be empty.
    ```

    If a prerequisites check fails, correct the problem and restart the installation.
11. Run the following command to start the installation program:

Operating system	Command
AIX	`./setup -rsp` *`response_file`* `-verbose`
Linux	`./setup -rsp` *`response_file`* `-verbose`
Windows	`setup.exe -rsp` *`response_file`* `-verbose`

    Where *response_file* is the path and name of your response file.

    The `-verbose` option is optional. The option causes the installation program to display log messages to the console and also write them to the log file. If the `-verbose` option is not specified, log messages are written to the log file only.

    When you start the installation program, the program begins to display messages in the window if the `-verbose` option is specified.

    Be mindful of the following conditions while the installation program is running:

    - Monitor the installation as described in "Monitoring the InfoSphere Information Server installation" on page 247.
    - Leave the terminal window open until the installation is complete.
    - The system might occasionally request you to grant permissions to execute a process, such as `Java.exe` or `db2setup.exe`. When this prompt is displayed, click **Yes**. The system might also display a firewall warning if the firewall is not shutdown or disabled. Click through the warning to allow the installation program to continue.
12. Microsoft Windows: When the installation process completes, restart the computer.

Taking action if the installation program encounters an error

After you start the installation process, the installation program might encounter an error that requires your action to resolve.

Symptoms

The installation program encounters an error and an Installation Exception window is displayed.

User response

Complete these steps in the Installation Exception window:

1. If the error window is still displayed with the option to Abort, Retry, or Continue the installation leave the window up. Likewise, if you are running the installation program as a wizard and you have not closed it, do not close it.
2. Review the last statements in the log file to determine what caused the error. Investigate and fix the error, if possible. Refer to the troubleshooting topics or the support tech notes online and take the specified user action.
3. Retry the installation:
   - If you did not close the installation program, click **Retry failed action**. If the issue is resolved, then it will proceed. Otherwise, if it runs into the same exception, you will see the exception window again.
   - If you did close the installation program, run it again. If you were able to save the results in a response file, run the installation program with the response file so that you do not need to re-enter the parameters. Use the following command if you want to run the installation program with the wizard and response file:

     ```
     setup -rsp response_file_name -uiMode graphical
     ```
4. If the subsequent attempt fails, manually uninstall the product files and try to install again.
5. If failure still occurs, contact IBM support for assistance.

setup command

You use the **setup** command to run the IBM InfoSphere Information Server installation program.

Command syntax

setup [**-help**] [**-uimode** *mode*] [**-rsp** *rspfile*] [**-reportOnly** | **-force**] [**-reportFile** *rptfile*] [**-language** *languagecode*] [**-verbose**] [**-secure** [**-passwordFile** *filename*]]

Parameters

Entering the **setup** command without parameters causes the installation program to start in graphical mode.

-help
Displays this usage information and then exits.

-uiMode *mode*
Specifies the mode to run in. Valid values for *mode* can be `graphical` (default) or `silent`.

-rsp *rspfile*
Specifies the response file to use for this installation. For a silent mode installation, you must specify a response file by using this parameter. For a graphical mode installation, the settings in the response file that you specify display in the wizard and you do not have to enter them again.

If you specify the **-rsp** parameter, the installation program runs in silent mode by default. To run in graphical mode with a response file, you must also specify the **-uiMode graphical** parameter.

-reportOnly
For a silent mode installation only. Runs the prerequisites check but does not run the installation.

Use with the **-reportFile** parameter to store the results of the prerequisites check in a text file.

If this parameter is specified for a software removal, the parameter is ignored.

-force
Use this parameter to cause the installation program to continue even if the prerequisites checks fail.

In graphical mode, if **-force** is specified on the command line and a prerequisites check fails, an **Ignore FAILED items and continue** check box is shown in the window. To continue the installation, click this check box and then click **Next**. If **-force** is not specified, the check box is not shown and the **Next** button is disabled.

In silent mode, if **-force** is specified on the command line and a prerequisites check fails, the installation process continues. If **-force** is not specified, the installation process halts if prerequisites check tests fail.

Important: FAILED and WARNING messages during the prerequisites check indicate items that might cause a failed installation or that might cause malfunctions in the installed software. Use the `-force` option only if directed by IBM Support.

-reportFile *rptfile*
For a silent mode installation only. Specifies a text file for the prerequisites check report.

If a **reportFile** value is not specified, the report is displayed to the console.

-language *languagecode*
Specifies the language that is displayed for the installation.

The following table lists valid language codes.

Table 80. Language codes

Language	Codes
Chinese (Simplified)	`zh_CN`
Chinese (Traditional)	`zh_TW`
English	`en` or `en_US`
French	`fr` or `fr_FR`
German	`de` or `de_DE`
Italian	`it` or `it_IT`
Japanese	`ja` or `ja_JP`
Korean	`ko` or `ko_KR`
Portuguese (Brazilian)	`pt_BR`
Spanish	`es` or `es_ES`

-verbose
Starts the installation session in trace mode. In this mode, log messages are displayed to the console and are also written to the log file. You can use this option in silent mode or graphical mode.

If you do not specify this parameter, log messages are written to the log file only.

-secure
For a graphical mode installation only. You are prompted to create a password when the installation program starts, unless you use the **-passwordFile** option to provide the password. The password is used to authenticate access to the installation program from the browser. The password can be any string of alphanumeric characters. The password can include white space characters.

-passwordFile
Specifies a file that contains a password to be used for this installation session. The file name can be specified as a full or relative path. The file must contain only the password on the first line. The password can be any string of alphanumeric characters. The password can include white space characters. You use this option with the **-secure** option to skip the password prompt.

Examples

The following command causes the installation program to prompt for a password and then start in graphical mode. The installation program uses the settings in the `rsp.txt` response file. It stores log messages in a log file.

```
./setup -secure -rsp rsp.txt -uiMode graphical
```

The following command causes the installation program to start in silent mode. The installation program uses the settings in the `rsp.txt` response file. It stores log messages in a log file. When the prerequisites checks run, results are stored in `prq.txt`. The installation program stops if a prerequisite check test fails.

```
./setup -rsp rsp.txt -reportFile prq.txt
```

The following command causes the installation program to start in silent mode. The installation program uses the settings in the `rsp.txt` response file. It stores log messages in a log file. When the prerequisites checks run, results are stored in `prq.txt`. The installation program stops after the prerequisites check without installing the software.

```
./setup -rsp rsp.txt -reportFile prq.txt -reportOnly
```

The following command causes the installation program to start in silent mode. The installation program uses the settings in the `rsp.txt` response file. It stores log messages in a log file and also displays them to the console. When the prerequisites check runs, results are stored in `prq.txt`.

```
./setup -rsp rsp.txt -reportFile prq.txt -verbose
```

Monitoring log files to follow installation progress

You can follow the progress of an installation by monitoring log files.

To monitor a log file, use the command `tail -f` *`logfilename`* where *logfilename* is the path and name of the log file.

The **tail** command is available on Microsoft Windows systems if you have installed the Microsoft Windows Resource Toolkit, MKS Toolkit, or a third-party

tail utility. The command is also available if you have installed IBM InfoSphere DataStage or server-side software. (The MKS Toolkit is installed with this product module.)

If the **tail** command is not available, open the log file in a text editor that does not lock the file for editing while it is open. To view updates, close the file and then reopen it.

Monitoring the InfoSphere Information Server installation

You can follow the progress of the IBM InfoSphere Information Server suite installation by monitoring the installation log files.

Before you begin

The monitoring procedures assume that you have installed InfoSphere Information Server in the default installation path. If your installation is different from the following path, substitute your path when you complete the procedures.

Operating system	Default installation path
AIX	`/opt/IBM/InformationServer`
Linux	`/opt/IBM/InformationServer`
Windows	`C:\IBM\InformationServer`

Procedure

1. On the computer on which the installation program is running, navigate to the `ibm_is_logs` directory that you specified as the temporary directory. The default temporary directory is:

Operating system	Directory
AIX	`/var/tmp`
Linux	`/tmp`
Windows	`%temp%`

2. Open the log file:

Option	Procedure
Use the `tail` command	`tail -f ISInstall*.log` `ISInstall*.log` is the log file with the most recent date timestamp.
Use a text editor that does not lock the file	Close the file and then reopen it

Important: Do not delete the installation log files or the log file directory during the installation process.

What to do next

When installation is complete, save the log files for troubleshooting and verification purposes.

Monitoring the WebSphere Application Server Network Deployment installation

You can follow the progress of the IBM WebSphere Application Server Network Deployment installation by monitoring the installation log files.

Before you begin

The monitoring procedures assume that you have installed InfoSphere Information Server in the default installation path. If your installation is different from the following path, substitute your path when you complete the procedures.

Operating system	Default installation path
AIX	`/opt/IBM/InformationServer`
Linux	`/opt/IBM/InformationServer`
Windows	`C:\IBM\InformationServer`

About this task

The installation process updates the WebSphere Application Server suite installation log file sporadically. There might be gaps of up to 20 minutes when no updates are displayed.

Procedure

1. On the computer on which the installation program is running, navigate to the directory that contains the log file:

Operating system	Directory
AIX	`cd /opt/IBM/InformationServer/logs`
Linux	`cd /opt/IBM/InformationServer/logs`
Windows	`cd c:\IBM\InformationServer\logs`

2. Open the log file:

Option	Procedure
Use the `tail` command	`tail -f WebSphereInstall.log` `WebSphereInstall.log` is the most recent log file.
Use a text editor that does not lock the file	Close the file and then reopen it

Important: Do not delete the installation log files or the log file directory during the installation process.

Monitoring the WebSphere Application Server Network Deployment application deployment

You can follow the progress of the IBM WebSphere Application Server Network Deployment application deployment by monitoring the installation log files.

Before you begin

The monitoring procedures assume that you have installed InfoSphere Information Server in the default installation path. If your installation is different from the following path, substitute your path when you complete the procedures.

Operating system	Default installation path
AIX	`/opt/IBM/InformationServer`
Linux	`/opt/IBM/InformationServer`
Windows	`C:\IBM\InformationServer`

About this task

Some application deployments can take up to 25 minutes, which can cause the installation program to seem as if it has stopped running.

Procedure

1. On the computer on which the installation program is running, navigate to the directory that contains the log file:

Operating system	Directory
AIX	`cd /usr/IBM/WebSphere/AppServer/profiles/InfoSphere/logs/server1`
Linux	`cd /opt/IBM/WebSphere/AppServer/profiles/InfoSphere/logs/server1`
Windows	`cd c:\IBM\WebSphere\AppServer\Profiles\InfoSphere\logs\server1`

2. Open the log file:

Option	Procedure
Use the `tail` command	`tail -f SystemOut.log` `SystemOut.log` is the most recent log file.
Use a text editor that does not lock the file	Close the file and then reopen it

Note: Do not delete the installation log files or the log file directory during the installation process.

Monitoring the WebSphere Application Server Liberty Profile application deployment

You can follow the progress of the IBM WebSphere Application Server Liberty Profile application deployment by monitoring the installation log files.

Before you begin

The monitoring procedures assume that you have installed InfoSphere Information Server in the default installation path. If your installation is different from the following path, substitute your path when you complete the procedures.

Operating system	Default installation path
AIX	/opt/IBM/InformationServer
Linux	/opt/IBM/InformationServer
Windows	C:\IBM\InformationServer

Procedure

1. On the computer on which the installation program is running, navigate to the directory that contains the log file:

Operating system	Directory
AIX	cd *IS_install_path*/wlp/usr/servers/iis/logs
Linux	cd *IS_install_path*/wlp/usr/servers/iis/logs
Windows	cd *IS_install_path*\wlp\usr\servers\iis\logs

2. Open the log file:

Option	Procedure
Use the `tail` command	`tail -f messages.log` `messages.log` is the most recent log file.
Use a text editor that does not lock the file	Close the file and then reopen it

Note: Do not delete the installation log files or the log file directory during the installation process.

Monitoring InfoSphere DataStage component installations

You can follow the progress of individual IBM InfoSphere DataStage component installations by monitoring the component installation log files.

Before you begin

The monitoring procedures assume that you have installed InfoSphere Information Server in the default installation path. If your installation is different from the following path, substitute your path when you complete the procedures.

Operating system	Default installation path
AIX	/opt/IBM/InformationServer
Linux	/opt/IBM/InformationServer
Windows	C:\IBM\InformationServer

About this task

You can monitor the installation of the following InfoSphere DataStage components:

- InfoSphere DataStage plug-ins
- Parallel engine operators

- Connectors

Each component must be installed into each project on each host system. If you have many projects, the plug-in installation can take considerable time.

Procedure

1. On the computer on which the installation program is running, navigate to the directory that contains the log file:

Operating system	Directory
AIX	`cd /opt/IBM/InformationServer/Server/StagingArea/Log`
Linux	`cd /opt/IBM/InformationServer/Server/StagingArea/Log`
Windows	`cd c:\IBM\InformationServer\Server\StagingArea\Log`

2. Open the log file. Log files are in the following format

Option	Procedure
Use the `tail` command	`tail -f *.log` The component log files are named with timestamps. When tail is run against `*.log`, the file with the most recent timestamp is displayed.
Use a text editor that does not lock the file	Close the file and then reopen it

Important: Do not delete the installation log files or the log file directory during the installation process.

Verifying and testing the installation

After you install IBM InfoSphere Information Server, verify that your installation was successful by running the IBM Support Assistant (ISA) Lite for InfoSphere Information Server tool. You can also verify the installation of individual components, such as the IBM InfoSphere Information Services Director.

Running the ISALite tool to verify the installation and troubleshoot

The ISALite for InfoSphere Information Server tool provides functionality to collect data, files, and artifacts from the local installation of IBM InfoSphere Information Server. It focuses on performing non-destructive tests and passive collection of data to report the system health and verify the correct configuration of the installation.

What the tool offers

The ISALite for InfoSphere Information Server tool helps you troubleshoot InfoSphere Information Server problems. The tool focuses on automatic collection of problem data and provides diagnostic verifications of an InfoSphere Information Server installation. Information pertinent to a problem is collected and analyzed to identify the origin of the problem. The ISALite for InfoSphere Information Server tool reduces the effort and amount of time it takes you to send the appropriate

diagnostic data to IBM Customer Support by automating the gathering and collection process. Using FTP, the tool sends the results to IBM Support and attaches them to an existing PMR (problem management record). Additionally, the tool performs symptom analysis to help streamline the problem determination process.

Where to find the tool and documentation

The ISALite for InfoSphere Information Server tool and documentation are installed with the product in the following directory locations:

Operating system	Directories
Linux AIX	**Tool:** *IS_install_path*/ISALite **Documentation:** *IS_install_path*/ISALite/doc
Windows	**Tool:** *IS_install_path*\ISALite **Documentation:** *IS_install_path*\ISALite\doc

IS_install_path is the directory where you installed InfoSphere Information Server.

For information about how to get the tool and documentation for previous releases, and for any updates that are available, see the ISALite download document on the support site (http://www.ibm.com/support/docview.wss?uid=swg24022700).

To verify the installation

Run the General Diagnostic Health Checker and other product-specific health checkers. The health checkers are available in the main menu, as well as in the **Additional Diagnostic Tools** menu.

To troubleshoot InfoSphere Information Server activities that fail to run or have errors

Run the InfoSphere Information Server Collector tool to collect logs and configuration data, which provides details about why activities fail to run or have errors. The InfoSphere Information Server Collector tool is available in the main menu.

Testing the installation of IBM InfoSphere DataStage and IBM InfoSphere QualityStage

After you install and configure InfoSphere DataStage and InfoSphere QualityStage, test the installation by logging in to the Administrator and Designer clients.

Before you begin

To run these tests, especially to connect to the IBM InfoSphere DataStage and QualityStage Designer client, you must be either an InfoSphere DataStage administrator or a lower-level user with specific project access granted by an administrator (in the Administrator client).

If the services and engine tier do not share a user registry, then you must set up credential mapping. For more information, refer to the *IBM InfoSphere Information Server Administration Guide*.

Procedure

1. Start the IBM InfoSphere DataStage and QualityStage Administrator client:
   a. On the computer where you installed the Administrator client, click **Start** > **All Programs** > **IBM InfoSphere Information Server** > **IBM InfoSphere DataStage and QualityStage Administrator**.
   b. When you are prompted for the server where the services tier is installed, select a URL in the **Host name of the services tier** list. If the URL is not in the list, type the URL into the field. The URL is in the form *server_name*:*port_number*. The server name and port values depend on whether IBM WebSphere Application Server clustering is set up within your services tier configuration.

Table 81. Host and port values for different configurations

WebSphere Application Server configuration	Host value	Port value (HTTPS protocol)	
Clustering is configured	The host name or IP address of the front-end dispatcher (either the Web server or the load balancer). Do not use the host name of a particular cluster member.	HTTPS secure port of the front-end dispatcher (for example, 443). Do not use the port number of a particular cluster member.	
Clustering is not configured	The host name or IP address of the computer where WebSphere Application Server is installed.	HTTPS transport secure port (configured as WC_defaulthost_secure in WebSphere Application Server). Default: 9443	

   c. Select the engine server name in the **Host name of the Information Server engine** list, or type it into the field if it is not in the list.
   d. Use your suite credentials to log in to the Administrator client.
   e. If you get a message that a security certificate from the server is not trusted, accept the certificate and store it in your list of trusted certificates:
      1) To view the security certificate, click **View Certificate**.
      2) Click the **Certification Path** tab, and then select the root certificate.
      3) Click the **General** tab.
      4) Click **Install Certificate**, and then click **Next**.
      5) Select **Place all certificates in the following store**.
      6) Click **Browse**, and then select **Trusted Root Certification Authorities**.
      7) Click **Next**, and then click **Finish** to import the certificate.
2. Start the IBM InfoSphere DataStage and QualityStage Designer client:
   a. On the computer where you installed the Designer client, click **Start** > **All Programs** > **IBM InfoSphere Information Server** > **IBM InfoSphere DataStage and QualityStage Designer**.
   b. Use your suite credentials to log in to the Designer client. If you completed step 1 and installed the certificate in the trusted store on this same computer, the Designer client will reuse that certificate when authenticating.

c. When you are prompted for the server where the services tier is installed, select a URL in the **Hostname of the services tier** list. If the URL is not in the list, type the URL into the field. The URL is in the form *server_name*:*port_number*. Use the same server name and port number that you used in the previous step.

d. When you are prompted for the server and project, select the server and project in the **Project** list. If the values are not in the list, type them into the field. Specify the server and project as follows: *server_name*/*project_name*. For example, if the engine is installed on the xyz server (in a non-clustered application server configuration) and the project is named my-project, type xyz/my-project.

Testing the IBM InfoSphere Information Services Director installation

After you install and configure IBM InfoSphere Information Services Director, test the installation.

Before you begin

You must have suite administrator credentials or have already set up credentials that have user authority.

Procedure

1. Start the IBM InfoSphere DataStage and QualityStage Administrator client:
   a. On the computer where you installed the InfoSphere Information Services Director client, click **Start** > **All Programs** > **IBM InfoSphere Information Server** > **IBM InfoSphere Information Server console**.
   b. Log in to the console with credentials that have the Information Services Director Administrator role or suite administrator credentials. If it asks you to install a certificate, install it in the trusted store.
   c. When you are prompted for the server where the services tier is installed, select a URL in the **Host name of the services tier** list. If the URL is not in the list, type the URL into the field. The URL is in the form *server_name*:*port_number*. The server name and port values depend on whether IBM WebSphere Application Server clustering is set up within your services tier configuration.:

Table 82. Host and port values for different configurations

WebSphere Application Server configuration	Host value	Port value (HTTPS protocol)	
Clustering is configured	The host name or IP address of the front-end dispatcher (either the Web server or the load balancer). Do not use the host name of a particular cluster member.	HTTPS secure port of the front-end dispatcher (for example, 443). Do not use the port number of a particular cluster member.	

Table 82. Host and port values for different configurations (continued)

WebSphere Application Server configuration	Host value	Port value (HTTPS protocol)	
Clustering is not configured	The host name or IP address of the computer where WebSphere Application Server is installed.	HTTPS transport secure port (configured as WC_defaulthost_secure in WebSphere Application Server). Default: 9443	

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.

Chapter 5. Installing components that are not in the suite installation program

Some components have their own installation programs. You install these components after you install IBM InfoSphere Information Server.

Installing the IBM InfoSphere Glossary Anywhere client

To use InfoSphere Glossary Anywhere, you must install and configure the client.

Preparing for enterprise-wide distribution

To prepare for installing the IBM InfoSphere Glossary Anywhere client on many computers in an enterprise, administrators can edit a configuration file so that the same settings are installed for each user. Administrators can also use a customizable silent installation script, and install the client with application management and deployment software.

Before you begin

- You must know the name and port number of the system where InfoSphere Glossary Anywhere is installed. For detailed information about the name and port number requirements, see "Configuration parameters in the config.ini file" on page 258.
- Uninstall previous versions of InfoSphere Glossary Anywhere that are on any target computers. To learn how to retain existing InfoSphere Glossary Anywhere client configuration settings, see "Retaining settings for InfoSphere Glossary Anywhere" on page 270.
- Target computers must have Microsoft Windows and Microsoft .NET Framework 2, Service Pack 1 or later service pack installed. See the InfoSphere Information Server system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html for the specific versions of Windows that are supported.
- The client computers that InfoSphere Glossary Anywhere will be installed on must have a verified certificate.

About this task

You must install and configure the InfoSphere Glossary Anywhere client separately from the installation of IBM InfoSphere Information Governance Catalog and other components of InfoSphere Information Server. The client runs on individual user computers running Microsoft Windows operating systems.

If you are an administrator and must deploy InfoSphere Glossary Anywhere on multiple client computers, you can take advantage of the following features:

- You can edit the `config.ini` configuration file so that InfoSphere Glossary Anywhere is installed with the same settings throughout the enterprise. The `config.ini` file specifies parameters such as connection information to the server that hosts InfoSphere Information Governance Catalog, language settings, text capture method, and which asset types users can search for. Users can later change these settings on their individual installed version of InfoSphere Glossary Anywhere.

© Copyright IBM Corp. 2005, 2014

- You can install InfoSphere Glossary Anywhere with a silent installation script. You can edit the installation script to modify the source and target installation directories and other parameters.
- You can deploy the client to many target computers at one time by using application management and deployment software such as IBM Endpoint Manager.

Procedure

1. Copy the file *file_name*.zip from the root of the installation media to a local computer, where *file_name* includes the version of the client that you want to install. This file is a compressed file that contains all of the required installation files.
2. Uncompress the *file_name*.zip file.
3. Optional: In the root of the installation directory, open config.ini with a text editor. Change the configuration values for your site.

   Note: The values that you configured in the config.ini file take effect the first time you install the InfoSphere Glossary Anywhere client on each user computer. For updates of the same version and release to a new modification level, configuration changes must be made manually within each client, from the **Settings** screen. For updates to a new version or release level, you must uninstall the existing version of InfoSphere Glossary Anywhere and install the new version.

What to do next

Next, install the InfoSphere Glossary Anywhere client in either of the following ways:

- Manually install the client by using a graphical user interface. See Installing the client manually.
- Silently install the client by using a script. See Installing the client silently.
- Deploy the client to many target computers at one time by using application management and deployment software. See "Deploying the client to large groups" on page 269.

Configuration parameters in the config.ini file

The config.ini configuration file contains parameters for the IBM InfoSphere Glossary Anywhere client.

Purpose

The parameters in the config.ini file configure settings such as configuration to the server that hosts IBM InfoSphere Information Governance Catalog and preferences such as text capture method, display language, and what types of items to display. InfoSphere Glossary Anywhere uses the settings in the config.ini file the first time that the client is installed. If you install an update, the settings that exist on the client are retained.

For most parameters, you can set the value for each parameter or accept the default value, if one exists. **Host Configuration** settings must be specified as appropriate to your environment.

Parameters that are not specified in the config.ini file are set to the defaults. If the config.ini file is not found during the installation process, then the defaults are used for all parameters.

Note: To change these parameters after the initial installation of the client, edit the parameters in the **Settings** tabs of each client.

General settings

LogFileDirectory
Writes log messages to the log file in this directory.

Optional: Type the complete directory path on the local disk of the client computer.

The default directory is %APPDATA%\IBM\BusinessGlossaryAnywhere\logs.

CustomizedText
Text you enter here is displayed on the Welcome screen beneath the product title and at the bottom of the InfoSphere Glossary Anywhere.

Optional: Type in any alphanumeric text. The text does not extend beyond the width of the window and is truncated at the end of the last complete word. Do not include graphics such as company logos.

The default text is blank.

Text capture settings

You can specify the method that users can use to capture text. You can choose from mouse and keyboard capture, keyboard capture, and keyboard from selected text capture. All or any combination of these methods can be enabled at the same time.

MouseCapture
If MouseCapture is true, users can click a mouse button in combination with pressing keyboard keys to specify a word or string they want to search for in the business glossary. Mouse capture can be specified to be from the left, middle, or right mouse button with the MouseCaptureModifiers parameter. The keyboard key or keys are specified with the MouseCaptureModifiers parameter. Users select the text with the mouse button, and the string is captured and placed in the search window when the user presses the specified keyboard keys.

The default is true.

MouseCaptureModifiers
When mouse and keyboard capture is enabled (MouseCapture is true), this parameter specifies the keys that can be used for capture of the string that the user selects with the mouse. Possible values are: MOD_CONTROL: Control (Ctrl) key; MOD_SHIFT: Shift key; or MOD_ALT: Alternate (Alt) key or any combination of these keys. For combinations, separate each value with a comma.

The default is the Shift key. The value takes effect only if MouseCapture is true.

MouseCaptureButton
Specifies the left, middle, or right mouse button for mouse capture of a specified string to search for in the business glossary. Possible values are: 1: Left mouse button; 2: Right mouse button; 4: Middle mouse button.

The default is the right mouse button. The value takes effect only if MouseCapture is true.

CaptureFromClipboard
If CaptureFromClipboard is true, users can select a string of text in a document and the captured text is searched for in the business glossary. The text can be single or multiple words. The default is true.

CaptureFromClipboardKeyCode
Specifies an alphanumeric or function key to be used to capture text in combination with the Alternate (Alt), Shift, or Control (Ctrl) keys (or combinations of these keys) when keyboard from selected text capture is specified (CaptureFromClipboard is true).

The default is the alphanumeric character L. The value takes effect only if CaptureFromClipboard is true.

CaptureFromClipboardModifiers
Specifies the keyboard keys to be used when keyboard from selected text capture is specified (CaptureFromClipboard is true). Possible values are: MOD_CONTROL: Control (Ctrl) key; MOD_SHIFT: Shift key; or MOD_ALT: Alternate (Alt) key or any combination of these keys. Separate each key specification with a comma (,) for multiple-key capture.

The default is Ctrl+Shift (Control and Shift keys pressed simultaneously.) The value takes effect only if CaptureFromClipboard is true.

Host configuration settings

HostName
The value of **HostName** depends on whether WebSphere Application Server clustering is set up within your services tier configuration.

Table 83. Value of HostName parameter according to services tier configuration

Services tier configuration	HostName value
Clustered	**HostName** is the name or IP address and the port of the front-end dispatcher (either the Web server or the load balancer). Do not use the host name of a particular cluster member.
Non-clustered	**HostName** is the host name or IP address of the computer where WebSphere Application Server is installed.

Port
Specifies the port number of the Web console of the server or front-end dispatcher that is specified in the **ServerName** parameter.

The value of **Port** depends on whether IBM WebSphere Application Server clustering is set up within your services tier configuration.

Table 84. Value of Port parameter according to services tier configuration

Services tier configuration	Port value
Clustered	**Port** is the port of the front-end dispatcher (either the Web server or the load balancer). Do not use the port of a particular cluster member.
Non-clustered	**Port** is the port that is assigned to the IBM InfoSphere Information Server Web console.

UserName
Specifies the name of a user account that has access to InfoSphere Information Governance Catalog on the server in the **HostName** parameter.

The user must have an Information Governance Catalog Glossary Basic User role or higher.

There is no default value.

Password
Specifies the password of the user account in the **UserName** parameter.

The password is not encrypted and can be seen by anyone who opens the `config.ini` file.

There is no default value.

EnableDesktopSSOConfiguration
When **EnableDesktopSSOConfiguration** is `true`, a check box is displayed on the InfoSphere Glossary Anywhere Settings page that users can select if they want to use Windows desktop single sign-on. InfoSphere Information Server must be configured to support Windows desktop single sign-on for InfoSphere Information Governance Catalog for this parameter to take effect. If **EnableDesktopSSOConfiguration** is `true`, users can override the **EnableDesktopSSO** setting in the `config.ini` file.

The default is `false`.

EnableDesktopSSO
Specifies whether InfoSphere Glossary Anywhere uses Windows desktop credentials when communicating with the server, if **EnableDesktopSSOConfiguration** is `false`. When set to `true`, user credential fields are not displayed in the InfoSphere Glossary Anywhere Settings tab. InfoSphere Information Server must be configured to support Windows desktop single sign-on for InfoSphere Information Governance Catalog for this parameter to take effect.

The default is `false`.

Note: If **EnableDesktopSSOConfiguration** is set to `true`, then users can override the **EnableDesktopSSO** setting.

User preference settings

UILangageCode
Defines the language for the client to use. The default setting is English (United States).

Note: If the `config.ini` file is not found, the default is still English (United States).
If the server is set to a different language than the client, then when the client connects to the server, the client language setting is changed to that of the server. The user is prompted to restart the client for the new setting to take effect.

The following values are supported:

Table 85. Language values

Value	Language
ar	Arabic
en-US	English (United States)
de-DE	German
es-ES	Spanish
fr-FR	French
he-IL	Hebrew

Table 85. Language values (continued)

Value	Language
it-IT	Italian
ja-JP	Japanese
ko-KR	Korean
pt-BR	Portuguese (Brazil)
ru-RU	Russian
zh-CN	Simplified Chinese
zh-TW	Traditional Chinese

ShowWelcomescreen

Specifies whether the InfoSphere Glossary Anywhere welcome screen displays at startup. Possible values are `true` (the screen displays) and `false` (the screen does not display). Users can also select a checkbox to disable or enable this display.The default is `true`.

InitialResultCount

This parameter sets the number of items that are displayed in the InfoSphere Glossary Anywhere window at one time. The available choices are 20, 50, or 100.

The default is `20`.

Search items settings

SelectedSearchClasses

Specifies the types of assets in the catalo that are searched for specified text. Choices are terms, categories, BI (Business Intelligence) reports, BI report fields, BI report models, BI report collections, BI report members, and applications. If multiple assets are specified, separate each with a comma (,).

Table 86. Types of assets that you can search for

Asset type	Value for SelectedSearchClasses
term	eclass:ASCLModel*BusinessTerm
category	eclass:ASCLModel*BusinessCategory
BI report	eclass:ASCLBI*ReportDef
BI report field	eclass:ASCLBI*ReportField
BI model	eclass:ASCLBI*OLAPModel
BI collection	eclass:ASCLBI*OLAPCollection
BI collection member	eclass:ASCLBI*OLAPMember
application	eclass:MwbExtensions*EDS_Application

The default is terms and categories.

Sample config.ini file

This sample file configures InfoSphere Glossary Anywhere to write informational messages to a log file in `C:\IBM\BGA\LOGS`. Errors are also sent to the Microsoft Windows Event Viewer. The company name, Acme Enterprises, is included in the Welcome and Getting Started windows. The welcome screen is displayed on startup, and 20 search results are displayed at one time in the InfoSphere Glossary Anywhere window. InfoSphere Information Governance Catalog is installed on a server called `sample_server`. The display language is Japanese. Users can use either mouse and keyboard capture or keyboard from selected text capture. They cannot use keyboard capture because `CaptureAtCaret` is set to `false`. With mouse and keyboard capture, users press the Control and Shift keys simultaneously to capture

the word selected with the mouse. With keyboard from selected text capture, users press the Control key to capture the selected string. Searches will be made of terms, categories, and BI reports.

```
LogFileDirectory=C:\IBM\BGA\LOGS
HostName=sample_server
Port=9080
UserName=sample_user
Password=sample_password
EnableDesktopSSOConfiguration=false
EnableDesktopSSO=false
CustomizedText=JKK Enterprises
UILanguageCode=ja-JP
MouseCapture=true
MouseCaptureModifiers=MOD_SHIFT, MOD_ALT
MouseCaptureButton=2
CaptureAtCaret=false
CaptureFromClipboard=true
CaptureFromClipboardKeyCode=L
CaptureFromClipboardModifiers=MOD_CONTROL
ShowWelcomeScreen=true
InitialResultCount=20
SelectedSearchClasses=eclass:ASCLModel*BusinessTerm,eclass:
ASCLModel*BusinessCategory,eclass:ASCLBI*ReportDef
```

Installing the client

You can install the InfoSphere Glossary Anywhere client with an installation wizard or with a silent installation script.

Installing the IBM InfoSphere Glossary Anywhere client manually

You can use an installation wizard to install the IBM InfoSphere Glossary Anywhere client on your computer, for either an initial installation or an update.

Before you begin

- You must have a user account on IBM InfoSphere Information Server with an Information Governance Catalog Glossary Basic User role or higher.
- Make sure that InfoSphere Glossary Anywhere is not currently installed on the target computer. If an existing version is installed, remove it before installing the new version. To learn how to retain existing configuration settings, see Retaining settings for InfoSphere Glossary Anywhere.
- The target computer must have a supported version of Microsoft Windows and Microsoft .NET Framework 2, Service Pack 1 or later service pack installed. See the InfoSphere Information Server system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html for the specific versions of Windows that are supported.
- Make sure that you have write access to the target installation directory. To ensure that you have appropriate access, either log in to the target computer as a Microsoft Windows Administrator or make sure that the target installation directory is unlocked.

  Note: If you are not a Microsoft Windows Administrator, and the default target installation directory is locked, you can use the silent installation script.
- You must have Microsoft Windows Internet Explorer, Version 9.0 or later, installed.
- You must have a valid security certificate, either from the target computer or from the server where IBM InfoSphere Information Governance Catalog is installed.

About this task

You can install the InfoSphere Glossary Anywhere client and define user preferences and configuration parameters. When you install the client, the installation process uses the configuration parameters in the provided file `config.ini`. Settings for text capture, host configuration user preferences (including language), and search are included in this file. You can edit this file to change the parameters. After installation, you can change the settings manually from the **Settings** panel of the InfoSphere Glossary Anywhere user interface.

Procedure

1. Copy the InfoSphere Glossary Anywhere installation files to the target computer.
   - If an administrator has prepared a shared version of the installation files, copy these prepared files to your local hard disk.
   - If an administrator has not prepared a shared version of the installation files, copy the installation files to your computer directly from the installation media. To copy the files from the installation media:
     a. Locate the compressed file in the `BusinessGlossaryAnywhereClient` directory in the root of the installation media.
     b. Copy this file to your local hard disk.
     c. Uncompress the file.
2. Optional: Edit the file `config.ini` to change the initial settings.
3. Double-click the file `BGASetup.bat` to start the installation. Follow the prompts to complete the installation. When installation is complete, the welcome screen is displayed and an **Open Book** icon  is displayed in your Windows system tray. Click **OK** to close the welcome screen.
4. To confirm the connection to the glossary of IBM InfoSphere Information Governance Catalog, double-click the **Open Book** icon in the system tray. If the `config.ini` file used by the installation process contains valid server configuration settings, InfoSphere Glossary Anywhere displays its **Getting Started** window. If the `config.ini` file does not contain valid server configuration settings, an error message is displayed. Follow step 5 to correct the server configuration settings manually.
5. To change the server configuration settings, right-click the open book icon and click **Settings** > **Host Configuration**. Configure, or update if necessary, the server name, server port, user name, and password. Contact your IBM InfoSphere Information Governance Catalog administrator or IBM InfoSphere Information Server administrator for this information.
6. The first time that you access InfoSphere Glossary Anywhere, a general error message will be displayed if the security certificate from the server is not trusted. If you receive such a message, then you must install a security certificate on your target computer. One method of doing so is by following these steps with your Microsoft Windows Internet Explorer web browser:
   a. Open Windows Internet Explorer. In the address field, type the connection information to connect to InfoSphere Information Governance Catalog. Use the following format: `https://host_server:port/ibm/iis/igc/`

      where *host_server* is the Host Name and *port* is the Port Number as configured in the Host Configuration tab of the InfoSphere Glossary Anywhere Settings screen.

Note: The following steps describe the procedure to follow with Windows Internet Explorer. Other Internet Explorer versions might require slightly different steps.

b. On the page that warns you about a problem with the security certificate, click **Continue to this website (Not recommended.)**

c. Click **Certificate Error** at the top of the window, then click **View certificates**.

d. Click the Certification Path tab.

e. Click the root path that is displayed at the top of the Certification Path window.

f. Click **View Certificate**.

g. Click **Install Certificate** to launch the Certificate Import Wizard.

h. Follow the Certificate Import Wizard prompts, and click **Finish**. You can store the certificate in any security store.

i. On the Security Warning screen, click **Yes** to install the certificate.

j. Double click the open book icon to reopen InfoSphere Glossary Anywhere and connect to the business glossary.

7. Optional: Restore your original customized files.
   a. Navigate to the root of the installation directory and replace the config.ini file there with the file that you saved earlier.
   b. Navigate to the directory `%APPDATA%\IBM\BusinessGlossaryAnywhere` and replace the user.ext and user.ext.config files there with the files that you saved earlier.
8. Optional: Change other configuration settings.
   a. Right-click the **Open Book** icon and click **Settings** > **Text Capture**. Configure the method that you want to use to capture text.
   b. Click the **User Preferences** tab to define other user parameters. Click **OK** to save the settings and close the configuration window.
   c. Click the **Search Items** tab to define what business glossary content you want to include in searches.

What to do next

If you have no connection errors, you can use InfoSphere Glossary Anywhere. Right-click the **Open Book** icon in the system tray, then select **Open** to display the InfoSphere Glossary Anywhere window at any time. Type a string in the **Search** window to search for terms or categories.

If you have connection errors, check the log file. The log file is in the directory that is in the **`LogFileDirectory`** parameter in the `config.ini` file. Each time that the client is started, a log file is created on the client computer. The name of the log file is `BusinessGlossaryAnywhere`*`yyyymmddhhmmss`*`.log`, where *`yyyymmddhhmmss`* is the year, month, day, hour, minute, and second when the application was started. Verify with the Information Governance Catalog Glossary Administrator that the server name, port number, user name, and password are correct.

You can change the text capture method, user preferences, and configuration values at any time by using the options available from the **Settings** tab.

Installing the InfoSphere Glossary Anywhere client silently

You can use a modifiable installation script to install and configure the IBM InfoSphere Glossary Anywhere client, for either an initial installation or an update.

Before you begin

- You must have a user account on IBM InfoSphere Information Server with an Information Governance Catalog Basic User role or higher.
- Make sure that InfoSphere Glossary Anywhere is not currently installed on the target computer. If an existing version is installed, remove it before installing the new version. To learn how to retain existing configuration settings, see Retaining settings for InfoSphere Glossary Anywhere.
- The target computer must have a supported version of Microsoft Windows and Microsoft .NET Framework 2, Service Pack 1 or later service pack installed. See the InfoSphere Information Server system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html for the specific versions of Windows that are supported.
- Make sure that you have write access to the target installation directory. To ensure this, either log in to the computer that you are installing on as a Microsoft Windows Administrator or make sure that the target installation directory is unlocked.

  Note: If you are not a Microsoft Windows Administrator, and the default target installation directory is locked, you can edit the silent installation script to change the target installation directory to one that is unlocked.
- You must have Microsoft Windows Internet Explorer, Version 9.0 or later, installed.
- You must have a valid security certificate, either from the target computer or from the server where IBM InfoSphere Information Governance Catalog is installed.

About this task

With silent installation, you can customize the installation process by editing the example script file that is provided. For example, you can change the target installation directory or add error codes and messages that you want displayed. The script file also works if you make no changes to it.

You can install the InfoSphere Glossary Anywhere client and define user preferences and configuration parameters. When you install the client, the installation process uses the configuration parameters in the provided file `config.ini`. Settings for text capture, host configuration user preferences (including language), and search are included in this file. You can edit this file to change the parameters. After installation, you can change the settings manually from the **Settings** panel of the InfoSphere Glossary Anywhere user interface.

Procedure

1. Copy the InfoSphere Glossary Anywhere installation files to the target computer.
   - If an administrator has prepared a shared version of the installation files, copy these prepared files to your local hard disk.
   - If an administrator has not prepared a shared version of the installation files, copy the installation files to your computer directly from the installation media. To copy the files from the installation media:
     a. Locate the compressed file in the `BusinessGlossaryAnywhereClient` directory in the root of the installation media.
     b. Copy this file to your local hard disk.
     c. Uncompress the file.

2. Open the file `ExampleSilentInstallBGA.bat` with a text editor, make any required changes, and save the file. By default InfoSphere Glossary Anywhere is installed in the directory `C:\program files\IBM\BusinessGlossaryAnywhere`. You can change this directory, and make any other changes, by editing the file.
3. Open a Windows command prompt window.
4. In the command prompt window, navigate to the directory containing the file.
5. Type the name of the file (as provided, `ExampleSilentInstallBGA.bat`) in the command prompt window.
6. Follow steps 7 through 10 to confirm the connection on each user's computer, and make any required configuration changes or additions.
7. To confirm the connection to the glossary of IBM InfoSphere Information Governance Catalog, double-click the **Open Book** icon in the system tray. If the `config.ini` file used by the installation process contains valid server configuration settings, InfoSphere Glossary Anywhere displays its **Getting Started** window. If the `config.ini` file does not contain valid server configuration settings, an error message is displayed. Follow step 5 to correct the server configuration settings manually.
8. To change the server configuration settings, right-click the open book icon and click **Settings** > **Host Configuration**. Configure, or update if necessary, the server name, server port, user name, and password. Contact your IBM InfoSphere Information Governance Catalog administrator or IBM InfoSphere Information Server administrator for this information.
9. The first time that you access InfoSphere Glossary Anywhere, a general error message will be displayed if the security certificate from the server is not trusted. If you receive such a message, then you must install a security certificate on your target computer. One method of doing so is by following these steps with your Microsoft Windows Internet Explorer web browser:
   a. Open Windows Internet Explorer. In the address field, type the connection information to connect to InfoSphere Information Governance Catalog. Use the following format: `https://`*`host_server`*`:`*`port`*`/ibm/iis/igc/`

      where *`host_server`* is the Host Name and *`port`* is the Port Number as configured in the Host Configuration tab of the InfoSphere Glossary Anywhere Settings screen.

      Note: The following steps describe the procedure to follow with Windows Internet Explorer. Other Internet Explorer versions might require slightly different steps.
   b. On the page that warns you about a problem with the security certificate, click **Continue to this website (Not recommended.)**
   c. Click **Certificate Error** at the top of the window, then click **View certificates**.
   d. Click the Certification Path tab.
   e. Click the root path that is displayed at the top of the Certification Path window.
   f. Click **View Certificate**.
   g. Click **Install Certificate** to launch the Certificate Import Wizard.
   h. Follow the Certificate Import Wizard prompts, and click **Finish**. You can store the certificate in any security store.
   i. On the Security Warning screen, click **Yes** to install the certificate.

j. Double click the open book icon to reopen InfoSphere Glossary Anywhere and connect to the business glossary.

10. Optional: Restore your original customized files.
    a. Navigate to the root of the installation directory and replace the config.ini file there with the file that you saved earlier.
    b. Navigate to the directory `%APPDATA%\IBM\BusinessGlossaryAnywhere` and replace the user.ext and user.ext.config files there with the files that you saved earlier.
11. Optional: Change other configuration settings.
    a. Right-click the **Open Book** icon and click **Settings** > **Text Capture**. Configure the method that you want to use to capture text.
    b. Click the **User Preferences** tab to define other user parameters. Click **OK** to save the settings and close the configuration window.
    c. Click the **Search Items** tab to define what business glossary content you want to include in searches.

What to do next

If users have no connection errors, they can use InfoSphere Glossary Anywhere.

If users have connection errors, check the log file. Each time that the client is started, a log file is created on the client computer. The name of the log file is `BusinessGlossaryAnywhere`*`yyyymmddhhmmss`*`.log`, where *yyyymmddhhmmss* is the year, month, day, hour, minute, and second when the application was started. The log file is in the directory as defined in the **LogFileDirectory** parameter in the `config.ini` file. Verify with the InfoSphere Information Governance Catalog administrator that the server name, port number, user name, and password are correct.

Users can change the text capture method, user preferences, and configuration values at any time by using the options available from the **Settings** tab.

Uninstalling IBM InfoSphere Glossary Anywhere

You can uninstall IBM InfoSphere Glossary Anywhere by using either the Microsoft Windows program removal utility or a silent uninstallation script.

About this task

You can uninstall InfoSphere Glossary Anywhere by using the program removal utility provided with the Windows operating system. Or, you can uninstall by using a customized script. The file `ExampleSilentUninstallBGA.bat` is a sample file that you can add to or change to create a customized script. The file also works if you make no changes or additions.

Procedure

Uninstall by using one of two methods:

- Use the program removal utility that is provided with your version of the Windows operating system.
- Use the customizable silent installation script. To use this script:
  1. Navigate to the directory from which you installed InfoSphere Glossary Anywhere.

2. (Optional) Edit the unininstallation script file `ExampleSilentUnInstallBGA.bat` to customize it. For example, you might want to add to the error codes and messages that are reported. Save the changes.
3. Open a Windows command prompt window, and navigate to the directory containing the file.
4. Type the name of the file (`ExampleSilentUnInstallBGA.bat`) in the command prompt window.
5. After the script has executed, close the command prompt window.

If an error occurs during execution of the script, check the log file `BGAInstall.log`. This file is located in the same directory as the uninstallation script.

Deploying the client to large groups

You can deploy the InfoSphere Glossary Anywhere client to a large set of users by using enterprise application management and deployment software.

About this task

In large enterprises, system administrators might want to deploy the client to large numbers of users. To do so, you can use the `BGAInstaller.msi` (Windows Installer package) or `.bat` files that are provided with InfoSphere Glossary Anywhere. For example, you can deploy these files with enterprise application management and deployment software such as IBM Endpoint Manager. You can also deploy the files with customized scripts.

Follow these high-level steps:

1. Copy and extract the InfoSphere Glossary Anywhere client files.
2. Edit the `config.ini` file, which contains the settings for the client.
3. Deploy the provided `.bat` file or `.msi` file by using your enterprise distribution software or scripts.

The `.msi` or `.bat` file contains instructions to install the client on a single computer. You specify the set of target computers from within the script or deployment tool that you use.

Procedure

1. Locate the client files. The client files are in a compressed file named *`file_name`*`.zip`, where *file_name* includes the client version.
2. Copy and extract the client files. The following files are extracted:
   - `config.ini`
   - `BGASetup.bat`
   - `ExampleSilentInstallBGA.bat`
   - `ExampleSilentUninstallBGA.bat`
   - `BGAInstaller.msi`
3. Edit the `config.ini` file. See “Configuration parameters in the config.ini file” on page 258 for an explanation of the values that you can change in this file.
4. Modify and deploy the `.msi` or `.bat` file.

Option	Steps
To use the .bat file	1. Copy the `ExampleSilentInstallBGA.bat` file, and give this file a new name with the `.bat` extension. 2. Open the new `.bat` file with a text editor, and edit it as needed to customize it to your enterprise. The file contains the following information that you can change or add to: • The location of the `config.ini` file. By default this file is bypassed. You must specify a directory location for `config.ini` for this file to be read during installation. • The installation directory on the client computers. By default, the directory is `%APPDATA%\IBM\BusinessGlossaryAnywhere`. • Responses to result codes, including codes for successful installation, unsuccessful installation, and detection of an existing installation of InfoSphere Glossary Anywhere . The `.bat` file includes detailed instructions that you can refer to for more information. 3. Configure your enterprise deployment software or scripts to run the `.bat` file to install the client on target computers.
To use the .msi file	1. Modify the following values in the `BGAInstaller.msi` file if you do not want to use the defaults: • The location of the `config.ini` file. By default this file is bypassed. You must specify a directory location for `config.ini` for this file to be read during installation. • The installation directory on the client computers. By default, this directory is `%APPDATA%\IBM\BusinessGlossaryAnywhere`. • Responses to result codes, including codes for successful installation, unsuccessful installation, and detection of an existing installation of InfoSphere Glossary Anywhere . 2. Configure your enterprise application management and deployment software to deploy the `BGAInstaller.msi` file.

Retaining settings for InfoSphere Glossary Anywhere

You can save the configuration and customization files from the current installation of the IBM InfoSphere Glossary Anywhere client to replace the default version of those files that are provided with a new installation.

About this task

You might need to uninstall the InfoSphere Glossary Anywhere client when you upgrade to a new version or for other reasons. When you uninstall, the client configuration settings are also removed. You can save these settings so that you do not have to configure them again.

InfoSphere Glossary Anywhere uses a `config.ini` file to store default configuration settings, which is installed when you install the client. This file does not store individual user customization after installation.

Users can configure their InfoSphere Glossary Anywhere client in the Settings window. The settings are saved in the `user.ext` and the `user.ext.config` files, which are stored in the `%APPDATA%\IBM\BusinessGlossaryAnywhere` directory.

Procedure

1. Retain the default configuration settings for your clients.
   a. Locate the `config.ini` file that was used for the previous version. The default location is root of the installation directory.
   b. Save a copy of that file to a temporary directory. Specify a directory that you can access when installing the new version of the client.
2. Type or paste the following path into the address bar of a Microsoft Windows folder window: `%APPDATA%\IBM\BusinessGlossaryAnywhere`. Entering this path brings you to the directory in which the relevant InfoSphere Glossary Anywhere files are located. The actual directory differs depending on which version of Microsoft Windows is installed.
3. Retain the user customized settings.
   a. In the `%APPDATA%\IBM\BusinessGlossaryAnywhere` directory, locate the user.ext and user.ext.config files.
   b. Save a copy of these files to the same temporary directory as in 1

Installing IBM InfoSphere Information Governance Catalog for Eclipse

InfoSphere Information Governance Catalog for Eclipse provides an Eclipse plug-in component so that you can access business glossary content directly from Eclipse-based applications.

Before you begin

Make sure that your system meets the prerequisites described for InfoSphere Information Governance Catalog for Eclipse in IBM InfoSphere Information Server.

The version of InfoSphere Information Governance Catalog for Eclipse must be the same as the version of IBM InfoSphere Information Governance Catalog on the InfoSphere Information Server host that you connect to. If the versions are different, the business glossary content might fail to download.

If you are working with Eclipse applications other than IBM InfoSphere Data Architect or a Rational® Software Architect product, make sure that your Eclipse application is set up to automatically install publicly available Eclipse components. These Eclipse components are provided with InfoSphere Data Architect and Rational Software Architect products. If your installation of Eclipse does not

include them, you might need to update the Eclipse plug-ins that correspond to your version of Eclipse. You must follow different steps depending on the version of Eclipse that is installed.

About this task

The installation file, BGE_ALL_11_3.zip, is on the software media.

You can install InfoSphere Information Governance Catalog for Eclipse from either of the following locations:

Web update site
When you install InfoSphere Information Governance Catalog, a web update site with the InfoSphere Information Governance Catalog for Eclipse installation files is created. The web update site for BGE_ALL_11_3.zip is `/ibm/iis/igc/bge/update`.

Archive update site
Copy the installation file that you need from the software media to a directory on your hard disk. The archive update site is that directory.

Tip: If you are an administrator who provides the client to multiple users, copy the file to a shared directory that can be accessed by all users.

The client consists of the Eclipse features that are described in the following table. Depending on your needs, you can install all of the features or a subset of them.

Feature name	Description	Feature dependencies
IBM InfoSphere Information Governance Catalog Core	Users can view business glossary from within their Eclipse-based application. The feature also includes the InfoSphere Information Governance Catalog for Eclipse Core API.	None
IBM InfoSphere Information Governance Catalog UML Profile	Users can apply the Business Glossary profile to their UML model and to view preexisting term assignments. The feature also includes the InfoSphere Information Governance Catalog for Eclipse UML Profile API.	None
IBM InfoSphere Information Governance Catalog UML Integration	Users can incorporate term information into UML model elements, including assigning terms to model elements.	• InfoSphere Information Governance Catalog Core • InfoSphere Information Governance Catalog UML Profile
IBM InfoSphere Information Governance Catalog Data Model Integration	Users can incorporate term information into logical and physical data model elements, including assigning terms to model elements.	• InfoSphere Information Governance Catalog Core • InfoSphere Information Governance Catalog UML Profile

Feature name	Description	Feature dependencies
IBM InfoSphere Information Governance Catalog Physical Model Integration	Users can incorporate term information and can assign terms to physical data model elements.	InfoSphere Information Governance Catalog Core
IBM InfoSphere Information Governance Catalog Business Process Model Integration	Users can define the process workflow for users, create logic inside a business process, and integrate with other applications and data sources.	InfoSphere Information Governance Catalog Core

Procedure

1. Start the Eclipse application to which you are adding InfoSphere Information Governance Catalog for Eclipse.
2. Configure Eclipse to update publicly available components. Do the following steps, based on the Eclipse version:

   Eclipse 3.4.2

   a. Click **Help** > **Software Updates**.
   b. Click the **Available Software** tab.
   c. Click **Manage Sites**. Make sure that the Ganymede Update Site is in the list of Available Software Sites, and selected. If the Ganymede Update Site is not in the list, click **Add** and enter `http://download.eclipse.org/releases/ganymede` in the **Location** field.
   d. Click **OK**.

   Eclipse 3.5.1

   a. Click **Help** > **Install New Software**.
   b. Click the **Available Software Sites** hyperlink. Make sure that the Galileo Update Site is in the list of Available Software Sites and is enabled. If the Galileo Update Site is not in the list, click **Add** and enter `http://download.eclipse.org/releases/galileo` in the **Location** field.
   c. Click **OK**.
   d. Make sure that **Contact all update sites during install to find required software** is checked.
3. Add the InfoSphere Information Governance Catalog for Eclipse update site.
   a. For Eclipse version 3.4.2, click **Help** > **Software Updates** > **Add Site**.

      Alternatively, for Eclipse version 3.5.1, click **Help** > **Install New Software** > **Add**.
   b. Install either from the web update site or from the archive update site.

      Install from the web update site

      1) In the Add Site window, type the installation URL: `https://`*host_server*`:9443/bge/update`

         host_server
         If your installation has separate front-end dispatcher, use the host name of the front-end dispatcher.

If your installation does not have a separate front-end dispatcher, use the host name of the services tier computer.

2) Click **OK**.

Install from the archive update site

1) In the Add Site window, click **Archive**.
2) Navigate to the installation file in the archive update site. Click **OK**.

   The InfoSphere Information Governance Catalog for Eclipse features are added to the list of available software.

c. Navigate to the site that you added and expand the contents. Select the check boxes for each of the dependent features you want to install. For details about feature dependencies, see the "Feature dependencies" column in the table.

   If you use Eclipse 3.5.1, make sure that **Group items by category** check box is not selected.

d. Click **Install**. Follow the remaining prompts to finish the installation.

4. The first time that you connect to the server, if the certificate from the server is not trusted, a message about HTTPS might be displayed. If you receive such a message, then you must install a security certificate by following these steps:

   a. Open Microsoft Windows Internet Explorer. In the address field, type the connection information to connect to InfoSphere Information Governance Catalog. Use the format: `https://host_server:port/ibm/iis/igc/`

      host_server
      If your installation has a separate front-end dispatcher, use the host name of the front-end dispatcher.

      If your installation does not have a separate front-end dispatcher, use the host name of the services tier computer.

      port If your installation has a separate front-end dispatcher, use the HTTPS port of the front-end dispatcher.

      If your installation does not have a separate front-end dispatcher, use the WebSphere Application Server secure transport port (HTTPS).

      Note: The following steps describe the procedure to follow with Windows Internet Explorer. Internet Explorer versions might require slightly different steps.

   b. On the page that warns you about a problem with the security certificate, click **Continue to this website (Not recommended)**.

   c. Click **Certificate Error** at the top of the window, then click **View certificates**.

   d. Click the **Details** tab.

   e. Click **Copy to File** and then click **Next**.

   f. Select **Base-64 encoded X.509 (.CER)** and then click **Next**.

   g. Store the certificate in any directory and then click **Finish**. The certificate name with its full directory path is needed in a later step.

   h. In a command prompt window, go to the directory of the JRE binary that is associated with the Eclipse application such as InfoSphere Data Architect or

IBM Rational Software Architect. An example of the directory might be `C:\Program Files\IBM\SDP\jdk\jre\bin`.

i. Run the command:

```
keytool -import -v -file "certificate_name" -keystore
..\lib\security\cacerts
```

The file *certificate_name* is the certificate name with its full directory path. You must enclose the file in quotation marks (").

An example of the command might be

```
keytool -import -v -file "C:\SSL\bge_certificate" -keystore
   ..\lib\security\cacerts
```

The keystore utility accesses the certificate and prepares to import it. Enter the following information:

1) At the **Enter keystore password** prompt, type `changeit`.
2) At the **Trust this certificate** prompt, type `Y`.

The certificate is installed and added to the keystore.

What to do next

Validate that InfoSphere Information Governance Catalog for Eclipse is installed, and then configure the client to connect to InfoSphere Information Governance Catalog.

To validate that the InfoSphere Information Governance Catalog for Eclipse is installed:

1. Start your Eclipse application.
2. From the menu bar, select **Window** > **Preferences** to see the list of available Preference page titles.
3. Check that **Business Glossary** is one of the Preference page titles that are listed in the left pane. If it is visible, then InfoSphere Information Governance Catalog for Eclipse was installed.

To connect to InfoSphere Information Governance Catalog:

1. Start your Eclipse application if it is not already started.
2. From the menu bar, select **Window** > **Preferences** to see the list of available Preference page titles.
3. From the list of Preference page titles, click **Business Glossary**.
4. Press **F1** or click the Help icon **?** to display the online help for the connection settings page.
5. Follow the instructions in the online help to fill out the connection settings.

To learn how to use the client, review the documentation by clicking **Help** > **Help Contents** > **IBM InfoSphere Information Governance Catalog for Eclipse** > **Overview**.

Part 3. Configuring InfoSphere Information Server software

To finish setting up IBM InfoSphere Information Server, complete the tasks that correspond to the product modules that you installed or the features that you want to enable. For additional setup steps, see the documentation for your specific product modules.

© Copyright IBM Corp. 2005, 2014

Chapter 6. Configuring IBM InfoSphere Information Server

After you install the product, you must configure connections to your application server and databases, enable security, and configure other products in the suite. Completing these configuration tasks ensures that your installation of InfoSphere Information Server is set up to run optimally with minimal downtime.

Configuring security for IBM InfoSphere Information Server

To configure user access and control, you choose a user registry, create users and groups in the registry, and assign security roles to define your security profile.

About this task

The following list outlines the tasks that you do to set up security for IBM InfoSphere Information Server. For details on completing each task, refer to the *IBM InfoSphere Information Server Administration Guide*.

Procedure

1. Choose a user registry and configure it for InfoSphere Information Server. A user registry contains valid user names and passwords. To log in to InfoSphere Information Server, a user must have a user name and password in the user registry. The installation program configures InfoSphere Information Server to use its internal user registry. As part of security setup, you can configure InfoSphere Information Server to use an external user registry such as the local operating system user registry or a Lightweight Directory Access Protocol (LDAP) user registry.
2. Create users and groups in the user registry. If InfoSphere Information Server is configured to use the internal user registry, create users and groups by using the InfoSphere Information Server console or the InfoSphere Information Server Web console. If InfoSphere Information Server is configured to use an external user registry, use standard operating system utilities or user registry utilities.
3. Assign security roles to users and groups. Security roles determine which suite components each user or group has access to and what level of access that user or group has in the suite component.
4. Configure InfoSphere Information Server engine security. The InfoSphere Information Server engine performs user authentication separately from other InfoSphere Information Server components. Depending upon your user registry configuration, you might have to map credentials between the InfoSphere Information Server user registry and the local operating system user registry on the computer where the engine is installed.
5. Assign project roles to users. Assign project roles if required by the installed suite components.
6. Optional: Configure IBM WebSphere Application Server for non-root administration. By default, WebSphere Application Server runs as root. However, it can also be run by using a non-root user ID. You can configure and set appropriate file system permissions for WebSphere Application Server to "run-as" a non-root user ID.
7. Optional: Configure InfoSphere Information Server agents for non-root administration. By default, the InfoSphere Information Server agents (such as

© Copyright IBM Corp. 2005, 2014

the ASB agent) run as root. However, they can also be run by using a non-root user ID. You can configure and set appropriate file system permissions for the agents to "run-as" a non-root user ID.

8. Optional: Configure the Auditing service. The Auditing service creates an audit trail of security-related events. These events include all security-related settings changes and user login and logout operations. You can configure which audit events to log and how much information to include based on your auditing requirements.

WebSphere Application Server cluster configuration

If you have a clustered application server environment, some post-installation configuration steps are required. For example, you might have already configured a front-end web server if you did not configure it before installation, but you must provide web server information and keys to IBM InfoSphere Information Server.

Configuring a single front-end web server

If you installed IBM InfoSphere Information Server by using a cluster member rather than a front-end web server, you now configure a front-end web server so that it can be mapped to the various InfoSphere Information Server applications.

Before you begin

Review the documentation:

- For Version 8.5.5.1, go to the WebSphere Application Server information center and read “Implementing a web server plug-in.”

About this task

The following steps describe how to install the IBM HTTP Server, Versions 8.5.5.1. Refer to the Apache HTTP Server documentation for information about how to install the Apache HTTP Server.

Procedure

1. Install the web server plug-in module:
   - For Version 8.5.5.1, go to the WebSphere Application Server information center and follow the steps for “Installing web server plug-ins”.

   After you install the Workload Manager plug-in binary files, run the `configure<webserver_name>` script to create a web server definition in WebSphere Application Server.

Operating system	Script
AIX Linux	`configure<webserver_name>.sh`
Windows	`configure<webserver_name>.bat`

   Note: For remote installations, running the script also creates an unmanaged node.
2. Configure the web server to accept encoded slashes. See “Cannot view report result from InfoSphere Information Server console or InfoSphere Information Server Web console” on page 429.
3. Configure SSL for the IBM HTTP Server, and configure SSL between the front-end IBM HTTP Server and the application server:

- For Version 8.5.5.1, go to the WebSphere Application Server information center under “Securing IBM HTTP Server” and complete the following tasks:
  - “Securing with SSL communications”
  - “Configure SSL between the front-end IBM HTTP Server and the deployment manager”

4. If you are using a port other than the default 80 or 443 for your front-end web server, add a virtual host entry with the same port value as your front-end web server.
   - For Version 8.5.5.1, see Configuring virtual hosts in the WebSphere Application Server Version 8.5 information center.
5. If you would like to add more nodes and servers to the existing cluster, see Adding a new cluster member (Vertical Cluster Setup) or Adding a new managed node (Horizontal Cluster Setup).
6. Restart all the cluster members You can stop and start the cluster from the WebSphere Application Server Administrative console and navigating to **Server** > **Clusters** > **WebSphere application server clusters**. Or, on each node, you can run the following commands

```
cd /opt/IBM/WebSphere/AppServer/bin
./stopServer.sh server_name
./startServer.sh server_name
```

Connecting the front-end web server with InfoSphere Information Server

After configuring the front-end web server, you must configure it with IBM InfoSphere Information Server.

Before you begin

Before running the post installation steps for configuring the front-end web server with InfoSphere Information Server you must have already installed and configured the front-end web server with SSL. See “Configuring a single front-end web server” on page 280.

Procedure

Complete these tasks on all InfoSphere Information Server tier computers. Complete them on the services tier first.

1. Edit the following file and change the value of `isf.server.host` and `isf.server.port` to the host name and SSL port of the front-end HTTP server:

   AIX Linux

   IS_install_path/Version.xml

   Windows

   IS_install_path\Version.xml

```
<PersistedVariable encrypted="false" name="isf.server.host"
persistent="true" readonly="true" value="frontend_hostname"/>

<PersistedVariable encrypted="false" name="is.console.port"
persistent="true" readonly="true" value="frontend_sslport"/>
```

2. Edit each of the `registered_servers.xml` files in the *IS_install_path*. You can use the following commands to determine the locations of the files:

   AIX Linux

```
cd /opt/IBM/InformationServer
find . -name registered_servers.xml
```

Example results:

```
./ASBNode/eclipse/plugins/com.ibm.iis.client/registered-servers.xml
./ASBServer/conf/registered-servers.xml
```

Windows

```
cd C:\IBM\InformationServer
dir /b /s /a-d | find "registered_servers.xml"
```

Example results:

```
C:\IBM\InformationServer\ASBNode\eclipse\plugins\com.ibm.iis.client\
registered-servers.xml
C:\IBM\InformationServer\ASBServer\conf\registered-servers.xml
```

Locate the following element in the `registered-servers.xml` file and provide the host name and SSL port number to the front-end HTTP server.

```
<asb-server name="services_computer_hostname"
host-name="frontend_hostname" port="frontend_sslport"
is-primary-server="true" />
```

3. Update the client side InfoSphere Information Server trust stores with the certificates from the front-end HTTP server. (The ⇒ symbol indicates a line continuation.)

AIX Linux

```
cd IS_install_path/ASBServer/bin
./UpdateSignerCerts.sh -url https://frontend_hostname:frontend_sslport ⇒
-user wasadmin_user -password wasadmin_password

cd IS_install_path/ASBNode/bin
./UpdateSignerCerts.sh -url https://frontend_hostname:frontend_sslport ⇒
-user wasadmin_user -password wasadmin_password
```

Windows

```
cd IS_install_path\ASBServer\bin
UpdateSignerCerts.bat -url https://frontend_hostname:frontend_sslport ⇒
-user wasadmin_user -password wasadmin_password

cd IS_install_path/ASBNode/bin
UpdateSignerCerts.bat -url https://frontend_hostname:frontend_sslport ⇒
-user wasadmin_user -password wasadmin_password
```

4. Verify that the certificates are properly accepted and imported to the client side trust store by running the **SessionAdmin** command and checking that a session is found.

AIX Linux

```
cd IS_install_path/ASBNode/bin
./SessionAdmin.sh -url https://frontend_hostname:frontend_sslport ⇒
-user IS_admin_user -password IS_admin_password ⇒
-list-user-sessions

cd IS_install_path/ASBServer/bin
./UpdateSignerCerts.sh -url https://frontend_hostname:frontend_sslport ⇒
-user IS_admin_user -password IS_admin_password ⇒
-list-user-sessions
```

Windows

```
cd IS_install_path/ASBNode/bin
SessionAdmin.bat -url https://frontend_hostname:frontend_sslport ⇒
-user IS_admin_user -password IS_admin_password ⇒
-list-user-sessions

cd IS_install_path\ASBServer\bin
SessionAdmin.bat -url https://frontend_hostname:frontend_sslport ⇒
-user IS_admin_user -password IS_admin_password ⇒
-list-user-sessions
```

5. On Engine tier computers, restart the agent and check the log files to ensure that the agent connected to the services tier.

Changing the front-end dispatcher location

If you change the front-end dispatcher location or configuration, you must complete a few additional steps. This post-configuration task involves updating two files that exist on all tiers.

About this task

During the installation process for a high availability installation, you enter the host name and port number of the front-end dispatcher. This information is used by the installation program to configure IBM InfoSphere Information Server. If you later modify the front-end dispatcher (for example, if you relocate the dispatcher to another machine, replace the Web server with a load balancer, and so on), you must update the following files with the new host name and port number: `registered-servers.xml` and `Version.xml`.

You must update these files on the services tier that hosts the deployment manager and every engine tier computer that is registered to that services tier.

In each of the paths indicated in this task, *is_install_dir* is the InfoSphere Information Server installation directory. The following location is the default for the installation directory:

Operating system	File path
AIX Linux	`/opt/IBM/InformationServer/`
Windows	`C:\IBM\InformationServer\`

Procedure

1. On the services tier computer and every engine tier computer, update the `registered-servers.xml` file with the correct front-end dispatcher host name and port number. The file is in the following directory:

Operating system	Directory
AIX Linux	Engine tier: *is_install_dir*`/ASBNode/eclipse/plugins/com.ibm.iis.client` Services tier: *is_install_dir*`/ASBServer/conf`
Windows	Engine tier: *is_install_dir*`\ASBNode\eclipse\plugins\com.ibm.iis.client` Services tier: *is_install_dir*`\ASBServer\conf`

   In the `registered-servers.xml` file, update the following <asb-server> element, where **dispatcher_hostname** is the new host name of the front-end dispatcher, and **dispatcher_port** is the new HTTP port of the front-end dispatcher.

```
<asb-server name="domain_name"
 host-name="dispatcher_hostname"
 port="dispatcher_port"
 is-primary-server="true"/>
```

Important: On engine tier systems that have clients installed, you might see multiple <asb-server> elements in the registered-servers.xml file. These instances represent the host name and port number of each services tier that has been successfully logged into from this client system. This is the source of the list of services tiers shown on client login pages. Update the host names and port numbers of only the services tiers that are affected by the front-end dispatcher changes.

Note: In the file you are examining, if there is no <asb-server> element with is-primary-server="true" in the registered-servers.xml file, then you are not on the engine tier. Refer to the last step in this topic for information about configuring files on the client tier.

2. On the services tier computer and every engine tier computer, update the Version.xml file with the correct front-end dispatcher host name and port number. The file is in the InfoSphere Information Server installation directory (*is_install_dir*).

   In the Version.xml file, update the following <PersistedVariable> elements, where **dispatcher_hostname** is the new host name of the front-end dispatcher, and **dispatcher_port** is the new HTTP port of the front-end dispatcher:

```
<PersistedVariable encrypted="false" name="was.default.host.port"
  persistent="true" readonly="true" value="dispatcher_port"/>
<PersistedVariable encrypted="false" name="is.console.port"
  persistent="true" readonly="true" value="dispatcher_port"/>
<PersistedVariable encrypted="false" name="isf.server.host"
  persistent="true" readonly="false" value="dispatcher_hostname"/>
...
<PersistedVariable encrypted="false" name="isf.http.host"
  persistent="true" readonly="true" value="dispatcher_hostname"/>
<PersistedVariable encrypted="false" name="isf.http.port"
  persistent="true" readonly="true" value="dispatcher_port"/>
```

   Note: The was.default.host.port parameter appears if you are using the embedded version of IBM WebSphere Application Server. If you preinstalled WebSphere Application Server, this parameter is not listed in your Version.xml file.

3. On the engine tier, recycle the agents and engine processes. Refer to the *IBM InfoSphere Information Server Administration Guide*.
4. Restart the cluster and cluster members. Refer to the *IBM InfoSphere Information Server Administration Guide*.
5. Go to http://dispatcher_hostname:*dispatcher_port* to verify whether you can log in to the InfoSphere Information Server Web console.
6. Optional: When logging into the client with a new host name and port number, the information will be added to the registered-servers.xml file on the client. However, you can optionally update the list of services tier host names and port numbers that appears on the client login page by using one of the following methods:
   - If you are comfortable with editing the XML file across the client tier, you can update the registered-servers.xml file on each client tier computer:

Operating system	File path
AIX Linux	*is_install_dir*/ASBNode/eclipse/plugins/com.ibm.iis.client/registered-servers.xml
Windows	*is_install_dir*\ASBNode\eclipse\plugin\com.ibm.iis.client\registered-servers.xml

   After editing, restart all clients on the client tier.

- Alternatively, if you do not want to edit the XML file, or if you prefer to use a command, you can use the **RegistrationCommand** tool. To remove the host name and port numbers that are no longer in use, run the following commands:

  Note: The **RegistrationCommand** tool is in the following directory: *is_install_dir*\ASBNode\bin

Operating system	Command
AIX Linux	./RegistrationCommand.sh **-remove -hostname** *old_host_name* **-port** *old_port_number* ./RegistrationCommand.sh **-add -name** *new_host_name* **-hostname** *new_host_name* **-port** *new_port_number* ./RegistrationCommand.sh **-list**
Windows	RegistrationCommand.bat **-remove -hostname** *old_host_name* **-port** *old_port_number* RegistrationCommand.bat **-add -name** *new_host_name* **-hostname** *new_host_name* **-port** *new_port_number* RegistrationCommand.bat **-list**

-remove
The **remove** command parameter removes the old host name and old port number so that they no longer appear in the list of services tiers shown during client login.

-add
The **add** command parameter adds the new host name and port number in advance. Alternatively, you can enter the new host name and port number at the next login.

-list
The **list** command parameter lists the servers that you currently have in your registered-servers.xml file, which reflects the list of servers that appear on the client login page.

Registering and creating data stores after installation

Some product modules make use of data stores for their supporting information. The IBM InfoSphere Information Analyzer database can be created before installation or by the installation program. Other data stores--the operations database, standardization rules designer database, and the exceptions database--can be created by the installation program or you can create them afterward. Depending on the installation options you chose, you might need to register them, create them, or both before you can use them.

Procedure

- **The operations database:**

  The operations database is a data store that stores the monitoring data that is displayed by the IBM InfoSphere DataStage and QualityStage Operations Console. This data store is offered for installation with IBM InfoSphere DataStage and QualityStage.

- If you used the installation program to install the repository tier in a DB2 database on the local computer and you did not skip the creation of the operations database, then it would have been registered and created for you during installation.
- If you connected to a remote DB2 database while installing the services tier, then the operations database would have been registered, but it would not have been created. You must create the operations database on the remote computer by using the scripts that the installation program has created for you.

  The scripts are in the following location on the services tier computer:

  AIX Linux

  IS_install_path/SQLScripts/dsodb

  Windows

  IS_install_path\SQLScripts\dsodb

  For instructions on running the scripts, start at step 3 in Creating the operations database schema.
- If, during installation, you chose to skip the creation of this database in a DB2 database on the local computer, then you must register and create the operations database schema. Complete all of the steps in Creating the operations database schema.

- **The standardization rules designer database:**

  The standardization rules designer database is a data store that stores data that is used by the IBM InfoSphere QualityStage Standardization Rules Designer. This data store is offered for installation with IBM InfoSphere QualityStage.
  - If you used the installation program to install the repository tier in a DB2 database on the local computer, then the standardization rules designer database would have been registered and created for you during installation.
  - If you connected to a remote DB2 database while installing the services tier, then the standardization rules designer database would have been registered, but it would not have been created. You must create the standardization rules designer database on the remote computer by using the scripts that the installation program has created for you.

    The scripts are in the following location on the services tier computer:

    AIX Linux

    IS_install_path/SQLScripts/qssrd

    Windows

    IS_install_path\SQLScripts\qssrd

    For instructions on running the scripts, start at step 5 in Example: Manually registering a repository.
- **The exceptions database:**

  The exceptions database is a data store that stores exception descriptors for the exception stages. This data is used by the IBM InfoSphere Data Quality Console. The exceptions database is not offered for installation by the installation program and must be registered and created after installation if you want to use it.

  To register and create the exceptions database, see Creating the exceptions database schema.

Changing configuration of databases and data stores

You might need to change the configuration of databases and data stores (repositories) that are used to run InfoSphere Information Server or its components. This might be required for a variety of reasons; for example, a change in the physical computer that hosts the data store, a change in host name or port number of the computer that hosts the data store, a change to the high-availability configuration, or a password change.

Some data stores might be deployed as separate schemas within the same database. This is always the case with the active metadata repository and the staging area. Other data stores can also be created as schemas in this same database or they can be created in their own separate database. Such a separate database can also hold multiple data stores.

If the configuration is changed for one of the data stores that is in a shared database, then some configuration details for the other data stores that are schemas within that shared database must also be changed. If any of the following changes are made to a database connection, the configuration for the data stores that use the connection information must be updated:

- host name
- port number
- whether regular or high-availability

Passwords are an exception and are unique for each schema in a shared database.

The following table lists the internal data stores and databases used by InfoSphere Information Server and its components, including links to reconfiguration instructions for each.

Table 87. Reconfiguring internal data stores and databases used by InfoSphere Information Server

Repository or database	Description	Default database and schema	Reconfiguration instructions
Active InfoSphere Information Server metadata repository	Stores the metadata about external data sources that are governed, managed, and analyzed by InfoSphere Information Server components. Normally referred to as the metadata repository.	Database: XMETA Schema: XMETA Database must be the same database that is used for the staging metadata repository.	Follow these instructions if you are moving the active metadata repository to a new computer or changing the host name and port: "Changing the metadata repository database host name and port" on page 290 Follow these instructions to configure for high-availability: Metadata repository tier clustering for high availability Follow these instructions to change database passwords: Metadata repository database and staging area repository owner password changing

Table 87. Reconfiguring internal data stores and databases used by InfoSphere Information Server (continued)

Repository or database	Description	Default database and schema	Reconfiguration instructions
InfoSphere Information Server staging area	Stores metadata that is imported from external data sources so that it can be examined before it is moved to the active metadata repository.	Database: XMETA Schema: XMETASR Database must be the same database that is used for the active metadata repository.	Follow these instructions if you are moving the staging area to a new computer or changing the host name and port: Follow these instructions to configure for high-availability: Metadata repository tier clustering for high availability. Follow these instructions to change the database password: Metadata repository database and staging area owner password changing
InfoSphere Information Analyzer analysis database	Stores results of information analyis by InfoSphere Information Analyzer.	Database: IADB Schema: IAUSER Database cannot be the same database that is used for the active metadata repository or staging area.	Follow these instructions if you are moving the analysis database to a new computer: Relocating an IBM InfoSphere Information Analyzer analysis database to a new location Follow these instructions to change the database password: Changing the analysis database owner account credentials
InfoSphere DataStage and QualityStage operations database	Stores monitoring data that is displayed by the InfoSphere DataStage and QualityStage Operations Console.	Database: XMETA Schema: User-defined repository user name, typically DSODB Database can be the same or different as the database that is used for the metadata repository.	Follow these instructions if you are moving the data store to a new computer: Relocating DSODB Follow these instructions if you are changing the host name of the server that hosts the database for the data store: Changing host name configuration for data stores Follow these instructions to reconfigure for high-availability: "Changing configuration of other data stores to high availability" on page 297. Follow these instructions to change the database password: Changing the user password for the operations database

Table 87. Reconfiguring internal data stores and databases used by InfoSphere Information Server (continued)

Repository or database	Description	Default database and schema	Reconfiguration instructions
IBM InfoSphere QualityStage Standardization Rules Designer database	Stores a copy of revisions to InfoSphere QualityStage rule sets that have been made in the IBM InfoSphere QualityStage Standardization Rules Designer .	Database: XMETA Schema: User-defined data store user name, typically SRDUSER Database can be the same or different as the database that is used for the metadata repository.	Follow these instructions if you are moving the data store to a new computer: Relocating QSSRDDB Follow these instructions if you are changing the host name of the server that hosts the database for the data store: Changing host name configuration for data stores Follow these instructions to reconfigure for high-availability: "Changing configuration of other data stores to high availability" on page 297. Follow these instructions to change the password in the connection configuration after you have changed the database password: Example: Changing connection properties
IBM InfoSphere Data Quality Console exceptions database	Stores exceptions that are generated by InfoSphere Information Server products and components.	Database: XMETA Schema: User-defined repository user name, typically ESDB Database can be the same or different as the database that is used for the metadata repository.	Follow these instructions if you are moving the data store to a new computer: Relocating ESDB Follow these instructions if you are changing the host name of the server that hosts the database for the data store: Changing host name configuration for data stores Follow these instructions to reconfigure for high-availability: "Changing configuration of other data stores to high availability" on page 297 To change the password in the connection configuration after you have changed the database password, follow the instructions in the following topic, but substitute the correct parameters for your IBM InfoSphere Data Quality Console exceptions database repository. Changing the user password for the exceptions database

Table 87. Reconfiguring internal data stores and databases used by InfoSphere Information Server (continued)

Repository or database	Description	Default database and schema	Reconfiguration instructions
InfoSphere QualityStage Match Designer database	Stores the results of match test passes by InfoSphere QualityStage Match Designer, a component of InfoSphere QualityStage. This data store is an ODBC data source that is used as a staging area before match designs are checked in to the active metadata repository.	User-defined database name and schema name. No default, but typically MDDB. Database cannot be the same database as that used for the metadata repository.	To change the configuration, you must redefine the DSNs (data source names) of the data source on the client and engine tier computers. "Configuring IBM InfoSphere QualityStage Match Designer" on page 346 Updating Match Designer configuration information

Changing the metadata repository database host name and port

You can change the metadata repository database host name or port by modifying configuration files and then changing host name and port values for IBM WebSphere Application Server data sources to use.

Updating the database properties URL

To update the host and port names for your metadata repository database, you specify the values in the `database.properties` file located on the services tier computer and run the **`iisAdmin -deploy`** command to deploy the changes.

Procedure

1. On the services tier computer, open the `database.properties` file that is used by the application server in a text editor. Ensure you edit the correct file because there are two copies of the file. The file to edit first is located in one of the following directories:
   - **For IBM WebSphere Application Server Network Deployment:**

Operating System	File path
AIX Linux	*IS_install_path*`/ASBServer/apps/lib/iis/classes`
Windows	*IS_install_path*`\ASBServer\apps\lib\iis\classes`

   - **For IBM WebSphere Application Server Liberty Profile:**

Operating System	File path
AIX Linux	*IS_install_path*`/ASBServer/apps/lib/iis/15properties`

Operating System	File path
Windows	*IS_install_path*\ASBServer\apps\lib\iis\15properties

2. In the file, locate the `url` keyword and change the *host* and *port* values to match the host name and port values of the new metadata repository database host.

Table 88. URL keyword and value format for different database systems

Database system	**url keyword and value**
IBM DB2 database system	`url=jdbc\:db2\://host\:port/dbname`
IBM DB2 cluster or HADR database system	Enter the following on one line. The → symbol indicates a line continuation. `url=jdbc\:db2\://host\:port/dbname;→` `clientRerouteAlternateServerName\=alternate_host;→` `clientRerouteAlternatePortNumber\=alternate_host_port;→` `queryCloseImplicit\=2;` **Note:** If you have changed the alternate host and port, update those as well; otherwise, keep the *alternate_host* and *alternate_port* values the same.

3. On the services tier computer, back up the following file and copy the file that you changed in step 2 to this one, overwriting it.

Operating System	File path
AIX Linux	*IS_install_path*/ASBServer/conf/database.properties
Windows	*IS_install_path*\ASBServer\conf\database.properties

4. Propagating the changes can require 1.5 GB or more of temporary disk space to run. Ensure that the following directories have at least 1.5 GB of unused disk space.

Table 89. Temporary system directories

Operating system	Temporary directory
AIX	/tmp
Linux	/tmp
Windows	%USERPROFILE%\AppData\Local\Temp

5. Propagate the changes through the application server by entering one of the following commands.
   - **For WebSphere Application Server Liberty Profile:**

Operating System	File path
AIX Linux	*IS_install_path*/ASBServer/bin/iisAdmin.sh -deploy -srcDir *IS_install_path*/ASBServer/apps/lib/iis/15properties -libName iis -libModule 15properties

Operating System	File path
Windows	*IS_install_path*/ASBServer\bin\ iisAdmin.bat -deploy -srcDir *IS_install_path*\ASBServer\apps\lib\iis\ 15properties -libName iis -libModule 15properties

- **For WebSphere Application Server Network Deployment stand-alone configuration:**

Operating System	File path
AIX Linux	*IS_install_path*/ASBServer/bin/iisAdmin.sh -deploy -srcDir *IS_install_path*/ ASBServer/apps/lib/iis/classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password*
Windows	*IS_install_path*/ASBServer\bin\ iisAdmin.bat -deploy -srcDir *IS_install_path*\ASBServer\apps\lib\iis\ classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password*

- **For WebSphere Application Server Network Deployment clustered configuration:**

Operating System	File path
AIX Linux	*IS_install_path*/ASBServer/bin/ iisAdmin.sh -deploy -srcDir *IS_install_path*/ASBServer/apps/lib/iis/ classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password* -cluster *cluster_name*
Windows	*IS_install_path*/ASBServer\bin\ iisAdmin.bat -deploy -srcDir *IS_install_path*\ASBServer\apps\lib\iis\ classes -libName iis -libModule classes -user *wasadmin* -password *wasadmin_password* -cluster *cluster_name*

6. Restart the application server.

Application server configuration	Procedure
Stand-alone	Restart WebSphere Application Server
Clustered	Restart the WebSphere Application Server cluster members

Configuring automatic client reroute

If you have a DB2 clustered or HADR configuration, configure automatic client reroute for IBM WebSphere Application Server to reconnect to the standby server smoothly. This task is required only if the host name and port change is the result of relocating the metadata repository database to a new computer, or if the alternate server host name or port has changed.

About this task

In an IBM InfoSphere Information Server installation with a clustered DB2 database system setup, automatic client reroute can redirect WebSphere Application Server to connect to an alternate server when the primary server fails. The alternate server is preregistered within automatic client reroute. When the failure occurs, automatic client reroute first tries to reconnect to the primary database. If the connection fails, automatic client reroute establishes a connection with the standby server.

If an update, deletion, or insertion transaction was taking place and the failure occurred before the transaction was committed, the user receives an error message. When the user tries the transaction again, the standby server handles the transaction. If a read transaction was taking place, the read transaction is automatically run on the standby server, and no error message is displayed.

Procedure

1. Log in to the primary node.
2. Run the following command to configure automatic client reroute for the database:

   ```
   db2 update alternate server for database database using
     hostname standby_IP port port
   ```

   The *standby_IP* can be an IPv4 address or an IPv6 address, as shown in the following examples:

   IPv4:

   ```
   db2 update alternate server for database xmeta using
     hostname 192.0.2.7 port 60000
   ```

   IPv6:

   ```
   db2 update alternate server for database xmeta using
     hostname ::ffff:192.0.2.7 port 60000
   ```

Modifying the connection properties for data sources for the metadata repository database

When you change the configuration of the active metadata repository and the staging area, you must modify the data source connection properties such as host name, port, or database, in the IBM WebSphere Application Server configuration so that the metadata repository database can be found when WebSphere Application Server starts.

Procedure

- **For IBM WebSphere Application Server Liberty Profile:**
  1. Edit the following file and change the `iis.db.host` and `iis.db.port` values to the new address and port values.

Operating System	File path
AIX Linux	*IS_install_path*`/wlp/usr/servers/iis/bootstrap.properties`
Windows	*IS_install_path*`\wlp\usr\servers\iis\bootstrap.properties`

When you complete this step, WebSphere Application Server Liberty Profile is able to use the host name and port values in the `bootstrap.properties` file to determine the location of the data sources indicated in the `datasources.xml` file.

2. For clustered metadata repository configurations, if the alternate host name and port have changed, then also update the alternate database host name and port information for the data sources for WebSphere Application Server Liberty Profile.
   - For DB2 clustering, complete the following task: WebSphere Application Server Liberty Profile: Specifying the alternate address and port for a DB2 clustered database.
   - For DB2 HADR, complete the following task: WebSphere Application Server Liberty Profile: Specifying the alternate address and port for a DB2 HADR database.

- **For IBM WebSphere Application Server Network Deployment:**
  1. Log in to the WebSphere Application Server administrative console. Use the WebSphere Application Server administrator user name and password.
  2. Click **Resources** > **JDBC** > **Data sources**.
  3. Complete the following steps for the following data sources for the active metadata repository and staging area that you are changing: **ASB JDBC DataSource**, **ASB JDBC XA DataSource**, **ASB Staging Repository JDBC DS**, and **Jreport JDBC DataSource**:
     a. Under **Common and required data source properties**, change the **hostname**, **port**, or **database** values to the new values for the database.
     b. Click **Apply** in the **Common and required data source properties** and wait a few seconds for the Message to display at the top of the page.
     c. In the Messages area at the top of the page, click **Save** to apply changes to the master configuration. You are returned to the Data sources page.
     d. On the Data sources page, click the data source that you modified, and click **Test Connection**.
  4. For clustered metadata repository configurations, if the alternate host name and port have changed, also update the alternate database host name and port information in WebSphere Application Server.
     - For DB2 clustering, complete the following task: Specifying the alternate address and port for a DB2 clustered database to connect to IBM WebSphere Application Server.
     - For DB2 HADR, complete the following task: Specifying the alternate address and port for a DB2 HADR database to connect to IBM WebSphere Application Server.
  5. Log out of the console.
  6. Restart WebSphere Application Server.

Application server configuration	Procedure
Stand-alone	Restart WebSphere Application Server
Clustered	Restart the WebSphere Application Server cluster members

Updating the installation records with the new metadata repository hostname and port

The `Version.xml` file contains the installation records that are used by the installation program. Updating the values in this file helps to avoid problems with future installations.

About this task

The `Version.xml` file is located in the IBM InfoSphere Information Server installation directory on the services tier computer.

Procedure

1. Use a text editor to open the `Version.xml` file.
2. In the file, locate the following XML element:

```
<PersistedVariable encrypted="false" name="xmeta.db.hostname"
    persistent="true" readonly="true" value="myserver"/>
```

3. Modify the `value` attribute to the host name of your metadata repository database.
4. In the file, locate the following XML element:

```
<PersistedVariable encrypted="false" name="xmeta.db.port"
    persistent="true" readonly="true" value="1433"/>
```

5. Modify the `value` attribute to the port number for your metadata repository database.
6. Save your changes and close the `Version.xml` file.

Registering data stores that are part of a newly moved metadata repository database

If you have moved your metadata repository database to a new host, and the database contains schemas other than the metadata repository and staging area, such as the operations database, then you must reregister these schemas with the new metadata repository on the new host. If the metadata repository database contains only the metadata repository and staging area schemas, you do not need to complete this task.

Before you begin

Before completing this task, you must have already migrated the metadata repository database to a new host and completed the steps to change the host name and port of the metadata repository database.

The paths and command names below are for a Windows file system. For UNIX and Linux, use forward slashes with the command name `ReporitoryAdmin.sh`. The `is_install_path` variable is the installation path for IBM InfoSphere Information Server. The default path is:

Operating system	**Default** *is_install_path* **value**
AIX Linux	`/opt/IBM/InformationServer`
Windows	`C:\IBM\InformationServer`

Procedure

1. On the services tier computer, list the data stores (repositories) that are currently registered.

```
cd is_install_path\ASBServer\bin
RepositoryAdmin.bat -listRepositories

DSODB
QSSRDDB
ESDB
```

2. Create a properties file for each of the data stores (repositories) that exist as schemas within the metadata repository database. If you are not sure which data stores are registered to the metadata repository database, create a properties file for each of the data stores using the **-displayRepository** command as shown below for each data store (repository) listed in step 1 on page 295. Examine the `Database.name` property to determine which data stores are schemas in the metadata repository database. Notice that they are still currently registered with the metadata repository database on the old host. The old host name is the information that you will change and register in the new metadata repository. If you find other data stores that are not in the metadata repository database that also show the old host name in the `DatabaseServer.host` property, these will need to be reregistered as well. For further details on reregistering these data stores, see Relocating repositories examples. For this example, assume that DSODB exists in the metadata repository database, but ESDB and QSSRDB exist in their own databases that have not been moved. You would create a properties file for DSODB.

   ```
   RepositoryAdmin.bat -displayRepository -rn DSODB -res dsodb.properties
   ```

   Example output of a properties file:

   ```
   DatabasePlatform.databaseType=DB2
   DatabasePlatform.version=10.5
   DatabaseServer.host=server1
   DatabaseServer.port=50000
   Database.name=XMETA
   Database.alias=XMETA
   Database.location=C:\
   Repository.name=DSODB
   Repository.description=Operations Database.
   Repository.tool=DataStage
   Repository.context=operations
   Repository.schema=DSODB
   RepositoryConnection.userName=DSODB
   RepositoryConnection.password={iisenc}gwFQseoj24I/SnCFEH+cWg==
   RepositoryConnection.connectionURL=jdbc:db2://server1:50000/XMETA
   ```

3. Unregister the metadata repository database. This action will also unregister each of the data stores in the database (except for the staging area and metadata repository itself, which are not registered in this manner). In this example, DSODB is the only one that would be unregistered with the database.

   ```
   RepositoryAdmin.bat -unregisterDatabase -rn XMETA
   ```

4. Edit each of the properties files, changing the old metadata repository database host name to the new one. In this example, there is only the properties file for DSODB. If the database is not configured for high availability, remove the connection URL property. The tool can derive the URL from other properties in the file.

   ```
   DatabasePlatform.databaseType=DB2
   DatabasePlatform.version=10.5
   DatabaseServer.host=server2
   DatabaseServer.port=50000
   Database.name=XMETA
   Database.alias=XMETA
   Database.location=C:\
   Repository.name=DSODB
   Repository.description=Operations Database.
   Repository.tool=DataStage
   ```

```
Repository.context=operations
Repository.schema=DSODB
RepositoryConnection.userName=DSODB
RepositoryConnection.password={iisenc}gwFQseoj24I/SnCFEH+cWg==
```

5. Register each of the data stores (repositories) that you unregistered, using the respective edited properties files that contain new host information. This will automatically register the new metadata repository database and the database server. In this example, DSODB is the only data store being registered.

   ```
   RepositoryAdmin.bat -registerRepository -pf dsodb.properties
   ```

6. The operations database and exceptions databases require connection files to connect to the engine. If one or both of these database schemas are in your newly moved metadata repository database, log into each engine tier computer and create the connection file for the database schema. See Generating the repository connection file for the operations database and Generating the repository connection file for the exceptions database.

Changing configuration of other data stores to high availability

You can change the InfoSphere DataStage and QualityStage operations database, the IBM InfoSphere QualityStage Standardization Rules Designer database, and the IBM InfoSphere Data Quality Console exceptions database to high-availability configurations.

About this task

You can change data stores that support InfoSphere Information Server components to high-availability configurations. These instructions apply to the following data stores:

- InfoSphere DataStage and QualityStage operations database (DSODB)
- IBM InfoSphere QualityStage Standardization Rules Designer database (QSSRDDB)
- IBM InfoSphere Data Quality Console exceptions database (ESDB)

For information about configuring the active metadata repository and staging area for high-availability, see Metadata repository tier clustering for high availability instead.

The following instructions assume that the data store already exists and is configured for non-high availability operation.

Procedure

1. Use the RepositoryAdmin tool to change the RepositoryConnection.connectionURL property of the data store to a new URL. Use the following commands, depending on the type of database system in use. See RepositoryAdmin tool reference for more information.

   DB2 cluster or DB2 HADR

   (The following is a single command although shown in this document on multiple lines.)

   ```
   RepositoryAdmin.sh -uc -rn reposName -cn reposConnName -cr
   'jdbc:db2://host:port/databaseName;
   clientRerouteAlternateServerName=alternate_host;
   clientRerouteAlternatePortNumber=alternate_port;
   enableSeamlessFailover=1;maxRetriesForClientReroute=2;
   retryIntervalForClientReroute=15;'
   ```

2. For ESDB and DSODB, use the **RegistrationCommand** tool to update the connection configuration file for the data store on each engine tier computer. You must run this command on each engine tier computer that has a *install_dir*/Server/DSODB or *install_dir*/Server/ESDB directory that corresponds to the relevant repository. The connection configuration file is used to connect the engine tier to the data store.

   Note: This step is not required for QSSRDDB. For QSSRDDB, go to step 3. The values of connection properties in the configuration file that are also specified by tokens in the template file are replaced. The result is written to the output file that is specified with the **-results** parameter. The following parameters will be replaced if found in the specified template file: **@DBTYPE@**, **@DRIVER_CLASS@**, **@DRIVER_JARFILES@**, **@CONNECTION_URL@**, **@DATABASE_USER@**, **@PASSWORD@**.

   a. Change to the directory that is associated with the data store. For example, for DSODB, change to the directory *install_dir*/Server/DSODB. For ESDB, change to the directory *install_dir*/Server/ESDB.

   b. Run the **RegistrationCommand** tool with the following syntax, where:
      - *install_dir* is the directory where you installed IBM InfoSphere Information Server.
      - *repos_dir* is the directory for the data store.

      (The ⇒ character indicates a line continuation):

Operating system	Command
AIX Linux	cd *install_dir*/Server/*repos_dir* ../../ASBNode/bin/RegistrationCommand.sh -get_config ⇒ -authfile *credentials_file* -repository ⇒ *repository_name* -configfile ⇒ *config_file_template* -results ⇒ *output_file* For example: cd /opt/IBM/InformationServer/Server/DSODB ../../ASBNode/bin/RegistrationCommand.sh -get_config ⇒ -authfile isadminCredentials.txt -repository_name DSODB ⇒ -configfile DSODBConnect.tmpl -results DSODBConnect.cfg
Windows	cd C:*install_dir*\Server*repos_dir* ..\..\ASBNode\bin\RegistrationCommand.bat -get_config ⇒ -authfile *credentials_file* -repository ⇒ *repository_name* -configfile ⇒ *config_file_template* -results ⇒ *output_file* For example: cd C:\IBM\InformationServer\Server\DSODB ..\..\ASBNode/bin/RegistrationCommand.bat -get_config ⇒ -authfile isadminCredentials.txt -repository_name DSODB ⇒ -configfile DSODBConnect.tmpl -results DSODBConnect.cfg

3. For QSSRDDB, reconfigure the application server data source if you are using WebSphere Application Server Network Deployment. Refer to Modifying the connection properties for data sources for the InfoSphere QualityStage Standardization Rules Designer database for specific instructions.

Modifying the connection properties for data sources for the InfoSphere QualityStage Standardization Rules Designer database

Complete this task only if your installation uses IBM WebSphere Application Server Network Deployment. When you change the configuration of the IBM InfoSphere QualityStage Standardization Rules Designer database, you must change the data source connection properties, such as hostname, port, or database, in the IBM WebSphere Application Server administrative console so that the connections can be found when WebSphere Application Server starts.

Procedure

1. Log in to the WebSphere Application Server administrative console. Use the WebSphere Application Server administrator username and password.
2. Click **Resources** > **JDBC** > **Data sources**.
3. Complete the following steps for both the **QSSRD DataSource** and **QSSRD Global XA DataSource** data sources for the IBM InfoSphere QualityStage Standardization Rules Designer database repository:
   a. Under **Common and required data source properties**, change the **hostname**, **port**, or **database** values to the new values for the database.
   b. Click **Apply** in the **Common and required data source properties** and wait a few seconds for the Message to display at the top of the page.
   c. In the Messages area at the top of the page, click **Save** to apply changes to the master configuration. You are returned to the Data sources page.
   d. On the Data sources page, click the data source that you modified, and click **Test Connection**.
4. Log out of the console.
5. Restart WebSphere Application Server.

Application server configuration	Procedure
Stand-alone	Restart WebSphere Application Server
Clustered	Restart the WebSphere Application Server cluster members

Changing application server host names and ports

IBM InfoSphere Information Server comes with default host name and port settings. If there are conflicts with your existing environment, you can change them.

Changing WebSphere Application Server Network Deployment HTTP server and transport ports

To change the application HTTP server and transport ports, use the WebSphere Application Server administrative console, restart WebSphere Application Server, and then edit other configuration files.

About this task

This task applies to stand-alone (non-cluster) environments only.

Procedure

1. Change the port and virtual port values in the WebSphere Application Server administrative console.
   a. Start the WebSphere Application Server administrative console.
   b. Click **Servers** > **Server Types** > **WebSphere application servers**.
   c. Select the server for which you want to change the port.
   d. Click **Ports**.
   e. Change the value of the port that you want to change by clicking the port name link, changing the port number, and clicking **OK**.
      - To change the value of the HTTP server port or HTTP transport port, change the `WC_defaulthost` value.
      - To change the value of the HTTPS port, change the `WC_defaulthost_secure` value.
   f. Click **Environment** > **Virtual Hosts**.
   g. Select **default_host**.
   h. Select **Host Aliases** under Additional Properties.
   i. Change the old values to the new ones that you specified earlier by clicking the asterisk (*) link next to the port number, changing the port number, and clicking **OK**.
   j. Click **Save** directly to the master configuration.
2. Stop and restart WebSphere Application Server.
3. Update the port value in the `Version.xml` files. These files contain the installation records that are used by the installation program. Keeping these files current avoids problems with future installations. Update the following `Version.xml` files:
   - The file that is on the server that hosts the engine tier
   - The file that is on the server that hosts the services tier

   On each server, the `Version.xml` file is located in the IBM InfoSphere Information Server installation directory.Edit each file as follows:
   a. Use a text editor to open the `Version.xml` file.
   b. In the file, locate the following XML element:
      ```
      <PersistedVariable encrypted="false" name="is.console.port"
          persistent="true" readonly="true" value="port" />
      ```
   c. Change the `value` attribute to the new HTTPS port if HTTPS is configured. If HTTPS is not configured, specify the new HTTP port.
   d. In the file, locate the following XML element:
      ```
      <PersistedVariable encrypted="false" name="isf.http.port"
          persistent="true" readonly="true" value="port" />
      ```
   e. Change the `value` attribute to the new HTTPS port if HTTPS is configured. If HTTPS is not configured, specify the new HTTP port.
   f. If you changed the `WC_defaulthost` value in step 1, specify the same value in the following XML element:
      ```
      <PersistedVariable encrypted="false" name="was.default.host.port"
          persistent="true" readonly="true" value="port" />
      ```
   g. If you changed the `WC_defaulthost_secure` value in step 1, specify the same value in the following XML element:
      ```
      <PersistedVariable encrypted="false" name="was.secure.default.host.port"
          persistent="true" readonly="true" value="port" />
      ```

4. Modify the registered-servers.xml file on each computer that hosts the services tier. This file is located in the following location, where *IS_install_path* is the location of your InfoSphere Information Server installation directory.

Operating system	Installation directory
AIX	*IS_install_path*/ASBServer/conf/ **Note:** The default installation location is /opt/IBM/InformationServer.
Linux	*IS_install_path*/ASBServer/conf/ **Note:** The default installation location is /opt/IBM/InformationServer.
Windows	*IS_install_path*\ASBServer\conf\ **Note:** The default installation location is C:\IBM\InformationServer.

   Edit the file as follows:

   a. Open the file in a text editor.
   b. In the file, locate the following XML element:

```
<asb-server name="host"
            host-name="host"
            port="port"
            is-primary-server="true" />
```

      There might be multiple <asb-server> elements in the registered-servers.xml file. Modify the <asb-server> elements for the servers whose port numbers are changing.
   c. Update the port value.
   d. Save your changes.

5. Modify the registered-servers.xml file on each computer that hosts the engine tier. This file is located in the following location, where *IS_install_path* is the location of your InfoSphere Information Server installation directory.

Operating system	Installation directory
AIX	*IS_install_path*/ASBNode/eclipse/plugins/com.ibm.iis.client
Linux	*IS_install_path*/ASBNode/eclipse/plugins/com.ibm.iis.client
Windows	*IS_install_path*\ASBNode\eclipse\plugins\com.ibm.iis.client

   Edit the file as follows:

   a. Open the file in a text editor.
   b. In the file, locate the following XML element:

```
<asb-server name="host"
            host-name="host"
            port="port"
            is-primary-server="true" />
```

      There might be multiple <asb-server> elements in the registered-servers.xml file. Modify the <asb-server> elements for the servers whose port numbers are changing.
   c. Update the port value.
   d. Save your changes.

6. Optional: Modify the registered-servers.xml file on each computer that hosts the client tier. The list of servers in these files is used to populate the server

lists in the login pages for some client applications. You can update these files if you want to have the updated port values shown on the next login attempt. Alternatively, you can enter the new value in the login page the next time that you log in. However, the old host name and port continue to display in the lists in some client application login pages unless you modify the `registered-servers.xml` file on the computer. Modify the file in one of the following ways:

- Change the file in the same way that you modified the `registered-servers.xml` files on the engine tier computers.

  Note: On client-only computers, the `is-primary-server="true"` attribute does not appear in the `<asb-server>` element.
- Run the **RegistrationCommand** command:

Operating system	Command
AIX	`cd /opt/IBM/InformationServer/ASBNode/bin` `./RegistrationCommand -remove -hostname` *host* `-port` *port*
Linux	`cd /opt/IBM/InformationServer/ASBNode/bin` `./RegistrationCommand -remove -hostname` *host* `-port` *port*
Windows	`cd C:\IBM\InformationServer\ASBNode\bin` `RegistrationCommand.bat -remove -hostname` *host* `-port` *port*

Note: These commands use example installation directories, which might be different from your own setup.

Tip: To use the **RegistrationCommand** command to add the new host name and port number for a new server entry, run the command:

Operating system	Command
AIX	`./RegistrationCommand.sh -add -name` *host* `-hostname` *host* `-port` *port*
Linux	`./RegistrationCommand.sh -add -name` *host* `-hostname` *host* `-port` *port*
Windows	`RegistrationCommand.bat -add -name` *host* `-hostname` *host* `-port` *port*

Changing other WebSphere Application Server Network Deployment ports

After you change ports for IBM WebSphere Application Server Network Deployment in the WebSphere Application Server administrative console, you must restart WebSphere Application Server Network Deployment, and then edit other configuration files.

Procedure

1. Change a port value in the WebSphere Application Server administrative console.
   a. Start the WebSphere Application Server administrative console.
   b. Click **Servers** > **Server Types** > **WebSphere application servers**.

c. Select the server for the port that you want to change.
d. Click **Ports**.
e. Change the value of the port that you want to change.

2. Stop and restart the application server.

Application server environment	Procedure
Stand-alone	Stop and restart WebSphere Application Server
Clustered	Stop and restart the WebSphere Application Server cluster member that you changed the port for

3. If you changed the `SOAP_CONNECTOR_ADDRESS` port value in the WebSphere port settings, you must update the `wsadmin.properties` file.

   Note: For a cluster environment, change `wsadmin.properties` file only if the `SOAP_CONNECTOR_ADDRESS` of the deployment manager is changed. Make sure to update the `wsadmin.properties` files of all the nodes and profiles of the cluster so that they point to the deployment manager. You do not need to update any file if you changed the `SOAP_CONNECTOR_ADDRESS` of a cluster member only.

   a. Use a text editor to open the `wsadmin.properties` file. This file is on the services tier in the following directory, where *<profile>* is the name of the profile where IBM InfoSphere Information Server is installed. In a clustered environment, this is the name of the Custom profile where the cluster member, for which the port has been changed, is deployed.

Operating system	Directory
AIX	`opt/IBM/WebSphere/AppServer/profiles/<profile>/properties/`
Linux	`opt/IBM/WebSphere/AppServer/profiles/<profile>/properties/`
Windows	`C:\IBM\WebSphere\AppServer\profiles\<profile>\properties\`

   Note: These file paths use example installation directories, which might be different from your own setup.

   b. Open the file by using a text editor, and locate the `com.ibm.ws.scripting.port` parameter.
   c. Change the parameter value to match value that you specified for `SOAP_CONNECTOR_ADDRESS`.
   d. Save any changes. Restart WebSphere Application Server after all other changes are complete.

   In a stand-alone environment, the default value of `SOAP_CONNECTOR_ADDRESS` is 8880.

   In a cluster environment, the default value of `SOAP_CONNECTOR_ADDRESS` for the deployment manager is 8879. The default value for `SOAP_CONNECTOR_ADDRESS` for a cluster member is 8880.

   Remember: For a cluster environment, repeat step 3 for all nodes and profiles in the cluster.

Changing WebSphere Application Server Liberty Profile ports

The application server ports are set during installation. If you discover conflicts with your existing environment, you can change the port settings.

About this task

The installation program sets the application server HTTPS port to 9443. If this port is not available, the installation program increments the value by one until an available port is found. It sets the HTTP port in the same way, first attempting with port 9080. If you discover that a chosen port value conflicts with another application in your environment, you can change the port setting.

Procedure

1. Stop the application server:

   AIX Linux

   ```
   <IS_install_path>/ASBServer/bin/MetadataServer.sh stop
   ```

   Windows

   ```
   net stop InfoSvr
   ```

2. Edit the WebSphere Application Server Liberty Profile *IS_install_path*/wlp/usr/servers/iis/bootstrap.properties file. Change the iis.https, iis.http, or both property values to your desired ports. The values must be different from each other. The application server cannot listen for two kinds of protocols on the same port.
3. Start the application server:

   AIX Linux

   ```
   <IS_install_path>/ASBServer/bin/MetadataServer.sh run
   ```

   Windows

   ```
   net start InfoSvr
   ```

4. Update the port values in the Version.xml file on the engine tier computer and services tier computer. The file is located in the IBM InfoSphere Information Server installation directory. This file contains the installation records that are used by the installation program. Keeping this file current avoids problems with future installations. Edit the file with a text editor:
   a. If you want to change the HTTP port, locate the PersistedVariable element for iis.console.http.port. Change the *port* value to the new HTTP port.

      ```
      <PersistedVariable encrypted="false" name="iis.console.http.port"
          persistent="true" readonly="true" value="port" />
      ```

   b. If you want to change the HTTPS port, locate the PersistedVariable elements for iis.console.https.port, is.console.port, and isf.http.port. Change the *port* values to the new HTTPS port.

      ```
      <PersistedVariable encrypted="false" name="iis.console.https.port"
          persistent="true" readonly="true" value="port" />
      <PersistedVariable encrypted="false" name="is.console.port"
          persistent="true" readonly="true" value="port" />

      ...

      <PersistedVariable encrypted="false" name="isf.http.port"
          persistent="true" readonly="true" value="port" />
      ```

5. If you changed the HTTPS port, modify the registered-servers.xml file on each computer that hosts the services tier. This file is in the following location, where *installdir* is the location of your InfoSphere Information Server installation directory.

AIX Linux

<*IS_install_path*>/ASBServer/conf/

Note: The typical installation location is /opt/IBM/InformationServer.

Windows

<*IS_install_path*>\ASBServer\conf\

Note: The typical installation location is C:\IBM\InformationServer.

Edit the file in a text editor. Locate the asb-server element and update the port value to the new HTTPS port.

```
<asb-server name="host"
            host-name="host"
            port="port"
            is-primary-server="true" />
```

There might be multiple asb-server elements in the registered-servers.xml file. Modify the asb-server elements for the servers whose port numbers are changing.

6. If you changed the HTTPS port, modify the registered-servers.xml file on each computer that hosts the engine tier. This file is in the following location, where *installdir* is the location of your InfoSphere Information Server installation directory.

   AIX Linux

   <*IS_install_path*>/ASBNode/eclipse/plugins/com.ibm.iis.client

   Windows

   <*IS_install_path*>\ASBNode\eclipse\plugins\com.ibm.iis.client

   Edit the file in a text editor. Locate the asb-server element and update the port value to the new HTTPS port.

   ```
   <asb-server name="host"
               host-name="host"
               port="port"
               is-primary-server="true" />
   ```

   There might be multiple asb-server elements in the registered-servers.xml file. Modify the asb-server elements for the servers whose port numbers are changing.

7. Optional: Modify the registered-servers.xml file on each computer that hosts the client tier. The list of servers in these files is used to populate the server lists in the login pages for some client applications. You can update these files if you want to have the updated port values shown on the next login attempt. Alternatively, you can enter the new value in the login page the next time that you log in. However, the old host name and port continue to display in the lists in some client application login pages unless you modify the registered-servers.xml file on the computer. Modify the file in one of the following ways:

   - Change the file in the same way that you modified the registered-servers.xml files on the engine tier computers.

     Note: On client-only computers, the is-primary-server="true" attribute does not appear in the <asb-server> element.

   - Run the **RegistrationCommand** command:

AIX Linux

```
cd <IS_install_path>/ASBNode/bin ./RegistrationCommand -remove
-hostname host -port port
```

Windows

```
cd <IS_install_path>\ASBNode\bin RegistrationCommand.bat
-remove -hostname host -port port
```

Note: These commands use example installation directories, which might be different from your own setup.

Tip: To use the **RegistrationCommand** command to add the new host name and port number for a new server entry, run the command:

AIX Linux

```
./RegistrationCommand.sh -add -name host -hostname host -port
port
```

Windows

```
RegistrationCommand.bat -add -name host -hostname host -port
port
```

Setting up firewall exceptions

In most cases, you can avoid specifying firewall exceptions by opening ports when you prepare a Microsoft Windows computer for installation. However, you might prefer to set up explicit firewall exceptions for client-side executable files that require network access.

Procedure

1. Use your firewall software to set up firewall exceptions for the following executable files. The default location of these files is `C:\IBM\InformationServer\Clients\Classic\`.

Table 90. Client-side executable files that require network access

Executable file	Description
`admin.exe`	IBM InfoSphere DataStage and QualityStage Administrator
`director.exe`	IBM InfoSphere DataStage and QualityStage Director
`DSCompilerWizard.exe`	Multi-job compiler used by IBM InfoSphere DataStage and QualityStage Designer
`DSDesign.exe`	IBM InfoSphere DataStage and QualityStage Designer
`dsviewer.exe`	Data set management for IBM InfoSphere DataStage and QualityStage Designer
`C:\IBM\InformationServer\jdk\jre\bin\java.exe` `C:\IBM\InformationServer\jdk\jre\bin\javaw.exe`	Processes used to run Java programs. These processes are needed by all clients.
`C:\IBM\InformationServer\Clients\Classic\ISC\isc.exe`	Program that runs the IBM InfoSphere Information Services Director console and IBM InfoSphere Information Analyzer console

Table 90. Client-side executable files that require network access (continued)

Executable file	Description
C:\IBM\InformationServer\Clients\istools\manager\manager.exe	IBM InfoSphere Information Server Manager
C:\IBM\InformationServer\Clients\istools\cli\istool.bat	InfoSphere Information Server Manager command-line interface

2. Use your firewall software to set up firewall exceptions for the following executable files, which start stand-alone applications. The default location of these files is C:\IBM\InformationServer\Clients\Classic\.

Table 91. Executable files for stand-alone tools that require network access

Executable file	Description
diffapicmdline.exe dsdiffexporter.exe	Compare command-line tool
dsadmin.exe	Command-line administration tool
dscc.exe	Command-line compiler tool
dscmdexport.exe	Command-line export tool
dscmdimport.exe	Command-line import tool
dsexport.exe	Export tool
dsimport.exe	Import tool
dsjob.exe	Command-line job administration tool
NLVAdmin.exe	Support tool for switching to English for a non-English installation
XMDICmd.exe	Metadata importer utility
XMLImporter.exe XMLMDImporter.exe	XML table definition importer

Chapter 7. Configuring a parallel processing environment

Some tasks are required to configure the parallel engine, and some tasks ensure the parallel engine is running smoothly. A well-performing parallel engine requires monitoring and adjustment to various settings, such as network, processor, and kernel parameter settings.

Adding computers to share engine processing

To configure an massively parallel processing (MPP) system that runs the parallel engine on multiple computers, copy the parallel engine components to each computer, then configure the parallel engine on each computer.

Each user who runs parallel jobs must have rsh or ssh access to each computer where the parallel engine components are copied.

Configuring computers to be compute nodes for in an MPP system (AIX Linux)

Follow this procedure to configure your computers to act as a compute nodes in a parallel processing cluster.

About this task

A parallel processing configuration consists of a conductor node that contains the IBM InfoSphere Information Server engine, and compute nodes that contain parallel engines.

Procedure

1. Make the parallel engine components available to each compute node.
2. Configure the parallel engine to run the remote shell (rsh) or the secure shell (ssh) without a password on all compute nodes
3. Set user privileges to run parallel jobs.
4. If you did not globally cross-mount the Projects directory in step 1, set up a method of distributing the transform libraries to the compute nodes.

Making the parallel engine components available to compute nodes (AIX Linux)

When you configure an MPP system, you must make the parallel engine components available to all computers in your MPP system.

About this task

Use one of the following methods to make the parallel engine available to all the compute nodes in the MPP system:

- You can globally cross-mount, typically by using network file system (NFS), a single directory on a single system containing the parallel engine software. This configuration makes software upgrades more convenient than if the parallel engine components are installed on all processing systems. If you are using NFS to globally mount the directory, mount it using the soft link option.

© Copyright IBM Corp. 2005, 2014

- You can use a script to copy the parallel engine components to a directory with the same path name on all processing systems that you designate for processing parallel jobs.

Procedure

1. Cross mount your directories. Ensure that all mounted directories are exactly the same path from the root directory on all nodes.

Option	Description
For non-NFS disks	On the computer where you installed the IBM InfoSphere Information Server engine, change the directory to the parallel engine directory (by default `/opt/IBM/InformationServer/Server/PXEngine`).
For NFS disks	1. On the computer where you installed the IBM InfoSphere Information Server engine, change the directory to the directory where you installed InfoSphere Information Server (by default `/opt/IBM/InformationServer/Server`). 2. Change the directory to your InfoSphere Information Server projects directory (by default `/opt/IBM/InformatonServer/Server/Projects`).

2. For non-NFS disks, enter the following command to copy
   `install_dir/install/copy-orchdist` *`node1 node2 ... nodeN`*

   For *node1* to *nodeN* specify the new node names, separated by spaces.
3. Add the nodes to the configuration files for jobs that will use the nodes. For information about the parallel engine configuration files, see the *IBM InfoSphere DataStage and QualityStage Parallel Job Developer's Guide*.

Configuring the parallel engine to use the secure shell (ssh) program (AIX Linux)

The secure shell (ssh) program provides strong authentication and secure communications between computers on a network.

Before you begin

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, the engine connects from the conductor node to all of the compute nodes, using RSA public key encryption for authentication. To set up ssh, do the following task for each user that runs parallel jobs. It is also necessary to log in to each computer before it can run parallel jobs.

Procedure

1. On the conductor node, run the following command to create public and private encryption keys. Press the Enter key twice when you run the command.
   `ssh-keygen -b 1024 -t rsa -f ~/.ssh/id_rsa`

The keys are saved in the ~/.ssh/id_rsa directory.

2. If the user home directories are not automatically mounted, run the following commands to send the public key to the other computers. Run **scp** for each user on each compute node.

```
cd ~/.ssh
scp id_rsa.pub user@secondary_computer:~/.ssh
```

3. Log in to each of the compute nodes and add the public key to the list of authorized keys. Repeat this process for each user that runs parallel jobs.

```
cd .ssh
cat id_rsa.pub >> authorized_keys2
chmod 640 authorized_keys2
rm -f id_rsa.pub
```

   The file that contains the authorized keys might be named authorized_keys or authorized_keys2, depending on the version of ssh.

4. On the conductor node, run the following command for each user and compute node to test the ssh configuration:

```
conductor_node_name% ssh computer_node_name ls
```

   If the command is successful, you should be able to log in to the remote shell without providing a password.

5. On the conductor node, create the remsh file in the /Server/PXEngine/etc/ directory. In the file, specify 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/.

6. Give the file execute permission:

```
chmod 755 remsh
```

Allowing users to run the remote shell program (rsh) without a password (AIX Linux)

If you do not want to use a secure shell connection, you can choose instead to use rsh for each parallel engine connection. As with ssh, the connections must be able to be established without a password prompt.

About this task

You achieve rsh connections differently according to your type of system. For example, you can edit the /etc/hosts.equiv file or create a .rhosts file for each user. In both cases, add the host name of each parallel processing node to the /etc/hosts.equiv file or the .rhosts file, specifying one host name per line.

The host name that is included in this file must correspond to the setting of the node **fastname** parameter in the parallel configuration file, which is located in the /opt/IBM/InformationServer/Server/PXEngine directory by default.

Procedure

1. Use one of these methods to allow users to run rsh without a password:
   - Edit the /etc/hosts.equiv file. The file must be owned by the root user, must grant read and write access to the root user, and grant no access to any other user (file mode of 600).

- Create a `.rhosts` file for each user. The file must be located in the home directory of each parallel user, be owned by the user, and grant read/write access to the user and no access to any other user (file mode of 600).

2. On each node, use this command to verify that users can use **rsh** without a password:

   ```
   rsh host_name uptime
   ```

   host_name is the name of a compute node. If *host_name* is accessible, this command prints a message displaying the time it has been up.

Specifying the location of the remote shell (rsh) program (AIX Linux):

When the parallel engine tries to establish a remote shell, it searches for the remote shell (rsh) program in a specific set of paths, in a specific order. If the rsh program is not in one of these paths, you must specify the path location explicitly in the `remsh` file.

About this task

The parallel engine searches the paths in the following order:

- *install_dir*`/etc/remsh`, where *install_dir* is the top-level directory where the engine tier is installed.
- `/usr/lpp/ssp/rcmd/bin/rsh`
- `/usr/ucb/rsh`
- `/usr/bin/remsh`
- `/bin/remsh`
- `/usr/bin/rsh`

If the rsh program is not in one of these paths, follow this procedure to specify the path to rsh in the first path that the engine searches: *install_dir*`/etc/remsh`.

Procedure

1. Copy or rename the file *install_dir*`/etc/remsh.example` to *install_dir*`/etc/remsh`.

   This file contains the following shell script, which invokes `/usr/bin/rsh`:

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

2. Edit the last line of the shell script to invoke the specific remote shell command that you are using. For example:

   ```
   #!/bin/sh
   exec /usr/local/remsh "$@"
   ```

3. To ensure that all users can run the command, run the **chmod** command:

   ```
   chmod 755 install_dir/etc/remsh
   ```

4. Test the shell script by running **rsh** on each processing node:

   ```
   rsh node_name uptime
   ```

Setting user privileges to run parallel jobs (AIX Linux)

All users who run parallel jobs must have the necessary privileges on all processing nodes.

Procedure

Ensure that each user who runs parallel jobs has the following privileges:

- Log in access.

- Read, write, and modify access to these directories:
  - /*ishome*/Server/Scratch
  - /*ishome*/Server/Datasets
  - /tmp
- Run permission on local copies of programs and scripts.
- Read access to the path specified in the $APT_ORCHHOME environment variable set in the project, or overridden in the parameters for the job. The installation default is /*ishome*/Server/PXEngine.

Configuring computers to be compute nodes for MPP (Windows)

Follow this procedure to configure Microsoft Windows computers to act as a compute nodes in a parallel processing cluster.

Procedure

1. Run the IBM InfoSphere Information Server installation program on each compute node to install the parallel processing components on the node. See "Installing parallel engine components on compute nodes (Windows)."
2. Set system environment variables on the conductor node and compute node. Also verify that a temporary directory exists on the C: drive. See "Making system environment changes (Windows)" on page 314.
3. Configure a common user account for the compute nodes to use. See "Configuring a common user account for parallel processing compute nodes (Windows)" on page 315.
4. Configure RSH by setting permissions in the hosts.equiv file. IBM InfoSphere DataStage uses remote shell (RSH) to communicate between the different physical computers in a parallel processing cluster. See "Setting remote shell (RSH) permissions (Windows)" on page 316.
5. If your users are Windows Domain users, configure the MKS Toolkit Rexecd and Rshd daemons to handle non-local user names that do not include a domain component. See "Configuring the MKS Rexecd and Rshd daemons (Windows)" on page 316.
6. To run parallel jobs on remote Windows hosts, define a remote shell script on the conductor node. See "Configuring a remote shell script for the parallel engine (Windows)" on page 317.
7. Set up a method of distributing the transform libraries to the compute nodes. See "Distributing transform libraries to compute nodes" on page 321.
8. After you have set up the compute nodes in a Microsoft Windows parallel engine configuration, test the configuration. You can test the configuration by using either an MKS Korn shell or the IBM InfoSphere DataStage client. See "Testing the configuration (Windows)" on page 318.

Installing parallel engine components on compute nodes (Windows)

When you set up a parallel processing configuration with Microsoft Windows nodes, you must run the IBM InfoSphere Information Server installation program to install the parallel engine components on all Windows compute nodes. The installation program installs the parallel engine components, and also installs the MKS Toolkit on each node.

Before you begin

Complete the installation of InfoSphere Information Server, including installation of the engine tier on the conductor node.

Procedure

1. On each computer that is a node, run the InfoSphere Information Server installation program.
2. In the Tier Selection page in the installation program wizard, select **Engine tier** only. You do not need to install other tiers on the compute nodes.
3. After you have installed the engine tier on each computer, add an entry in the parallel engine configuration file for each additional node. For more information about the parallel engine configuration file, see the *IBM InfoSphere DataStage and QualityStage Parallel Job Developer's Guide*.
4. Repeat this procedure for each Windows computer that is to be a compute node in your environment.

Making system environment changes (Windows)

To set up a parallel processing configuration, set system environment variables on the conductor node and compute node. Also verify that a temporary directory exists on the C: drive.

Before you begin

You must have Administrator privileges on the conductor node and compute node computers.

About this task

Do this procedure on the conductor node and on each compute node.

Procedure

1. On the conductor node and on each compute node, use the Microsoft Windows Control Panel to set the following environment variables. Ensure that you specify these environment variables in all capital letters:

Table 92. Environment variables to set on the conductor node and compute nodes

Environment variable	Setting
APT_PM_NOWAITPID	1
APT_DISABLE_TRACKER_STARTUP	1
APT_ORCHHOME	The fully qualified path to the `PXEngine` directory. Use forward slashes instead of backslashes in this specification. For example: `C:/IBM/InformationServer/Server/PXEngine`

2. On the compute nodes, update the system PATH environment variable value. You do not need to update this value on the conductor node.
   a. In the PATH value, add the paths to the parallel engine binaries (executables and DLLs). Add the new paths directly before or after the MKS Toolkit references. Use the standard Windows notation (backslashes) for these specifications. For example:

      ```
      C:\IBM\InformationServer\Server\PXEngine\bin;C:\PROGRA~1\MKSTOO~1\bin;C:\
      PROGRA~1\MKSTOO~1\bin\X11;C:\PROGRA~1\MKSTOO~1\mksnt; ...
      ```

b. Verify that the MKS Toolkit paths in the PATH value precede any DSEngine paths. If DSEngine paths precede MKS Toolkit paths, edit the variable value so that the DSEngine paths follow the MKS Toolkit paths. For example:

```
C:\IBM\InformationServer\Server\PXEngine\bin;C:\PROGRA~1\MKSTOO~1\bin;C:\
PROGRA~1\MKSTOO~1\bin\X11;C:\PROGRA~1\MKSTOO~1\mksnt;C:\IBM\InformationSe
rver\Server\DSEngine\bin ...
```

3. On the conductor node and on each compute node, if you installed InfoSphere Information Server on a drive other than C: (for example, D:), create a tmp directory on the C: drive if the directory does not exist on the C: drive. For example, C:\tmp.

Configuring a common user account for parallel processing compute nodes (Windows)

Do these steps to configure a common user account and privileges for the compute nodes in your Microsoft Windows parallel processing environment.

Before you begin

You must have Administrator privileges on the compute nodes.

About this task

Repeat this procedure for each Windows computer that is to be a compute node in your environment.

Procedure

1. Create a generic user account for the compute nodes to use.
   a. At a command prompt, create a new user account. All computers that participate in the parallel processing configuration must use this account name and password. For example, the following command creates a user named dstage:

      ```
      NET USER dstage dspassword /ADD /COMMENT:"InfoSphere DataStage
      multi-node clustered user account" /FULLNAME:"InfoSphere
      DataStage User" /ACTIVE:YES /EXPIRES:NEVER /PASSWORDCHG:NO
      ```

      dstage is the new user that you are creating

      dspassword is the password that is assigned to the new user
   b. Log off and then log in as the new user.
   c. Edit the User environment variable named HOME. If the entry contains other environment variables, for example %HOMEPATH%, replace the entire contents with the literal value of HOMEPATH. The default location is typically C:\Documents and Settings*newusername* where *newusername* is the name of the new user account. MKS cannot resolve the environment variable when connecting remotely.
   d. Log off as the new user and then log in as an Administrator user.
2. Grant the new user account Write, Read, and Modify permissions to the following directories at minimum. The user must have these access rights to run jobs:
   - *X:*\tmp
   - *X:*\Temp
   - *X:*\Windows\Temp
   - *X:*\IBM\InformationServer\Server\Scratch
   - *X:*\IBM\InformationServer\Server\Datasets

   X: is the drive where your directories are located

If individual users will be allowed to execute jobs under their own credentials, create a common user group and make each user a member of that group. Then, grant the group Write, Read, and Modify permissions to the directories.

Setting remote shell (RSH) permissions (Windows)

IBM InfoSphere DataStage uses remote shell (RSH) to communicate between the different physical computers in a parallel processing cluster. To configure RSH, set permissions in the `hosts.equiv` file.

About this task

The `hosts.equiv` file is used by the remote shell daemon (rshd) to determine which remote hosts and remote users are considered trusted by the local host. Trusted users can run commands on the local host as though they were logged in to the local host directly, without requiring a password.

Procedure

1. At a command prompt, type **ksh**.
2. Connect to the remote engine:
   a. From the ksh shell, run the **ssh** command: `ssh -l newusername hostname`
   b. If prompted for an RSA certificate, accept it.
   c. Enter the password for the new user.
3. Run the **rsetup** command to add the new user and password to the Windows Local Security Authority (LSA) database for use by the rshd service. Enter the same password when prompted.
4. Navigate to the `/etc` subdirectory of the MKS Toolkit installation directory. If you installed the MKS Toolkit in the default location, the directory is `C:\Program Files\MKS Toolkit\etc`.
5. Open the `hosts.equiv` file and add the following line to set permissions for the common user. If this file does not exist, create it.

   ```
   + +
   ```

   This designation sets permissions as wide as possible for the common user.
6. At the command prompt, enter `exit` to end the **ssh** session.
7. Test the configuration by attempting a remote directory listing:

   ```
   rsh -l newusername hostname ls
   ```

   If RSH is configured properly, this command returns a listing of the common user's home directory without prompting you for a password.

What to do next

When you are sure that InfoSphere DataStage is working correctly, modify the `hosts.equiv` file to make the permissions more secure. Alternatively, you can use `.rhosts` files instead of the `hosts.equiv` file. An `.rhosts` file must be set up in each user's home account. The syntax for the `.rhosts` files is the same as the `hosts.equiv` file.

Configuring the MKS Rexecd and Rshd daemons (Windows)

If your users are Microsoft Windows Domain users, configure the MKS Toolkit Rexecd and Rshd daemons to handle non-local user names that do not include a domain component.

About this task

If a user like EXAMPLECO\dstage (where EXAMPLECO is the Windows domain) logs in to the system and tries to run an RSH command without specifying a fully qualified domain name (for example: `rsh -l dstage remotehost ls`), RSH fails trying to obtain the password for the user dstage. This problem occurs even if the MKS **rsetup** utility was run on the remote host as described in "Setting remote shell (RSH) permissions (Windows)" on page 316.

To remedy this situation, use the Configure MKS Toolkit control panel to configure the MKS Rexecd and Rshd daemons to use a domain user if the local user does not exist.

Procedure

1. In the Windows Control Panel, start the **Configure MKS Toolkit** control panel.
2. Click **Rexecd/Rshd**.
3. In the **Rexecd** section, select **Fallback to a domain user if the local user does not exist**.
4. In the **Rshd** section, select **Fallback to a domain user if the local user does not exist**.
5. Click **Manage Services**.
6. In the Service list, click **MKS RSHD Service** and click **Restart**. The Rshd daemon restarts.
7. In the Service list, click **MKS REXECD Service** and click **Restart**. The Rexecd daemon restarts.
8. Click **OK**.

Configuring a remote shell script for the parallel engine (Windows)

To run parallel jobs on remote Microsoft Windows hosts, define a remote shell script on the conductor node.

About this task

A sample script named `remsh.example` is provided in the `PXEngine\etc` directory. Copy this script file and modify it for your configuration.

Also, modify the `remcp.example` file. The **remcp** script controls which remote copy utility (**rcp** or **scp**) is used by the IBM InfoSphere DataStage **distribute-component** script. The **distribute-component** script is also located in the `PXEngine\etc` directory. It is used when the transformer libraries in jobs that contain compiled PX transforms must be copied to the compute nodes.

Procedure

1. Copy the `remsh.example` script file, and rename it `remsh`. Save the file in the `PXEngine\etc` directory.
2. Edit the script file and locate the line that begins with `rsh`:

```
#!/bin/sh
# Example $APT_ORCHHOME/etc/remsh
rsh "$@"
```

3. Change the line as shown in the following example:

```
#!/bin/sh
# Example $APT_ORCHHOME/etc/remsh
$(ROOTDIR)/mksnt/rsh.exe "$@"
```

4. Copy the `remcp.example` script file, and rename it `remcp`. Store the file in the `PXEngine\etc` directory.
5. Edit the script file and locate the line that begins with `rcp`:

```
#!/bin/sh
# Example $APT_ORCHHOME/etc/rercp
rcp "$@"
```

6. Change the line as shown:

```
#!/bin/sh
# Example $APT_ORCHHOME/etc/rercp
$(ROOTDIR)/mksnt/rcp.exe "$@"
```

7. Save your changes. You are now set up to run parallel jobs that contained transform operators in a multi-node clustered SMP environment.

What to do next

Set up a method for distributing compiled transform libraries to all physical compute nodes in the cluster. See "Distributing transform libraries to compute nodes" on page 321.

Testing the configuration (Windows)

After you have set up the compute nodes in a Microsoft Windows parallel engine configuration, test the configuration.

You can test the configuration by using either an MKS Korn shell or the IBM InfoSphere DataStage client.

Testing the parallel engine cluster by using the MKS Korn shell:

To test a Microsoft Windows parallel engine cluster by using the MKS Korn shell, create an environment file for a user and run a simple parallel job.

Before you begin

Set up a method for distributing compiled transform libraries to all physical compute nodes in the cluster. See "Distributing transform libraries to compute nodes" on page 321.

Procedure

1. Log in to the computer as a common user.
2. Create an environment file for the user.
   a. Create an environment file in the user's home directory. For example, for a user named dstage, create the file in the `C:\Documents and Settings\dstage` directory. Give the file a meaningful name. For example, name the file `2node.env`, indicating that this environment file is set up to use a two-node configuration.
   b. In the file, include contents like the following lines:

```
#!/bin/sh

export APT_ORCHHOME="C:/IBM/InformationServer/Server/PXEngine"
export APT_CONFIG_FILE="C:/IBM/InformationServer/Server/
   Configurations/2node.apt"
export PATH="$APT_ORCHHOME/bin;$PATH"
export APT_PM_SHOWRSH=1
export APT_PM_NOWIATPID=1
```

In the `export APT_CONFIG_FILE=` line, specify the parallel engine configuration file to be used. In this example, a configuration file named `2node.apt` is referenced. This file defines two physical nodes: the conductor node and one compute node.

3. Click **Start** > **Run**. Type `ksh`, and click **OK**.
4. Source the environment file into the current environment context:
   ```
   . ./2node.env
   ```
5. Around a simple parallel job to verify the configuration settings. For example:
   ```
   osh "generator -schema record(a:int32) | peek"
   ```

   If the job is successful, 10 lines of integer values are displayed.

What to do next

To save the parallel job command for future testing, create a text file with a `.osh` extension that includes the text `generator -schema record(a:int32) | peek`. To run the job again, use the **`osh -f`** command. For example:

```
osh -f 2node.osh
```

If the number of nodes changes, you can add additional environment files. For example, if the number of physical nodes increases from two to three, create the following files:

- `3node.env`
- `3node.osh`
- `3node.apt`

Source the environment file and run a test job by using the following commands:

```
. ./3node.env
osh -f 3node.osh
```

Testing the parallel engine cluster by using the IBM InfoSphere DataStage Designer client (Windows):

To test a Microsoft Windows parallel engine cluster by using the InfoSphere DataStage Designer client, run a simple job.

Before you begin

Set up a method for distributing compiled transform libraries to all physical compute nodes in the cluster. See “Distributing transform libraries to compute nodes” on page 321.

Procedure

1. Log in to the computer as a common user.
2. Click **Start** > **Programs** > **IBM InfoSphere Information Server** > **IBM InfoSphere DataStage and QualityStage Designer**.
3. Log in to the Designer client.
4. Open an existing job. Alternatively, you can create a job:
   a. Drag a Row Generator stage and a Peek stage from the palette onto the canvas.
   b. Right-click the Row Generator stage and connect it to the Peek stage.
   c. Double-click the Row Generator icon and select the **Output** link tab.

d. In the column grid, define one column with an arbitrary name (for example, `col1`) and assign an SQL data type of `Integer`.
e. Save the job.
5. Click **Edit** > **Job Properties**.
6. Click **Parameters**.
7. Click **Add Environment Variable**.
8. Select the environment variable APT_ORCHHOME from the list. Change the value of the variable to the path to the configuration file to use. For example:

   ```
   C:/IBM/InformationServer/Server/Configurations/2node.apt
   ```

   Make sure to use forward slashes as shown in the example.
9. Click **OK**.
10. Save and compile the job.
11. Click **File** > **Run** and run the job. If the job is successful, 10 lines of integer values are displayed.

Strengthening remote shell permissions (Windows)

After you have verified that the parallel engine configuration operates correctly, strengthen the remote shell (RSH) permissions.

Procedure

Modify the `hosts.equiv` file on the conductor node to narrow the range of trusted machines. Alternatively, set up an `.rhosts` file in the home account directory of each user. Use the following syntax in the file:

Table 93. hosts.equiv and .rhosts syntax

Syntax	Meaning
+	All remote computers are to be trusted by the local computer.
+ *username*	*username* is to be trusted by the local computer. The plus sign and *username* are separated by whitespace in the file.
+ +	All remote users on all remote computers are to be trusted by the local computer. The plus signs are separated by whitespace in the file.

In all cases, the user ID that is specified by the remote user, either explicitly or implicitly, must have a valid account and a valid entry in the Windows LSA database on the local computer.
For example, user dstage issues the following command on remote computer exampleremote:

```
rsh -l dstage examplelocal ls
```

This command lists the contents of the dstage user home account on the local machine examplelocal if both of the following criteria are true:

- Computer exampleremote is resolvable in the `hosts.equiv` file, or hosts.equiv contains two plus signs (+ +).
- A valid entry exists for the local user dstage in the Windows LSA database on the local computer examplelocal.

Distributing transform libraries to compute nodes

If you cannot globally cross-mount the Projects directory (for example, in a Microsoft Windows MPP configuration), you must set up a method for distributing compiled transform libraries to all physical compute nodes in the cluster.

About this task

If the project directory is not shared among all servers in a cluster configuration or MPP configuration, then you must choose one of the following methods for distributing the transform libraries:

Table 94. Transform library distribution methods

Method	Configuration	Advantages	Disadvantages
Manually distribute the transform libraries.	Use operating system commands to copy the transform libraries to all compute nodes.	This is the simplest method.	This method is time-consuming and error-prone. Also, it requires a time window when jobs are not running on the system.
Copy the transform libraries automatically at runtime.	Set the APT_COPY_TRANSFORM_OPERATOR environment variable to True. The environment variable can be set at the job or project level.	This method is automatic and does not require manual intervention.	This method adds the startup overhead. Also, it works correctly only with single instance jobs.
Copy the transform libraries automatically at compile time.	Set the APT_DIST_TRANSFORM_OPERATOR environment variable to a list of all server names participating in the cluster. The environment variable can be set at the job or project level. The list of server names is merged with the list of fastnames referenced in the configuration file that is used to compile the jobs. If no configuration file is specified, the default configuration file is used.	This method is also automatic and does not require manual intervention. This method does not incur any startup overhead. The method works correctly with both single and multi-instance jobs. This method ensures that compiled transform libraries are distributed to all participating compute nodes regardless of the configuration file that is used at compile time.	This method requires that a compiler be present on the conductor node. This requirement could be difficult to meet in production environments.

Note: If you use compile-time distribution you do not need to use runtime distribution.

Setting up a C++ compiler

The computer where you install the IBM InfoSphere Information Server engine must have a C++ compiler. The job compilation process uses it to compile parallel jobs with transformer stages, buildop components, and some lookup stages. To function correctly, the transform feature of the parallel canvas requires the compiler. The parallel canvas requires that the compiler and SDK be installed in their default locations.

For some compilers, each developer must have a license at the time that the developer compiles the job with the Designer client. The maximum number of simultaneous processes that compile jobs determines the number of licenses.

Install only one compiler on each engine tier for development environments. For production environments, you do not need a C++ compiler unless you have set up a parallel processing environment in which transform libraries are distributed at compile time. Also, some operating systems require additional C++ runtime libraries.

For the list of supported C++ compilers, see the system requirements for your operating system: www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.

When InfoSphere Information Server is installed, the compiler settings are preset. However, for some compilers you might need to alter the settings after installation.

Setting up a compiler on IBM AIX

For AIX systems, install a supported compiler and then set compiler parameters before you create and run parallel jobs.

About this task

Complete this task for every engine tier computer where jobs will be recompiled. The compiler settings must be specified for each project that requires them.

Procedure

1. Install the compiler. For information about supported compilers, see the system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.
2. Log into the IBM InfoSphere DataStage and QualityStage Administrator using an account with administrator authority.
3. Select **Projects**.
4. For each project listed in the window, complete these steps to update compiler settings so that the transforms can be compiled successfully:
   a. Click the project name, and click **Properties**.
   b. Click **Environment**.
   c. Select **Parallel** > **Compiler**.
   d. Select **APT_COMPILEOPT**. In the **Value** column, change the value to `-c -O -q64`
   e. Select **APT_COMPILER**. In the **Value** column, change the value to `/usr/vacpp/bin/xlC_r`
   f. Select **APT_LINKER**. In the **Value** column, change the value to `/usr/vacpp/bin/xlC_r`
   g. Select **APT_LINKOPT**. In the **Value** column, change the value to `-G -q64`
   h. Click **OK**.
   i. Click **OK**.
5. Click **Close**.

Setting up a GCC compiler (Linux on 64-bit Intel systems)

For Linux on 64-bit Intel systems, install the GCC compiler and then set compiler parameters before you create and run parallel jobs.

About this task

Complete this task for every engine tier computer where jobs will be recompiled. The compiler settings must be specified for each project that requires them.

Procedure

1. Install the compiler. For information about supported compilers, see the system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.
2. Log into the IBM InfoSphere DataStage and QualityStage Administrator using an account with administrator authority.
3. Select **Projects**.
4. For each project listed in the window, complete these steps to update compiler settings so that the transforms can be compiled successfully:
   a. Click the project name, and click **Properties**.
   b. Click **Environment**.
   c. Select **Parallel** > **Compiler**.
   d. Select **APT_COMPILEOPT**. In the **Value** column, change the value to `-O -fPIC -Wno-deprecated -c -m64 -mtune=generic -mcmodel=small`
   e. Select **APT_COMPILER**. In the **Value** column, change the value to `g++`
   f. Select **APT_LINKER**. In the **Value** column, change the value to `g++`
   g. Select **APT_LINKOPT**. In the **Value** column, change the value to `-shared -m64`
   h. Click **OK**.
   i. Click **OK**.
5. Click **Close**.

Setting up a Microsoft Visual Studio .NET C++ compiler

For Microsoft Visual Studio .NET compilers, set LIB and INCLUDE environment variables before you create and run parallel jobs.

About this task

Complete this task for every engine tier computer where jobs will be recompiled.

Procedure

1. Install the compiler. For information about supported compilers, see the system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html. For more information about compiler installation, see your Microsoft documentation.
2. Within Microsoft Windows, set the LIB and INCLUDE environment variables to the appropriate locations for your compiler version. The following table is an example of what to set for Microsoft Visual Studio .NET 2010 or Microsoft Visual Studio .NET 2010 Express® Edition C++:

Table 95. Environment variables to set for Microsoft Visual Studio .NET 2010 or Microsoft Visual Studio .NET 2010 Express Edition C++

Environment variable	Setting
LIB	Set to the location of the 32-bit library directory for the SDK. For example, for Microsoft Visual Studio .NET 2010 Express Edition C++, a typical location is `C:\Program Files (x86)\Microsoft SDKs\Windows\v7.0A\Lib`.
INCLUDE	Set to the location of the 32-bit include directory for the SDK. For example, for Microsoft Visual Studio .NET 2010 Express Edition C++, a typical location is `C:\Program Files (x86)\Microsoft SDKs\Windows\v7.0A\Include`.

The version number in the path (`v7.0A` in the previous examples) might vary based upon the latest updates to the compiler.

Note: The LIB and INCLUDE environment variable names must be specified in uppercase characters.

3. Restart the computer for the environment variable settings to take effect.

Setting environment variables for the parallel engine

You set environment variables to ensure smooth operation of the parallel engine. Environment variables are set on a per-project basis.

About this task

You can set the environment variables by using the IBM InfoSphere DataStage and QualityStage Administrator or the `dsadmin` command line tool.

Procedure

Set the environment variables using one of the following methods:

- Setting environment variables by using the Administrator client:
  1. 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.
- Setting environment variables by using **`dsadmin`** tool:

  Use the **`-envadd`** option to add the environment variables mentioned in the next section.

```
cd $DSHOME
. ./dsenv
dsadmin -envadd EnvVarName -type STRING -prompt "PromptText"
[-value "Value"] ProjectName
```

Environment variables for the parallel engine

Set the listed environment variables depending on whether your environment meets the conditions stated in each variable.

See the chapter regarding environment variables in the *IBM InfoSphere DataStage and QualityStage Parallel Job Advanced Developer's Guide* for a complete listing of environment variables that you can set.

Network settings

APT_IO_MAXIMUM_OUTSTANDING

If the system connects to multiple processing nodes through a network, 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 the setting 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

If any of the stages within a job has a large number of communication links between nodes, specify this environment variable with the TCP/IP buffer space that is allocated for each connection. Specify the value 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.

APT_SENDBUFSIZE

If any of the stages within a job has a large number of communication links between nodes, specify this environment variable with the TCP/IP buffer space that is allocated for each connection. Specify the value 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.

Transform library

If you are working on a non-NFS MPP system, set the following environment variables:

- Set the APT_COPY_TRANSFORM_OPERATOR environment variable to true. This environment variable is set to false by default. IBM InfoSphere DataStage and QualityStage Administrator users must have the appropriate privileges to create project directory paths on all the remote nodes at runtime.
- Set the APT_DISABLE_TRACKER_STARTUP environment variable to 1. The value '1' allows the job to run without starting the Resource Tracker. You can set this environment variable in the Administrator client Projects tab, on the Parallel page. You must have the appropriate level of administrator privilege to set the environment variable.

Job monitoring

By default, the job monitor uses time-based monitoring. 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. If you override the default setting for the APT_MONITOR_TIME, the setting of the APT_MONITOR_SIZE environment variable is also overridden.

Detailed information about jobs

To produce detailed information about jobs as they run, set the APT_DUMP_SCORE value to True. 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 `/tmp` directory 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.

Configuring the parallel engine (AIX Linux)

Complete the following tasks to configure the parallel engine.

Setting the maximum number of processes for parallel jobs (AIX)

The setting of the **maxuproc** kernel parameter on each parallel engine node controls the maximum number of processes per user for that node. Use the **chdev** command to set an appropriate value.

About this task

A user who runs parallel jobs must be able to run at least 1000 processes on each processing node. Large systems and symmetric multiprocessing (SMP) servers require more than 1000 processes. Complete the following steps on each computer that acts as a processing node in your parallel environment.

Procedure

1. Issue this command to determine the current setting for the number of processes:

   ```
   /usr/sbin/lsattr -E -l sys0 | grep maxuproc
   ```

2. If the value of the **maxuproc** parameter is less than 1000, enter this command to increase it to 1000:

   ```
   /usr/sbin/chdev -l sys0 -a maxuproc=1000
   ```

This command must be issued on the node by a user with root privileges.

3. In an AIX cluster, you can use the dsh command to set the maximum number of processes per user to 1000 on all cluster nodes.

   ```
   dsh -a chdev -l sys0 -a maxuproc=1000
   ```

   This command must be run from the control workstation with root privileges. You can also use the **smit** or **smitty** commands to perform this action.

Setting per-user kernel parameters for parallel jobs (AIX)

Per-user kernel parameters ensure that users who run parallel jobs have sufficient resources. Use the **chuser** command to set appropriate values.

About this task

The following table describes the parameters that control memory resources and recommends a value for each parameter. Complete the following steps for each user on each processing node in your parallel environment.

Table 96. Per-user kernel parameter values

Parameter	Description	Value
fsize	The largest file that a user can create.	2 GB minus 512 bytes (4,194,303 512-byte blocks).
data	The largest data segment, or heap, that a program can have.	128 MB (262,144 512-byte blocks) or more. For no limit, specify -1.
stack	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 process can use.	64 MB (131,072 512-byte blocks) or more.

Procedure

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

   ```
   /usr/sbin/lsuser -f user_name
   ```

   This command displays various parameters, including fsize, data, stack, and rss. For example:

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

   These numbers represent the number of 512-byte blocks.

2. Log in as the root user.
3. Set the parameters for each user as defined by the values in Table 96. The command has the form:

   ```
   /usr/sbin/chuser param=value user_name
   ```

   For example:

   ```
   /usr/sbin/chuser fsize=4194303 operator2
   /usr/sbin/chuser data=262144 operator2
   /usr/sbin/chuser stack=65536 operator2
   /usr/sbin/chuser rss=131072 operator2
   ```

4. In an AIX cluster, you can use the dsh command to set the kernel parameters on all cluster nodes. This command must be issued from the control workstation.

```
dsh -a chuser param=value user_name
```

What to do next

Users whose per-user kernel parameters have changed must log out and log in again for the changes to take effect.

Setting the paging space for the parallel engine

To process parallel jobs that have large amounts of data, specify the minimum paging space settings.

Setting the paging space for the parallel engine (AIX)

To process parallel jobs that have large amounts of data, specify the minimum paging space settings.

Before you begin

You must be a system administrator logged in as the root user to reset the system paging space.

About this task

The minimum amount of paging space to run parallel jobs is 500 MB per processor for computers that have fewer than 20 processors and 1 GB per processor for computers that have 20 or more processors. The design of the parallel jobs and the number of jobs that you plan to run at one time affect the paging space requirement. Jobs that have many individual stages require more paging space than jobs that have fewer individual stages. Running more than one job at a time requires more paging space.

Complete this procedure on each computer in your parallel processing environment.

Procedure

1. Run the **lsps** command on each node or workstation to display the current size of the paging space:

   ```
   lsps -a
   ```

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

Table 97. Results of lsps command

Page space	Physical volume	Volume	Group size	% Used	Active	Auto	Type
paging01	hdisk1	pagevg	2000 MB	1	yes	yes	lv
paging00	hdisk2	rootvg	192 MB	5	yes	yes	lv
hd6	hdisk0	rootvg	64 MB	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 MB + 192 MB + 64 MB = 2256 MB (approximately 2.2 GB).

2. See the AIX documentation for additional information about setting the paging space.

Setting the paging space for the parallel engine (Linux)

To process parallel jobs that have large amounts of data, specify the minimum paging space settings.

Before you begin

You must be a system administrator logged in as the root user to reset the system paging space.

About this task

The minimum amount of paging space to run parallel jobs is 500 MB per processor for computers that have fewer than 20 processors and 1 GB per processor for computers that have 20 or more processors. The design of the parallel jobs and the number of jobs that you plan to run at one time affect the paging space requirement. Jobs that have many individual stages require more paging space than jobs that have fewer individual stages. Running more than one job at a time requires more paging space.

Procedure

See the documentation for your Linux distribution for information about setting the paging space for your system. Configure the paging space on each computer in your parallel processing environment.

Updating the /etc/magic file

Edit the `/etc/magic` file to include data set descriptor files for the parallel engine.

About this task

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

The *`install_dir`*`/etc/magic` file contains additions to the `/etc/magic` file that are specific to the parallel engine. The *install_dir* directory is the top-level directory of the parallel engine installation. The default path is `/opt/IBM/InformationServer/Server/PXEngine/`.

Procedure

1. Log in as the root user.
2. Back up the /etc/magic file.

   ```
   cp /etc/magic /etc/magic.bak
   ```
3. Edit the /etc/magic file and add the contents of the *`install_dir`*`/etc/magic` to the /etc/magic file. For example: `0 string Torrent IBM Information Server parallel engine dataset descriptor`
4. Save the /etc/magic file.
5. Test that the /etc/magic file has been updated correctly by running the file command on a data set descriptor file.

   `file` *`file_name`*

   where *file_name* is the name of a data set descriptor file. For example, if you enter `file t1.ds`, the following command returns:

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

Configuring the parallel engine (Windows Server)

You can configure the parallel engine for Microsoft Windows Server, 32-bit and 64-bit editions, by editing the Windows Registry and by changing system settings that control available memory. The Windows Registry includes keys that control desktop heap size, swap area size, the number of available TCP connections, and memory management.

Configuring the Windows registry: Enabling auto-tuning for paged pool size

If the parallel engine is installed on a computer that runs Microsoft Windows Server, Standard or Enterprise edition, enable auto-tuning for the paged pool size. Enabling auto-tuning ensures that sufficient memory is available for paged pools.

Before you begin

1. Log in to Windows Server as the Administrator.
2. Start the Registry Editor, and back up the Windows Registry. If you plan to make several changes (for example, to configure the parallel engine), create a single backup before you make all the changes.

About this task

The Windows Server kernel allocates memory in pools. These pools are known as the paged pool and the non-paged pool. Performance degradation and server instability might result if the memory for these pools is exhausted. To avoid this situation, you can enable auto-tuning at server startup by editing the **PagedPoolSize** registry value in the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management` registry subkey.

Procedure

1. In the Registry Editor, select the following registry subkey:

   ```
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\
   Memory Management
   ```

2. In the right pane of the Registry Editor, right-click **PagedPoolSize**, and click **Modify.**
3. In the **Base** area, select **Hexadecimal**.
4. In the **Value data** field, specify 0 (zero). If 0 is already specified, auto-tuning is enabled, and you do not need to change this value. If you must specify a value other than 0, the value data for **SystemPages** must be 0.
5. Click **OK**.

What to do next

After you make changes to the Windows Registry, restart Windows to apply your changes. If you plan to make several changes (for example, to configure the parallel engine), make all changes before you restart Windows.

Configuring the Windows registry: Enabling auto-tuning for system page size

If the parallel engine is installed on a computer that runs Microsoft Windows Server, Standard or Enterprise edition, enable auto-tuning for the system page size. Enabling auto-tuning ensures that sufficient memory is available for page tables and page table entries.

Before you begin

1. Log in to Windows Server as the Administrator.
2. Start the Registry Editor, and back up the Windows Registry. If you plan to make several changes (for example, to configure the parallel engine), create a single backup before you make all the changes.

About this task

To support virtual to physical address translation, the Windows operating system uses page tables and page table entries (PTEs). These tables reside in kernel memory. If not enough memory is allocated for the page tables, Windows might fail to create processes, threads, and I/O buffers. Because the parallel engine creates many processes to run a job, jobs will fail at startup or during run time if Windows does not have enough resources to create processes.

You can enable auto-tuning at server startup by editing the **SystemPages** registry value in the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management` registry subkey.

Procedure

1. In the Registry Editor, select the following registry subkey:

   ```
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\
   Memory Management
   ```

2. In the right pane of the Registry Editor, right-click **SystemPages**, and click **Modify.**
3. In the **Base** area, select **Hexadecimal**.
4. In the **Value data** field, specify 0 (zero). Typically, specifying 0 for this value provides sufficient memory. If the available PTEs are exhausted, however, you can specify the maximum value, `0XFFFFFFFF`. If you must specify a value other than 0, the value data for **PagedPoolSize** must be 0.
5. Click **OK**.

What to do next

After you make changes to the Windows Registry, restart Windows to apply your changes. If you plan to make several changes (for example, to configure the parallel engine), make all changes before you restart Windows.

Configuring the Windows registry: Setting the threshold for de-committing memory

If the parallel engine is installed on a computer that runs Microsoft Windows Server, Standard or Enterprise edition, set the threshold for de-committing memory. On computers that have 1 GB or more of RAM, setting the threshold to the suggested value improves memory management.

Before you begin

1. Log in to Windows Server as the Administrator.
2. Start the Registry Editor, and back up the Windows Registry. If you plan to make several changes (for example, to configure the parallel engine), create a single backup before you make all the changes.

About this task

When Windows frees memory at a specific address, the memory can remain committed, or it can be de-committed and marked as reserved. You can improve memory management on computers that have 1 GB or more of RAM by editing the **HeapDeCommitFreeBlockThreshold** registry value in the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager` registry subkey. This registry value specifies the number of contiguous bytes of memory (around the address of the freed memory) above which the block of memory is de-committed.

Procedure

1. In the Registry Editor, select the following registry subkey:
   `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager`
2. Specify the value data for **HeapDeCommitFreeBlockThreshold** registry value:
   a. If the registry value does not exist, click **Edit** > **New** > **DWORD Value**, and type the name of the registry value.
   b. In the right pane of the Registry Editor, right-click the registry value, and click **Modify**.
   c. In the **Base** area, select **Hexadecimal**.
   d. In the **Value data** field, specify `0x00040000`.
   e. Click **OK**.

What to do next

After you make changes to the Windows Registry, restart Windows to apply your changes. If you plan to make several changes (for example, to configure the parallel engine), make all changes before you restart Windows.

Configuring the Windows registry: Increasing the noninteractive desktop heap size

If the parallel engine is installed on a computer that runs Microsoft Windows Server, Standard or Enterprise edition, increase the noninteractive desktop heap size to ensure that a sufficient number of processes can be created and run concurrently.

Before you begin

1. Log in to Windows Server as the Administrator.
2. Start the Registry Editor, and back up the Windows Registry. If you plan to make several changes (for example, to configure the parallel engine), create a single backup before you make all the changes.

About this task

IBM InfoSphere DataStage processes are created in a specific desktop heap. Each process consumes a small amount of memory from the desktop heap. If you expect

to run many InfoSphere DataStage jobs concurrently, increase the size of the noninteractive desktop heap to allow more processes to be created and run concurrently.

You increase the size of the noninteractive desktop heap by editing the **SharedSection** parameter string in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\SubSystems registry subkey. The **SharedSection** parameter string is located in the **Windows** registry value and uses the following format to specify heap size:

SharedSection=*xxxx*,*yyyy*,*zzzz*

Where:

- *xxxx* specifies the maximum size of the system-wide heap (in kilobytes)
- *yyyy* specifies the size of each desktop heap
- *zzzz* specifies the size of the desktop heap that is associated with a noninteractive Windows instance

Table 98. The default values for interactive and noninteractive desktop heaps by operating system

Operating system	Interactive desktop heap size (SharedSection 2nd value)	Noninteractive desktop heap size (SharedSection 3rd value)
Windows Server 2008 (32 bit)	12288 KB	512 KB
Windows Server 2008 (64 bit)	20480 KB	768 KB

Procedure

1. In the Registry Editor, select the following registry subkey:
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\SubSystems
2. In the right pane of the Registry Editor, locate the **Windows** registry value.
   a. Right-click the **Windows** registry value and click **Modify**.
   b. In the **Value data** field, locate the **SharedSection** parameter string, and change the last **SharedSection** parameter from 512 to 2048. For example, specify the following values for this parameter on a 64-bit computer:
      SharedSection=1024,20480,2048
      If you cannot change the size of the desktop heap to 2048, try setting the value to 768. For example, specify the following values for this parameter:
      SharedSection=1024,20480,768
3. Click **OK** and close the Registry Editor.

What to do next

After you make changes to the Windows Registry, restart Windows to apply your changes. If you plan to make several changes (for example, to configure the parallel engine), make all changes before you restart Windows.

Configuring the Windows registry: Specifying TCP/IP settings (Windows Server 2003)

If the parallel engine is installed on a computer that runs Microsoft Windows Server 2003, Standard or Enterprise edition, specify settings for TCP/IP in the Windows Registry to ensure that the required number of parallel processes can run

concurrently. This task does not apply if the parallel engine is installed on a computer that runs Windows Server 2008.

Before you begin

1. Log in to Windows Server as the Administrator.
2. Start the Registry Editor, and back up the Windows Registry. If you plan to make several changes (for example, to configure the parallel engine), create a single backup before you make all the changes.

About this task

You specify TCP/IP settings by editing the following registry values in the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters` registry subkey:

`MaxUserPort`
Specifies the maximum port number for ephemeral TCP ports.

`TcpNumConnections`
Specifies the maximum number of concurrent connections that TCP can open. This value significantly affects the number of concurrent `osh.exe` processes that are allowed. If the value for **`TcpNumConnections`** is too low, Windows cannot assign TCP ports to stages in parallel jobs, and the parallel jobs cannot run.

These keys are not added to the registry by default.

Procedure

1. In the Registry Editor, select the following registry subkey:
   `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters`
2. Specify the following registry values: **`MaxUserPort`** and **`TcpNumConnections`**.
   a. If the registry value does not exist, click **Edit** > **New** > **DWORD Value**, and type the name of the registry value.
   b. In the right pane of the Registry Editor, right-click the registry value, and click **Modify**.
   c. In the **Base** area, select **Decimal**.
   d. In the **Value data** field, specify the following values:

Registry value name	Value data	Valid range
`MaxUserPort`	65534	5000 - 65534 (decimal)
`TcpNumConnections`	65534	0 - 65534 (decimal)

   e. Click **OK**.

What to do next

After you make changes to the Windows Registry, restart Windows to apply your changes. If you plan to make several changes (for example, to configure the parallel engine), make all changes before you restart Windows.

Configuring Windows: Editing the boot.ini file

If the parallel engine is installed on a computer that runs Microsoft Windows Server, Standard or Enterprise edition, edit the `boot.ini` file to remove the /3GB switch and, if necessary, add the /PAE switch.

Before you begin

Save a copy of the `boot.ini` file before you make any changes to the file.

About this task

By default, Windows reserves 2 GB of memory for the kernel and 2 GB of memory for user mode programs. When the /3GB switch is specified in the `boot.ini` file, 3 GB of memory is reserved for user mode programs, and the amount of kernel memory is reduced by 1 GB. This setting results in an insufficient amount of memory for the kernel. To resolve this problem, remove the /3GB switch, if present, from `boot.ini` file.

In addition, on servers that have more than 4 GB of physical memory installed, specify the /PAE switch in the `boot.ini` file. The /PAE switch specifies physical address extensions that allow a 32-bit processor to access memory that is above the 4 GB limit.

Procedure

1. Click **Start** > **Control Panel** > **System**.
2. In the System Properties window, click the **Advanced** tab.
3. In the **Startup and Recovery** area, click **Settings**.
4. Click **Edit** to edit the `boot.ini` file.
5. Remove the /3GB switch, if present.
6. On servers that have more than 4 GB of physical memory installed, add the /PAE switch to the `boot.ini` file.
7. Save and close the `boot.ini` file.
8. Click **OK** twice, and then close the Control Panel window.

What to do next

Restart Windows to apply the changes.

Configuring database connectivity in a parallel environment

You can use IBM InfoSphere Information Server to access many external data sources. You can access many data sources directly through the ODBC and JDBC support that is built into IBM InfoSphere DataStage. However, for some data sources, you must install client libraries and include the path to these installed libraries in the library path environment variable. For more information, see the section about configuring access to data sources in the connector documentation.

Chapter 8. Configuring product modules in the suite

Depending on which products you installed or the tasks that you want to complete, you must complete some configuration tasks, such as configuring access to ODBC data sources for IBM InfoSphere DataStage and IBM InfoSphere QualityStage, setting up application servers to run reports, and configuring additional product modules.

Configuring IBM InfoSphere DataStage and IBM InfoSphere QualityStage

Before you use IBM InfoSphere DataStage and IBM InfoSphere QualityStage, you must perform several configuration tasks.

Procedure

Complete the following tasks to configure and administer IBM InfoSphere DataStage and IBM InfoSphere QualityStage.

1. Configure security. For more information, see the topics on managing security in the *IBM InfoSphere Information Server Administration Guide*.
2. **AIX Linux**: Ensure that InfoSphere DataStage users have the correct localization settings. If you installed globalization support (NLS), ensure that IBM InfoSphere DataStage users have the correct settings specified for localization.
3. Configure the ODBC drivers to access ODBC data sources. The drivers are installed with the client tier or the engine tier.
4. Optional: Configure the parallel processing environment by adding computers to share engine processing, setting environment variables for the parallel engine, and configuring the parallel engine.
5. Optional: **Windows**: For systems that have large numbers of simultaneously connected users or large numbers of simultaneously running jobs, configure the server engine to support the required number of open files, lock table entries, and locks that are necessary for correct engine operation. To tune the engine, modify settings in the `uvconfig` file.
6. Optional: Migrate jobs that you created in an earlier version of InfoSphere Information Server to the current version. For more information, see the *Guide to Migrating to IBM InfoSphere Information Server, Version 11.3*.
7. Optional: If you intend to develop InfoSphere QualityStage matches, ensure that a new or existing supported database is available to receive test pass results for the Match Designer. The IBM InfoSphere QualityStage Match Designer requires an ODBC data source name (DSN) to connect to the database where the Match Designer stores match pass results.
8. Optional: Set up the application server to run SQA or investigation reports. The investigation reports and Standardization Quality Assessment (SQA) reports can read data from a sequential file or a database table. You must set up the application server to find the applicable data source.
9. Optional: Create and configure the Operations Database. If you plan to use the Operations Console to monitor your InfoSphere DataStage and QualityStage engines, projects, and jobs, you must configure Operations Database to provide information to the console.

© Copyright IBM Corp. 2005, 2014

10. After you install and configure InfoSphere DataStage and InfoSphere QualityStage, test the installation by logging in to the Administrator and Designer clients.

Configuring the engine to use PAM (Linux, UNIX)

Configuring PAM on the engine tier is required only if you want the engine to authenticate through an LDAP server. If you do, you must configure the engine for PAM before configuring IBM InfoSphere Information Server to use an LDAP user registry.

Before you begin

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

About this task

Perform this task on the computer that hosts the engine tier.

Use the following procedure to configure PAM on Linux and UNIX.

Procedure

1. From the `$DSHOME` directory, set up the environment by running the following command:

   ```
   . ./dsenv
   ```

2. Stop the InfoSphere Information Server engine by running the following command:

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

3. Edit the `uvconfig` file in the `DSHOME` directory to change the setting of the `AUTHENTICATION` tunable to 1. The following example shows the `AUTHENTICATION` tunable set to 1.

   ```
   # AUTHENTICATION - Specifies the method by which UNIX user
   #   authentication is done.  Currently, the following methods
   #   are supported:
   #
   #     0)  Standard O/S Authentication (default)
   #     1)  Pluggable Authentication Module (PAM)
   #
   #   This value should only be changed with a full understanding
   #   of the implications, as improper setting of this value can
   #   lead to the environment being unusable.
   AUTHENTICATION 1
   ```

4. As the root user, edit the PAM configuration file. The file path and contents of the PAM configuration file are platform-dependent.

Operating system	Procedure
AIX	1. Obtain the operating system level, where the nnnn-nn-nn-nnnn value returned tells you the operating system base level (the first 4 digits), the technology level (the next 2 digits), the service pack (the next 2 digits), and the date of release: `oslevel -s` 2. Edit the /etc/pam.conf file, and add four lines based on your version of the operating system: • Add the following lines when the oslevel is lower than 6100-07-02-*nnnn*: `dsepam auth required /usr/lib/security/64/pam_aix` `dsepam account required /usr/lib/security/64/pam_aix` `dsepam session required /usr/lib/security/64/pam_aix` `dsepam password required /usr/lib/security/security/64/pam_aix` • Add the following lines when the oslevel is higher than or equal to than 6100-07-02-*nnnn*: `dsepam auth required pam_aix` `dsepam account required pam_aix` `dsepam session required pam_aix` `dsepam password required pam_aix` **Note:** You cannot use PAM authentication with IMPERSONATION mode turned off. For more information, see http://www-01.ibm.com/support/docview.wss?uid=swg21516230
Linux	Add the following entry to the file /etc/pam.d/dsepam. `#%PAM-1.0` `# for engine PAM authentication` `auth include system-auth` `account include system-auth` `password include system-auth` `session include system-auth`

5. Regenerate the InfoSphere Information Server engine configuration file:
   `$DSHOME/bin/uv -admin -regen`
6. Restart the InfoSphere Information Server engine:
   `$DSHOME/bin/uv -admin -start`

What to do next

Set up the InfoSphere Information Server engine tier to use the Lightweight Directory Access Protocol (LDAP) user registry.

If you have configured PAM for both the engine and services tier by using an LDAP user registry, you can share the user registry with the tiers that you configured for PAM. Both PAM configurations must point to the same user registry by using the same set of PAM modules.

Ensuring that InfoSphere DataStage users have the correct localization settings (AIX Linux)

If you installed globalization support (NLS), ensure that IBM InfoSphere DataStage users have the correct settings specified for localization.

About this task

On most systems, the locale for individual users is set by using a system-wide default. For example, on Linux systems, this default is defined in the `/etc/sysconfig/i18n` file. If the user who performed the installation used this default and all InfoSphere DataStage users are expected to inherit the same default, then you do not need to perform the steps that are described in this task.

However, in some cases, the locale that is set for an InfoSphere DataStage user might be different from the locale that was set during installation. For example, this problem might occur if the user who performed the installation used a locale that was not the default, or if a user set a private locale (for example, in their `$HOME/.profile` directory).

To ensure that InfoSphere DataStage users have the correct locale and environment settings, you can edit the environment variables in the `dsenv` file. The environment variables in the `dsenv` file affect the entire InfoSphere DataStage environment.

To define the environment variables on a per-project or per-job basis, use the IBM InfoSphere DataStage and QualityStage Administrator client or the IBM InfoSphere DataStage and QualityStage Designer client.

Perform this task as soon as possible after installation, before any jobs are stored or compiled.

Procedure

1. Log in to the operating system as the user who ran the installation program.
2. Determine the locale and encoding name to use. The encoding must correspond to one of the encodings that InfoSphere DataStage supports.

   To determine the current locale that is set for the server operating system, use the **`locale`** command:

   ```
   locale
   ```

3. Locate the following section in the `dsenv` file. By default, this file is located in `/opt/IBM/InformationServer/Server/DSEngine`.

   ```
   #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
   ```

4. Replace the occurrences of *langdef* with the locale and encoding name that were specified during installation. For example, the following settings specify that the locale is ja_JP and the encoding is UTF-8:

   ```
   LANG="ja_JP.utf8";export LANG
   LC_ALL="ja_JP.utf8";export LC_ALL
   LC_CTYPE="ja_JP.utf8";export LC_CTYPE
   LC_COLLATE="ja_JP.utf8";export LC_COLLATE
   LC_MONETARY="ja_JP.utf8";export LC_MONETARY
   LC_NUMERIC="ja_JP.utf8";export LC_NUMERIC
   LC_TIME="ja_JP.utf8";export LC_TIME
   LC_MESSAGES="ja_JP.utf8"; export LC_MESSAGES
   ```

5. Remove the comment marker (#) that is at the beginning of each line. Save the file.
6. Restart the InfoSphere Information Server engine.

What to do next

Before performing any work from the InfoSphere DataStage command line, InfoSphere DataStage users must issue the following command:

```
. $DSHOME/dsenv
```

Configuring ODBC access (AIX Linux)

If you installed ODBC drivers, you must configure the drivers to access ODBC data sources.

The DataDirect Technologies Connect ODBC drivers are installed in *install_home*/Server/branded_odbc. Documentation for configuring the ODBC drivers is provided within this directory tree.

Edit the dsenv, odbc.ini, and uvodbc.config files to set up access to ODBC data sources.

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

For information about configuring the ODBC environment for a specific database, see the *DataDirect Connect and Connect XE for ODBC Reference* manual (a PDF file) in the $DSHOME/Server/branded_odbc/books/odbc directory. Also review the odbcread.me file in the branded_odbc directory and an HTML file in the branded_odbc/odbchelp directory.

Configuring the dsenv file

For some ODBC connections, plug-ins, and connectors, and for interactions with external applications such as IBM WebSphere MQ, you must add environment variables to enable interactive use of ODBC drivers to make a connection to an ODBC data source.

Before you begin

You must be logged in as an InfoSphere DataStage administrator with the operating system credentials on the server for the InfoSphere DataStage administrator.

Back up the dsenv file before you edit it.

About this task

The dsenv file contains a series of shell arguments that are used when the engine starts. Interactive users, other programs, and scripts can use the dsenv file. For a connection that uses a wire protocol driver, you do not have to modify the dsenv file.

Procedure

1. Edit the dsenv file. The file is located in the $DSHOME/DSEngine directory, where $DSHOME identifies the InfoSphere DataStage installation directory. The default directory is /opt/IBM/InformationServer/Server/DSEngine.
2. Specify the following information in the dsenv file:
   - Environment variables for the database client software

- Database home location
- Database library directory

Table 99. Names of the library path environment variable, by operating system

Operating system	Library path environment variable
AIX	LIBPATH
Linux	LD_LIBRARY_PATH

The following examples show typical entries for commonly used databases. The entries vary slightly depending on your operating system. See the data source documentation for more information.

Sybase 11

```
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
```

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 10.5

```
DB2DIR=/opt/IBM/db2/V10.5
export DB2DIR
DB2INSTANCE=db2inst1
export DB2INSTANCE
INSTHOME=/export/home/db2inst1
export INSTHOME
PATH=$PATH:$INSTHOME/sqllib/bin:$INSTHOME/sqllib/adm:
  $INSTHOME/sqllib/misc
export PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$INSTHOME/sqllib/lib
export LD_LIBRARY_PATH
THREADS_FLAG=native
export THREADS_FLAG
```

3. Save your changes.
4. Stop and restart the IBM InfoSphere Information Server engine.

Configuring the odbc.ini file

The `odbc.ini` file provides information about connecting to all the databases and database client applications that you use. You must configure this file so that IBM InfoSphere Information Server can access ODBC data sources.

About this task

The `/opt/IBM/InformationServer/Server/branded_odbc/IBM_Tools` directory contains example `odbc.ini` files for each supported platform. Use the contents of these files to create the `odbc.ini` file for the `$DSHOME` directory. The `$DSHOME` variable identifies the InfoSphere DataStage installation directory.

Procedure

1. Navigate to the <*installation-directory*>/Server/DSEngine directory, where <*installation-directory*> is the directory where you installed InfoSphere Information Server.
2. Run the following command to create a soft link for the sample .odbc.ini file to point to the odbc.ini file.

   ```
   ln -s odbc.ini.odbc.ini
   ```

3. Modify the .odbc.ini file to match your database environment.

Specifying data source names in the uvodbc.config file

The uvodbc.config file must specify the ODBC data source name (DSN) for each database.

About this task

A copy of the uvodbc.config file exists in the $DSHOME directory and each project directory. For example, the default path for projects on Linux is /opt/IBM/InformationServer/Server/Projects.

By default, IBM InfoSphere DataStage searches the current project directory for a uvodbc.config file and, if it finds one, uses this file instead of $DSHOME. This file arrangement allows you to configure data sources that are known to some projects but not others.

Procedure

Use this syntax to add entries to the uvodbc.config file:

```
<DSN_name>
DBMSTYPE = ODBC
```

DSN_name is the ODBC DSN. This name must be the same name that is specified in the .odbc.ini file. The space on each side of the equal sign is required.
In the uvodbc.config file, comment lines begin with a number sign (#). Comments that are not in the comment block at the top of the file cause processing errors. Add comments only to the comment block that is at the top of the uvodbc.config file.
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
```

What to do next

If you alter `uvodbc.config` after creating projects, copy the edited file to the project directories.

Creating and binding a DB2 package to a data source name

To use a DB2 package, you must bind it to each DB2 data source name (DSN).

Procedure

1. Open a terminal window, and change to the `$DSHOME/../branded_odbc/bin` directory. The `$DSHOME` variable identifies the InfoSphere DataStage installation directory. The default directory is `/opt/IBM/InformationServer/Server/DSEngine`.
2. Enter the following command:

   bind22 *dsn*

   dsn is the ODBC data source name.
3. Specify your user ID and password, if you are prompted to provide them.

Configuring other ODBC drivers

You can use ODBC drivers from other vendors to connect to data sources. You might be required to connect to a database not supported by the DataDirect driver pack, for example, Red Brick®. You can use only one ODBC Manager at a time. As a result, you cannot use drivers from other vendors at the same time that you use the IBM InfoSphere DataStage drivers, and you must perform some reconfiguration to switch between the two types.

Procedure

1. Install the ODBC manager and drivers. Use the documentation that the driver vendor provides.
2. Configure the computer to connect to the driver.
   a. In the `dsenv` file, add the required environment variables.
   b. In the `odbc.ini` file, set up the connection to the database or to the database client application. You must specify the Data Source Name (DSN), driver name, and login information.
   c. In the `uvodbc.config` file, specify the DSN to use to connect to the database.
3. Include the shared library `libodbc.sl` for the ODBC Manager in the InfoSphere DataStage load library path. Edit the `dsenv` file and ensure that the path to the ODBC Manager precedes the path to `$DSHOME/../branded_odbc/lib` in the shared library path environment variable.
4. After you install and configure the driver, test the connectivity.

Results

To use the DataDirect drivers again, remove the path to the ODBC Manager from the `dsenv` file, and ensure that the InfoSphere DataStage ODBC driver library precedes any other ODBC driver libraries.

Testing ODBC driver connectivity

After you configure ODBC connectivity that you defined for IBM InfoSphere DataStage, test the configuration to verify your connection to the data source.

Procedure

1. Change directories to the $DSHOME directory and set up the environment by running the following command:
   ```
   . ./dsenv
   ```
2. Enter the following command to start the IBM InfoSphere Information Server engine shell:
   ```
   ./bin/dssh
   ```
3. Enter the following command to log in to the project:
   ```
   LOGTO project_name
   ```
   where *project_name* is case-sensitive.
4. Enter the following command to obtain a list of available data source names:
   ```
   DS_CONNECT
   ```
5. Enter the following command to test the connection:
   ```
   DS_CONNECT dsn
   ```
   where *dsn* is the data source name of the connection to test.
6. Enter your user name and password to connect to the specified database.
7. Enter .Q to quit the connection.

Configuring ODBC access (Microsoft Windows)

If you installed ODBC drivers, you must configure the drivers to access ODBC data sources.

The Progress DataDirect Technologies ODBC drivers are installed in the *install_dir*\ODBCDrivers directory, where *install_dir* is the directory where you installed IBM InfoSphere Information Server. To configure the ODBC Data Source Name (DSN), use the Data Source Administrator application that is located in the following directory.

Table 100. Directory for the ODBC Data Source Administrator application

Version	Directory
Microsoft Windows 32-bit	C:\Windows\System32\odbcad32.exe
Microsoft Windows 64-bit	C:\Windows\sysWOW64\odbcad32.exe

Tuning the InfoSphere Information Server engine for large numbers of users or jobs (Windows Server)

For Microsoft Windows Server systems that have large numbers of simultaneously connected users or large numbers of simultaneously running jobs, you can configure the IBM InfoSphere Information Server engine to support the number of open files, lock table entries, and locks that are necessary for correct engine operation. To tune the engine, modify settings in the uvconfig file.

Before you begin

- Back up the uvconfig file. By default, this file is located in the following directory: C:\IBM\InformationServer\Server\DSEngine.
- Stop all IBM InfoSphere DataStage jobs and ask all users to log out of InfoSphere DataStage.

About this task

For new installations of IBM InfoSphere Information Server, change the settings described in this task if you are configuring production systems that have 80 or more simultaneously connected users or 50 or more simultaneously running jobs.

For existing installations, change these settings if you encounter resource problems while running InfoSphere DataStage jobs.

Procedure

1. Perform the following steps to verify that no InfoSphere DataStage jobs or `osh.exe` processes are running on the server and that no connections are active:
   a. From the Microsoft Windows Task Manager, click the **Processes** tab, and then click the **Image Name** column to sort the process list by name.
   b. Verify that the **Show processes from all users** check box is selected.
   c. In the list of processes, verify that the following processes are not running: `uvsh.exe` and `osh.exe`. If these processes are not running, there are no InfoSphere DataStage jobs running on the server.
   d. In the list of processes, verify that the following processes are not running: `dsapi_server.exe` and `dsapi_slave.exe`. If these processes are not running, there are no active InfoSphere DataStage connections.
2. Stop the InfoSphere Information Server engine.
3. Use a text editor to edit the `uvconfig` file. By default, this file is located in the following directory: `C:\IBM\InformationServer\Server\DSEngine`. Specify the following settings

   ```
   MFILES    200
   T30FILES  300
   RLTABSZ   300
   MAXRLOCK  299
   ```

4. Save the `uvconfig` file. Make sure that you do not save the file with a `.txt` extension. The `uvconfig` file cannot have a file extension.
5. From the Windows command prompt, change to the directory that contains the `uvconfig` file. For example, change to the following directory:

   ```
   C:\IBM\InformationServer\Server\DSEngine
   ```

6. Issue the following command to regenerate the binary `uvconfig` file:

   ```
   bin\uvregen.exe
   ```

7. Restart the InfoSphere Information Server engine.

Configuring IBM InfoSphere QualityStage Match Designer

The IBM InfoSphere QualityStage Match Designer requires an ODBC data source name (DSN) to connect to the database where the Match Designer stores match pass results.

Before you begin

Designate a new or existing database to be dedicated to the storage of Match Designer output. See “Creating the IBM InfoSphere QualityStage Match Designer database” on page 143.

Procedure

1. On each client tier and engine tier computer, define a System DSN that points to the Match Designer database. The ODBC DSN must have the same name on

each computer. You might want to select from the set of IBM Wire Protocol ODBC drivers installed as part of the IBM InfoSphere Information Server environment. These drivers are convenient because their functionality does not depend on the presence of any additional database-specific client software.

2. On each client tier computer, open the Match Designer and run a Test Environment Update for your match specification. When you run this function, the connection to the Match Designer database is established and your encrypted credentials are retained.

Setting up application servers to run investigation, SQA, or Match Statistics reports

The investigation, Standardization Quality Assessment (SQA), and Match Statistics reports can read data from a sequential file or a database table. You must set up the application server to find the applicable data source.

Procedure

1. If the stage writes the report data to a sequential file, perform the action that corresponds to the topology of your installation.

Topology	Action
IBM InfoSphere Information Server engine and the services tier are on the same computer.	No additional steps are necessary.
InfoSphere Information Server engine and the services tier are on different computers.	The sequential file exists only on the engine. Because the reports run on the application server, perform one of the following steps: • Create a network shared directory that both the engine and application server can access. • Copy the sequential file from the engine tier computer to the local file system of the services tier. Ensure that the file paths are identical on both the engine tier computer and the services tier computer.

2. If the stage writes the report data to a database table, perform the action that corresponds to the topology and operating system of your installation.

Topology and operating system	Action
• All operating systems • InfoSphere Information Server engine and the services tier are on different computers	Because the necessary ODBC drivers are not installed on the application server, perform the following steps: 1. Set up the JDBC data source on the application server. 2. Ensure that the JNDI name is the same as the ODBC DSN on the engine.

Topology and operating system	Action
• AIX • InfoSphere Information Server engine and the services tier are on the same computer	Because the ODBC DSN is not visible to the engine as the system DSN, perform the following steps: 1. Set up the Java Database Connectivity (JDBC) data source on the application server. 2. Ensure that the Java Naming and Directory Interface (JNDI) name is the same as the ODBC DSN on the engine.
• Linux • InfoSphere Information Server engine and the services tier are on the same computer	Because the ODBC DSN is not visible to the engine as the system DSN, perform the following steps: 1. Set up the Java Database Connectivity (JDBC) data source on the application server. 2. Ensure that the Java Naming and Directory Interface (JNDI) name is the same as the ODBC DSN on the engine.
• Microsoft Windows • InfoSphere Information Server engine and the services tier are on the same computer	No additional steps are necessary. The reports run on the application server. They can find the database table by using the Open Database Connectivity (ODBC) data source name (DSN) that was set up on the engine.

Configuring the Oozie Workflow Activity stage (Linux)

If you installed IBM InfoSphere BigInsights, you can use the Oozie Workflow Activity stage to invoke Oozie jobs from the Designer client. To use the Oozie Workflow Activity stage, you must extract the files that the Oozie client API requires to run the invoker utility.

About this task

The Oozie client API is installed when you install InfoSphere BigInsights. You must ensure that InfoSphere DataStage can access the following .jar files that the Oozie client API requires. The names of the .jar files might differ, depending on the version of Oozie that you installed.

- `oozie-client-2.3.1.jar`
- `json-simple-1.1.jar`
- `commons-cli-1.2.jar`

This procedure uses InfoSphere BigInsights as an example, but the procedure applies to any Hadoop installation. On other systems, such as Cloudera, you must configure Oozie before you can configure the Oozie Workflow Activity stage.

Procedure

1. Copy the Oozie client API files.
   a. On the machine that is running InfoSphere BigInsights, open a shell prompt and log in as the `root` user.

b. Navigate to the *$BIGINSIGHTS_HOME*/oozie directory, where *$BIGINSIGHTS_HOME* is the directory where you installed InfoSphere BigInsights.

c. Copy the Oozie client API files to the machine that is running the InfoSphere DataStage engine. The files are contained in compressed file like oozie-client-2.3.1.tar.gz. The file name might differ, depending on the version of Oozie that you installed.

2. Make the Oozie client API .jar file available to the InfoSphere DataStage engine.

a. On the machine that is running the InfoSphere DataStage engine, log in as the InfoSphere DataStage administrator user.

b. Extract the .jar files from the compressed Oozie client API file to a directory that is accessible by the InfoSphere DataStage administrator user. For example, /home/dsadm/oozie.

c. Run the **chmod** command to grant InfoSphere DataStage access to the .jar files.

```
chmod o+x *
```

d. Navigate to the opt/IBM/InformationServer/Server/DSEngine directory and open the dsenv file in an editor.

e. Add the following lines to the end of the dsenv file.

```
# Oozie client library location
OOZ_PATH=/home/dsadm/oozie

# Dependent Oozie .jar files
JAR1=$OOZ_PATH/oozie-client-2.3.1.jar
JAR2=$OOZ_PATH/json-simple-1.1.jar
JAR3=$OOZ_PATH/commons-cli-1.2.jar

CLASSPATH=${JAR1}:${JAR2}:${JAR3}:${CLASSPATH}
```

OOZ_PATH is the directory where you extracted the .jar files from the compressed Oozie client API file. The file names might differ, depending on the version of Oozie that you installed.

f. Save and close the dsenv file.

3. Restart the InfoSphere DataStage engine.

a. Navigate to the opt/IBM/InformationServer/Server/DSEngine directory.

b. Run the following command to start the InfoSphere DataStage engine.

```
bin/uv -admin -stop
```

c. Run the following command to restart the InfoSphere DataStage engine.

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

Configuring IBM InfoSphere Information Governance Catalog

You must configure some settings before you start using InfoSphere Information Governance Catalog.

Before you begin

The IBM InfoSphere Information Server suite administrator must assign security roles to InfoSphere Information Governance Catalog users.

You must have the InfoSphere Information Governance Catalog Glossary Administrator or Information Asset Administrator role.

Procedure

1. On the services tier, run the following command to prevent cross-frame scripting attacks:

   AIX Linux

   ```
   cd IS_install_path/ASBServer/bin
   ./iisAdmin.sh -set -key com.ibm.iis.gov.xFrameOptions -value DENY
   ```

   Windows

   ```
   cd IS_install_path\ASBServer\bin
   iisAdmin.bat -set -key com.ibm.iis.gov.xFrameOptions -value DENY
   ```

2. On the services tier, edit the following file and add the line `com.ibm.iis.http.soTimeout=600000` to enable you to import large files by using **istool**. If this parameter is not set, the import fails.

   AIX Linux

   ```
   IS_install_path/ASBNode/eclipse/plugins/com.ibm.iis.client/
   iis.client.site.properties
   ```

   Windows

   ```
   IS_install_path\ASBNode\eclipse\plugins\com.ibm.iis.client\
   iis.client.site.properties
   ```

3. If you intend to run lineage reports from IBM Cognos®, you must change the Cognos Metadata Information Service URI. See http://www.ibm.com/support/docview.wss?uid=swg21672175 for more information.
4. Assign security roles to InfoSphere Information Governance Catalog users. This step must be done by the InfoSphere Information Server suite administrator.
5. On Microsoft Windows Server 2008, disable the Enhanced Internet Security feature for Internet Explorer. This feature must be disabled for certain InfoSphere Information Governance Catalog features to run correctly.
   a. Click **Start** > **Administrative Tools** > **Server Manager**.
   b. In the Server Manager window, expand the Security Information section and click **Configure IE ESC**.
   c. In the Internet Explorer Enhanced Security Configuration window, under Administrators, click the **Off** radio button and then click **OK**.
6. Log in to InfoSphere Information Governance Catalog as a user with the Information Governance Catalog Glossary Administrator or Information Asset Administrator role.
   a. Click **Administration**.
   b. Expand **Setup**.
      - Click **Property Display** to configure how to display properties that have no values.
      - Click **Information Asset Type Display** to configure which asset types to display.
      - Click **General Settings** to configure your settings.
   a. Expand **Catalog Management**.
      - Click **Catalog Permissions** to restrict or grant access to glossary content in the catalog.
      - Click **Workflow** to support an approval process to update and create glossary content before it is published to the catalog.
7. Optional: Configure support for international languages and for bidirectional language support.

8. Replace the old URL address of IBM InfoSphere Business Glossary and of IBM InfoSphere Metadata Workbench with the URL for InfoSphere Information Governance Catalog.
   a. If the services tier uses IBM WebSphere Application Server Network Deployment, create an **iis.https** server variable that is set to the new port. Do this for each server. To set the server variable:
      1) Log in to the WebSphere Application Server Administrative console.
      2) Navigate to the location to add a new server variable: **Servers** > **Server Types** > **WebSphere Application Servers** > **<*server_name*>** > **Server Infrastructure** > **Java and Process Management** > **Process Definition** > **Additional Properties** > **Java Virtual Machine** > **Additional Properties** > **Custom Properties**
      3) Click **New**, add the name **iis.https** and the new port value, click **OK**, and save the changes to the master configuration.
   b. Update scripts, web browser Favorites, REST API programs, and so forth.

   Important: This console redirection will not be available after InfoSphere Information Governance Catalog, version 11.3.

Stopping and starting the IBM InfoSphere Information Server server engine

Restart the InfoSphere Information Server engine after you complete certain tasks such as editing the `dsenv` file or modifying the `uvconfig` file.

Procedure

- AIX Linux
  1. Log in to the engine tier computer as the IBM InfoSphere DataStage administrator (typically dsadm).
  2. Change to the engine directory and set the environment.
     ```
     cd $DSHOME
     . ./dsenv
     ```
     `$DSHOME` is the name of the engine directory. The default directory is `/opt/IBM/InformationServer/Server/DSEngine`.
  3. Stop the engine.
     ```
     ./bin/uv -admin -stop
     ```
     The command displays a message when the engine stops.
  4. Restart the engine.
     ```
     ./bin/uv -admin -start
     ```
- Microsoft Windows
  1. In the Windows control panel, open the DataStage Control Panel.
  2. Click **Stop all services**. Wait until the services have stopped.
  3. Click **Start all services**.

Chapter 9. Processes used by InfoSphere Information Server

The following processes are ones that you might see running in InfoSphere Information Server.

This list also includes processes for DB2 Database and WebSphere Application Server, among other processes.

Note: This list is not exhaustive. Other processes or services might run within your environment.

Windows

DB2 processes

db2dasrrm.exe
db2fmp64.exe
db2mgmtsvc.exe
db2rcmd.exe
db2syscs.exe

InfoSphere Information Server processes

Admin.exe
ASBAgent.exe
Director.exe
dsapi_server.exe
dsapi_slave.exe
DSDesign.exe
dsrpcd.exe
dsservice.exe
dstelnet.exe
fdbserver.exe
java.exe
NT_OshWrapper.exe
osh.exe
resource_tracker.exe
tl_dsservice.exe

AIX Linux

DB2 processes

```
root     19072     1  0 Jun07 ?        00:01:01 /opt/IBM/db2/V10.5/bin/db2fmcd

root     16647     1  0 Jun07 pts/2    00:00:00 db2wdog 0

db2inst1 16649 16647  0 Jun07 pts/2    00:17:56 db2sysc 0

root     16650 16649  0 Jun07 pts/2    00:00:21 db2ckpwd 0

root     16651 16649  0 Jun07 pts/2    00:00:21 db2ckpwd 0

root     16652 16649  0 Jun07 pts/2    00:00:21 db2ckpwd 0

db2inst1 16656 16649  0 Jun07 pts/2    00:00:00 db2vend (PD Vendor Process - 1)
```

© Copyright IBM Corp. 2005, 2014

```
db2inst1 16665 16647  0 Jun07 pts/2    00:01:23 db2acd
         0 ,0,0,0,1,0,0,0000,1,0,8a6740,14,200014,2,0,1,1ffc0,0x210000000,0x210000000,
         1600000,128005,2,4d0012

db2fenc1 19086 16647  0 Jun07 pts/2    00:00:03 db2fmp
         ( ,0,0,0,0,0,0,0000,1,0,8a6740,14,200014,2,0,1,64ffc0,0x210000000,0x210000000,
         1600000,128005,2,b3f004a

db2fenc1 19132 16647  0 Jun07 pts/2    00:00:02 db2fmp
         ( ,1,0,0,0,0,0,0000,1,0,8a6740,14,200014,2,0,1,85ffc0,0x210000000,0x210000000,
         1600000,128005,2,b598040
```

InfoSphere Information Server processes

```
root      3713     1  0 Jun07 pts/2
          00:00:00 /opt/IBM/InformationServer/Server/DSEngine/bin/dsrpcd

root      3786     1  0 Jun07 pts/2
          00:00:01 /opt/IBM/InformationServer/ASBNode/apps/jre/bin/java
          -classpath /opt/IBM/InformationServer/Server/PXEngine/java/JobMonApp.jar:
          /opt/IBM/InformationServer/Server/PXEngine/java/xerces/xercesImpl.jar:
          /opt/IBM/InformationServer/Server/PXEngine/java/xerces/xmlParserAPIs.jar
            JobMonApp 13400 13401

dsadm     3801     1  0 Jun07 ?
          00:09:07 /opt/IBM/InformationServer/Server/PXEngine/bin/resource_tracker

root     31294     1  0 Jun07 pts/2
         00:01:44 /opt/IBM/InformationServer/ASBNode/apps/jre/bin/java
         -Duser.language=en
         -Duser.country=US
         -Djava.security.auth.login.config=
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/auth.conf
         -Dcom.ibm.CORBA.ConfigURL=file:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/sas.client.props
         -Dcom.ibm.SSL.ConfigURL=file:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/ssl.client.props
         -Dcom.ibm.CORBA.enableClientCallbacks=true
         -Dcom.ibm.CORBA.FragmentSize=128000
         -classpath /opt/IBM/InformationServer/ASBNode/conf:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/
             iis_util.jar:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/
             ACS_client.jar:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/
             ACS_common.jar:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/
             ISF_util.jar:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/
             ISF_j2ee.jar:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/
             ASB_utils.jar:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/
             ASB_util_client.jar:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/
             com.ibm.ws.ejb.thinclient_7.0.0.jar:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/
             com.ibm.ws.security.crypto.jar

root     31324 31319 0 Jun07 pts/2 00:01:36 /opt/IBM/InformationServer/ASBNode/apps/jre/bin/java
         -Xbootclasspath/a:conf:eclipse/plugins/com.ibm.iis.client
         -Xss2M
         -Xoss2M
         -Duser.language=en
         -Duser.country=US
         -Djava.ext.dirs=apps/jre/lib/ext:lib/java:eclipse/plugins:eclipse/plugins/
           com.ibm.iis.client
         -classpath conf:eclipse/plugins/com.ibm.iis.client
         -Djava.security.auth.login.config=
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/auth.conf
         -Dcom.ibm.CORBA.ConfigURL=file:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/sas.client.props
         -Dcom.ibm.SSL.ConfigURL=file:
           /opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/ssl.client.props
         -Dcom.ibm.CORBA.enableClientCallbacks=true
         -Dcom.ibm.CORBA.FragmentSize=128000 com.ascential.asb.agent.impl.AgentImpl run
dsadm     4601     1  4 Jun07 pts/0    00:00:00 /opt/IBM/InformationServer/jdk/jre/java
         -DPSMARKER=zyxOMA2049047529xyz
         -Xms32m
         -Xmx384m
         -classpath /opt/IBM/InformationServer/ASBNode/lib/java/datastage_rt_omdmonapp.jar:
         /opt/IBM/InformationServer/ASBNode/lib/java/datastage_common_lang.jar:
         /opt/IBM/InformationServer/ASBNode/lib/java/datastage_common.jar:
         /opt/IBM/InformationServer/ASBNode/lib/java/datastage_rt_common.jar:
         /opt/IBM/InformationServer/ASBNode/lib/java/datastage_java_logger.jar:
```

```
/opt/IBM/InformationServer/ASBNode/lib/java/datastag_external.jar:
/opt/IBM/InformationServer/ASBNode/eclipse/plugins/com.ibm.iis.client/isf-core.jar
com.ibm.datastage.runtime.omdmonapp.OMDMonApp -start
```

WebSphere Application Server Network Deployment processes

```
root     30600     1  0 Jun07 pts/2
01:06:02 /opt/IBM/WebSphere/AppServer/java/bin/java
-Declipse.security
-Dwas.status.socket=44248
-Dosgi.install.area=/opt/IBM/WebSphere/AppServer
-Dosgi.configuration.area=/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/configuration
-Dosgi.framework.extensions=com.ibm.cds,com.ibm.ws.eclipse.adaptors
-Xshareclasses:name=webspherev70_%g,groupAccess,nonFatal
-Xscmx50M -Xbootclasspath/p:/opt/IBM/WebSphere/AppServer/java/jre/lib/ext/ibmorb.jar:
  /opt/IBM/WebSphere/AppServer/java/jre/lib/ext/ibmext.jar
-classpath /opt/IBM/WebSphere/AppServer/profiles/InfoSphere/properties:
  /opt/IBM/WebSphere/AppServer/properties:
    /opt/IBM/WebSphere/AppServer/lib/startup.jar:
  /opt/IBM/WebSphere/AppServer/lib/bootstrap.jar:
    /opt/IBM/WebSphere/AppServer/lib/jsf-nls.jar:
  /opt/IBM/WebSphere/AppServer/lib/lmproxy.jar:
    /opt/IBM/WebSphere/AppServer/lib/urlprotocols.jar:
  /opt/IBM/WebSphere/AppServer/deploytool/itp/batchboot.jar:
  /opt/IBM/WebSphere/AppServer/deploytool/itp/batch2.jar:
  /opt/IBM/WebSphere/AppServer/java/lib/tools.jar
-Dibm.websphere.internalClassAccessMode=allow
-verbose:gc
-Xms1280m -Xmx2048m
-Xcompressedrefs
-Dws.ext.dirs=/opt/IBM/WebSphere/AppServer/java/lib:
  /opt/IBM/WebSphere/AppServer/profiles/InfoSphere/classes:
  /opt/IBM/WebSphere/AppServer/classes:/opt/IBM/WebSphere/AppServer/lib:
  /opt/IBM/WebSphere/AppServer/installedChannels:/opt/IBM/WebSphere/AppServer/lib/ext:
  /opt/IBM/WebSphere/AppServer/web/help:
  /opt/IBM/WebSphere/AppServer/deploytool/itp/plugins/com.ibm.etools.ejbdeploy/runtime
-Dderby.system.home=/opt/IBM/WebSphere/AppServer/derby
-Dcom.ibm.itp.location=/opt/IBM/WebSphere/AppServer/bin
-Djava.util.logging.configureByServer=true
-Duser.install.root=/opt/IBM/WebSphere/AppServer/profiles/InfoSphere
-Djavax.management.builder.initial=com.ibm.ws.management.PlatformMBeanServerBuilder
-Dwas.install.root=/opt/IBM/WebSphere/AppServer
-Dpython.cachedir=/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/temp/cachedir
-Djava.util.logging.manager=com.ibm.ws.bootstrap.WsLogManager
-Dserver.root=/opt/IBM/WebSphere/AppServer/profiles/InfoSphere
-Dcom.ibm.security.jgss.debug=off
-Dcom.ibm.security.krb5.Krb5Debug=off
-Duser.language=en
-Duser.country=US
-Xdisableexplicitgc
-Djava.awt.headless=true
-Djava.security.policy=/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/properties
  /server.policy
-Djava.security.auth.login.config=
  /opt/IBM/WebSphere/AppServer/profiles/InfoSphere/properties/wsjaas.conf
  com.ibm.wsspi.bootstrap.WSPreLauncher
-nosplash
-application com.ibm.ws.bootstrap.WSLauncher com.ibm.ws.runtime.WsServer
  /opt/IBM/WebSphere/AppServer/profiles/InfoSphere/config issusex1Node01Cell
    issusex1Node01 server1
```

IBM WebSphere Application Server Liberty Profile process

```
root     20429     1  0 Apr11   ?
05:42:19   /opt/IBM/InformationServer/jdk/jre/bin/java
-XX:MaxPermSize=256m -Xms1280m -Xmx2048m -Djava.awt.headless=true
-Xdisableexplicitgc -Duser.language=en -Duser.country=US
-Dorg.eclipse.emf.ecore.EPackage.Registry.INSTANCE=
  com.ibm.xmeta.emf.util.PackageRegistry
-Dorg.eclipse.emf.common.util.URI.archiveSchemes=wsjar wszip jar zip
-javaagent:/opt/IBM/InformationServer/wlp/bin/tools/ws-javaagent.jar
-jar /opt/IBM/InformationServer/wlp/bin/tools/ws-server.jar iis
--clean
```

Part 4. Applying updates or adding tiers or products

After you have installed IBM InfoSphere Information Server, you can apply updates or add additional tiers or products to your installation. You use the update installer to apply fix pack or patch updates. You use the InfoSphere Information Server installation program to install additional tiers and most additional products.

© Copyright IBM Corp. 2005, 2014

Chapter 10. Updating software

Tools are provided to help you update your software by applying patches or fix packs.

Procedure

- **Updating IBM InfoSphere Information Server**

  You apply patches or fix packs to InfoSphere Information Server by using the InfoSphere Information Server Update Installer. You should always download the latest Update Installer from Fix Central before applying patches or fix packs. See www.ibm.com/support/docview.wss?uid=swg27041875 for information on the download and use of the Update Installer.

- **Updating IBM WebSphere Application Server Network Deployment**

  You apply patches or fix packs to WebSphere Application Server Network Deployment by using the IBM Installation Manager. Installation Manager is not installed by the InfoSphere Information Server suite installer when it installs WebSphere Application Server. You install IBM Installation Manager separately. Information on installing Installation Manager can be found in the WebSphere Application Server Information Center. For example, visit WebSphere Application Server in the Knowledge Center and search your version for the topic "Installing Installation Manager" and "Preparing to install the product." Use the option to "Download the files from the Passport Advantage® site, and use local installation." After you have installed Installation Manager, you can follow the instructions in "Installing and uninstalling interim fixes and fix packs on distributed operating systems."

- **Updating DB2 for Linux, UNIX, and Windows.**

  You apply patches or fix packs to DB2 by using the installation code that is packaged with the patch or fix pack. If you installed DB2 by using the InfoSphere Information Server suite installer, then it installed DB2 Enterprise Server Edition (ESE). Instructions on applying fix packs can be found in the DB2 Information Center. For example, visit DB2 for Linux, UNIX, and Windows in the Knowledge Center and search your version of DB2 for the topic "Fix pack installation methods."

© Copyright IBM Corp. 2005, 2014

Chapter 11. Adding product modules to an installation

You can add product modules to an existing IBM InfoSphere Information Server installation by rerunning the same installation program that was used for the installation of the most current version.

Before you begin

- Choose your installation topology, review the installation checklists, and prepare the tier computers. Obtain the product media to include the installation and entitlement packages. Configure your browser correctly so that the installation program runs smoothly.
- Ensure that you have a current backup of your InfoSphere Information Server installation.
- Do not use the InfoSphere Information Server installation program to add IBM InfoSphere DataStage editions or feature packs after initial installation. Instead, enable them by using the **iisAdmin** tool. Refer to the *IBM InfoSphere Information Server Administration Guide*.
- If the WebSphere Application Server processes or node agents are currently running under a non-root user, there is no longer a need to change them to run as root before running the installation or uninstallation program. However, after running the program, the application server will have been switched to run as root. Therefore, if you want to run it as non-root again, you will need to switch it to run as the non-root user afterward.
-

  Important: During the initial installation session, install all of the product modules that you plan to use either now or in the future. If you install only some of the product modules now, and then install fixes and service packs, you might need to reinstall the fixes and service packs if you add product modules later.

About this task

Depending on the product modules that you select, you might not have to run the installation program for all tiers. Not all product modules have components on all tiers. You only need to run the installation program on the tier computers that contain components for the product modules. You can skip tier computers that do not contain components for the product modules.

The following table lists product modules and indicates the tier computers on which you must run the installation program.

Table 101. Product modules and the tiers on which you must run the installation program

Product module	Services tier	Engine tier	Repository tier	Client tier
IBM InfoSphere FastTrack	X			X
IBM InfoSphere Information Analyzer	X	X	X	X
IBM InfoSphere Information Governance Catalog	X	X		X

© Copyright IBM Corp. 2005, 2014

Table 101. Product modules and the tiers on which you must run the installation program (continued)

Product module	Services tier	Engine tier	Repository tier	Client tier
IBM InfoSphere Glossary Anywhere (requires Information Governance Catalog). The client is installed separately.	X			
Metadata interchange agent and bridges (installs IBM InfoSphere Metadata Integration Bridges for use with InfoSphere Metadata Asset Manager.				X
IBM InfoSphere Data Quality Console	X	†	X	†
All others	X	X		X

† - When you install IBM InfoSphere Data Quality Console on the engine or client tier, only asset interchange is installed. The full data quality console is not installed on these tiers.

The installation program only installs the appropriate software for the tiers on a computer. You can run the installation program on all of the computers in your configuration, and select the products to be installed each time. The installation program determines automatically which tiers exist on each computer and adds the appropriate components.

Run the program on the computers in the following order:

1. Repository tier, if the product modules you are installing require changes to this tier.

   If you preinstalled the database system for the metadata repository and configured the databases within the database system by using the scripts that are supplied with the installation media, you do not run the installation program on the repository tier computers.
2. Services tier. All product module additions require changes to the services tier.
3. Engine tier, if needed. Refer to the table to determine whether you need to run the installation program on the engine tier.

   If the product modules you are installing require an engine tier installation, and the engine tier is not present, first install the engine tier as described in the procedure.

   Note: If the metadata repository contains host computer assets with names that are not specified in all upper case (that is, lower case or mixed case), then contact IBM Support before running the installation program to install additional products on the engine tier. You can determine the case of the host computer assets (data resources) by logging into InfoSphere Metadata Asset Manager and navigating to **Repository Management** > **Browse Assets** > **Implemented Data Resources**.
4. Client, if the product modules you are installing require changes to this tier.

   The client tier installation can take place at any point in the sequence.

For example, to install IBM InfoSphere Information Analyzer, first run the installation program on the repository tier computer. Then run it on the services tier computer. Lastly, run the program on the engine tier. You can run the installation program on the client tier computers at any point in the sequence.

InfoSphere DataStage and IBM InfoSphere QualityStage share many components. For this reason, if you add one of these product modules to a system on which the other product module is installed, the installation process runs faster.

IBM InfoSphere Information Analyzer shares many of the same components as InfoSphere DataStage and InfoSphere QualityStage. The installation process for IBM InfoSphere Information Analyzer runs faster if InfoSphere DataStage or InfoSphere QualityStage is already installed.

Procedure

1. Verify that the appropriate software tiers are present in your installation to install the product module. If a required tier is not present, install and configure the tier as described in the Chapter 12, "Adding additional tiers to an installation," on page 365 topic. For example, if you installed only IBM InfoSphere FastTrack, your installation does not have an engine tier. To add InfoSphere DataStage, first create an engine tier.
2. Run the installation program on the computers.
   - Follow the prompts. When the installation program detects the existing installation and prompts you to choose whether to add additional products or tiers, select **Add products**.
   - Continue to follow the prompts to select the products to install and to complete the installation. The product modules that are displayed are only those that are part of the entitlement package that you purchased.
3. Repeat the process for each target computer.
4. Run the IBM Support Assistant Lite for InfoSphere Information Server tool to verify the installation.

What to do next

Configure the newly installed product modules.

If you disabled scheduled tasks or deployed applications before running the installation program, restart the tasks and reenable the applications.

Chapter 12. Adding additional tiers to an installation

You can add a software tier to a computer that already hosts one or more tiers, by running the installation program.

Before you begin

- Choose your installation topology, review the installation checklists, and prepare the tier computers. Obtain the product media to include the installation and entitlement packages. Configure your browser correctly so that the installation program runs smoothly.
- Ensure that you have a current backup of your InfoSphere Information Server installation.
- If the WebSphere Application Server processes or node agents are currently running under a non-root user, there is no longer a need to change them to run as root before running the installation or uninstallation program. However, after running the program, the application server will have been switched to run as root. Therefore, if you want to run it as non-root again, you will need to switch it to run as the non-root user afterward.
-

  Important: During the initial installation session, install all of the product modules that you plan to use either now or in the future. If you install only some of the product modules now, and then install fixes and service packs, you might need to reinstall the fixes and service packs if you add product modules later.

About this task

To add an additional software tier to a computer that already hosts one or more tiers, run the installation program on the target computer.

You only need to run the installation program one time on each computer even if you are adding more than one tier. The installation program installs the tiers in the correct order on the computer.

The client tier has no dependencies and can be installed by itself on a computer. However, if you decide to install any of the other tiers on the same computer where a client has already been installed, the existing client must first be uninstalled and installed again along with the additional tiers.

For client tier installations, if you want to include multiple versions of the clients on the same computer, install the earlier version clients on the computer before you install the most recent version of the client tier. Install the clients in order from earliest version to latest version.

Procedure

1. Run the installation program on the computer.
   - Follow the prompts. When the installation program detects the existing installation and prompts you to choose whether to add additional products or tiers, select **Add Tiers**.

© Copyright IBM Corp. 2005, 2014

If you are installing a client tier on a computer where an engine tier is installed, the client tier and engine tier must share the same base directory (for example, `C:\IBM\InformationServer`).

Keep the following dependencies in mind when adding additional tiers to an installation:

Table 102. Tier dependencies

Tier	Dependency
Repository tier	None
Services tier	Repository tier
Engine tier	Services tier
Client tier	None

- Continue to follow the prompts to select the products to install and to complete the installation.

2. Run the IBM Support Assistant Lite for InfoSphere Information Server tool to verify the installation. Refer to the "Running the ISALite tool to verify the installation and troubleshoot" on page 251 topic for more information.

What to do next

Configure the newly installed product modules.

If you disabled scheduled tasks or deployed applications before running the installation program, restart the tasks and re-enable the applications.

Part 5. Removing IBM InfoSphere Information Server software

To remove InfoSphere Information Server from your computers, use the software removal utilities or remove the software manually.

© Copyright IBM Corp. 2005, 2014

Chapter 13. Preparing to remove software

Before you run the software removal program to remove an installation or to remove product modules or tiers, back up critical data on your computers and make sure that the computers are ready for the removal process.

Before you begin

If you are running the software removal program in graphical mode, configure your browser so that the program runs correctly.

For an installation that involves WebSphere Application Server clustering, perform the following tasks before you begin the software removal process:

1. Start management processes, including node agents. To ensure that sufficient file descriptor resources are allocated to the processes use the **ulimit** command.
2. Verify that you can start and stop the cluster and that you can restart the node agents. See "Administering IBM InfoSphere Information Server and IBM WebSphere Application Server services" in the *IBM InfoSphere Information Server Administration Guide*.
3. Verify that you can stop and start the Web server.

Procedure

Follow this procedure for each target computer:

1. Back up the following directories, your components, and critical data.

Operating system	Directory
AIX	• /etc/services • /etc/inittab • /etc/passwd • /etc/group
Linux	• /etc/services • /etc/inittab • /etc/passwd • /etc/group
Windows	• MicrosoftWindows registry • C:\Windows\System32\drivers\etc

   See "Backing up and restoring IBM InfoSphere Information Server" in the *IBM InfoSphere Information Server Administration Guide* for more information on how to back up your current IBM InfoSphere Information Server installation.
2. **Microsoft Windows**: The software removal program requires the Windows Script Engine to operate. To test your Windows Script Engine installation, see the script engine troubleshooting topic.
3. Disable all scheduled jobs, scheduled tasks, scheduled reports, and enabled IBM InfoSphere Information Services Director applications before running the software removal program.

© Copyright IBM Corp. 2005, 2014

4. Disable firewall software and antivirus software during the software removal session.
5. **Microsoft Windows**: When you run the software removal program in graphical mode, the message, "The publisher cannot be verified" might be displayed in the browser. You can safely ignore the message and continue the software removal. This message does not indicate a security risk when it appears during InfoSphere Information Server installation or software removal.
6. **Microsoft Windows**: Ensure that the following directories are in your system path:
   - `%SystemRoot%`
   - `%SystemRoot%\system32`
   - `%SystemRoot%\system32\wbem`
7. If the WebSphere Application Server processes or node agents are currently running under a non-root user, there is no longer a need to change them to run as root before running the installation or uninstallation program. However, after running the program, the application server will have been switched to run as root. Therefore, if you want to run it as non-root again, you will need to switch it to run as the non-root user afterward.

Shutting down tasks and applications before removing software (service tier and engine tier)

Before running the software removal program, shut down scheduled tasks and disable IBM InfoSphere Information Services Director applications on both the service tier computer and the engine tier computer.

Before you begin

For installations that include the Pack for SAP BW or the Pack for SAP R/3 on the engine tier, shut down the SAP BW and SAP R/3 listeners. For details on shutting down the listeners, see the following troubleshooting technote (www.ibm.com/support/docview.wss?uid=swg21409098).

Procedure

1. Shut down all scheduled tasks, such as reports and log purge activities. To shut down these tasks, create an unfiltered schedule view within the IBM InfoSphere Information Server Web console, and then stop all scheduled tasks that are displayed in the view.
   a. Log in to the IBM InfoSphere Information Server Web console with an account that has the suite administrator role.
   b. In the Web console, click the **Administration** tab.
   c. In the Navigation pane, select **Scheduling Management** > **Views of Schedules**.
   d. In the Scheduling Views pane, click **New Scheduling View**.
   e. Enter a name and description in the **Name** and **Description** fields.
   f. Click **Save** and **Close**.
   g. Select the schedule view that you created in the previous steps.
   h. Click **View Schedules**.
   i. Select a schedule in the list, and click **View Task**. Look at the **Status** value in the **Schedule** section:

- If the **Status** value is STARTED, click **Close**, and then click **Stop** to stop the task. If you are removing a product module or a tier, write down the name of the schedule so that you can restart it after the software removal is completed.
- If the **Status** value is not STARTED, click **Close**.

j. Complete the previous step for each schedule in the list. For a product module or tier removal, do not restart the scheduled tasks until both the services and engine tiers are updated.

2. Disable any InfoSphere Information Services Director applications that are currently running.
   a. Log in to the IBM InfoSphere Information Server console with an account that has the suite administrator role.
   b. Select **OPERATE** > **Deployed Information Services Applications**.
   c. From the View pane, select a deployed application.
   d. Click **Edit**.
   e. In the bottom of the View pane, click **Disable** and select **Disable** from the menu.
   f. Repeat these steps for each deployed application. For a product module or tier removal, do not reenable the deployed applications until both the services and engine tiers are updated.

Configuring your browser for software removal

To run the software removal program successfully, you must configure your browser.

All browsers

- Verify that you have a supported browser. For details, go to the system requirements at www.ibm.com/software/data/infosphere/info-server/overview/requirements.html.
- Make sure that JavaScript is enabled.

Microsoft Internet Explorer

Restriction: Graphics in the installation wizard might not display correctly in Microsoft Internet Explorer. To see the graphics in the wizard, use Mozilla Firefox if you have it installed.

Make sure that the security zone for the computer allows the installation program to run:

1. From Microsoft Internet Explorer, in the **Tools** menu, select **Internet Options**.
2. Click the **Security** tab.
3. Click the **Trusted Sites** icon.
4. Click **Sites**.
5. In the **Add this Website to the zone** field, type `https://`*hostname* where *hostname* is the host name of the computer on which the installation program will run.
6. Click **Add**.
7. Click **OK**.
8. Click **OK**.

Mozilla Firefox

Newer versions of Firefox enable JavaScript and load images automatically by default. If you have older versions of Mozilla Firefox and have disabled these options, enable them before running the installation program:

1. In the **Options**, menu select **Options**. The Options window displays.
2. Select **Content**.
3. Enable **Load images automatically**.
4. Enable **Enable JavaScript**.
5. Click **OK**.

Chapter 14. Removing IBM InfoSphere Information Server software by using the software removal program

To remove IBM InfoSphere Information Server software from your computers, you can run the InfoSphere Information Server software removal program.

You can remove the entire software suite, individual product modules, or individual software tiers.

Attention: Removing the entire software suite, individual product modules, or individual software tiers will result in data loss. Make sure that all file systems and databases are backed up before beginning a software removal.

The first time that you run the software removal program, use graphical mode. The program displays a graphical wizard that prompts you for options. After you complete the wizard, you save your settings in a response file, and then continue using the wizard to complete the software removal. Alternatively, you can exit the program after saving the response file. The next time that you run the program in graphical mode, you can load the settings in your response file instead of specifying them again.

You can also run the program in silent mode. In silent mode, the program uses the settings in a response file and does not display the graphical interface.

Run the software removal program on all target computers, in the following order. If you are removing individual software tiers, the program removes all InfoSphere Information Server product module components that belong to the selected tiers. Also, the program removes all InfoSphere Information Server common components that belong to the selected tiers.

Attention: For client tier removals, you might only want to remove some versions. It is not possible to do this without rendering the remaining versions unusable. Therefore, to remove a particular client tier version from a computer, you first remove the entire client tier from the computer with the software removal program and reinstall the versions that you want.

You run the software removal program only once on each computer even if you are removing more than one tier. The removal program removes the tiers in the correct order. The engine tier is dependent on the services tier, and the services tier is dependent on the repository tier. If more than one tier is installed on a computer, you cannot remove a tier without removing the tiers that are dependent on it.

1. Client tier

   You can run the removal program on any point in the sequence.
2. Engine tier

   Engine, if your installation includes an engine tier.
3. Services tier

   If you installed IBM WebSphere Application Server by using the installation program, WebSphere Application Server is removed when the services tier is removed.

© Copyright IBM Corp. 2005, 2014

If you preinstalled the application server, components that are specific to InfoSphere Information Server are removed. However, the application server is not removed.

4. Repository tier

   If you installed IBM DB2 by using the installation program, the metadata repository database (and analysis database, if present) are dropped and DB2 is removed.

   If you used the installation program to create the metadata repository database or analysis database in a preinstalled DB2 instance, the databases are dropped. However, DB2 is not removed.

   If you did not create the repository tier by using the installation program, do not use the installation program to remove the repository tier.

Removing individual product modules

Depending upon the product modules that you select, you might not have to run the software removal program for all tiers. Not all product modules have components on all tiers. You only need to run the program on the tier computers that contain components for the product modules.

The software removal program removes only the appropriate software for the tiers on a computer. You can run the software removal program on all of the computers in your configuration, and select all tiers each time. The removal program determines automatically which tiers exist on each computer and removes the appropriate components.

For example, to remove InfoSphere Information Analyzer, first run the software removal program on the engine tier computer. Then, run it on the services tier computer. Lastly, run the program on the metadata repository tier. You can run the software removal program on the client tier computers at any point in the sequence.

If you remove all product modules that require a particular tier, the tier is also removed. For example, in a system that includes InfoSphere DataStage and IBM InfoSphere Information Governance Catalog, if you remove InfoSphere DataStage, the entire engine tier is removed because InfoSphere Information Governance Catalog does not require it.

The following table indicates the tier computers on which you must run the software removal program for specific product modules:

Table 103. Product modules and the tiers on which you must run the software removal program

Product module	Services tier	Engine tier	Client tier
IBM InfoSphere FastTrack	X		X
IBM InfoSphere Information Governance Catalog	X	X	X

Table 103. Product modules and the tiers on which you must run the software removal program (continued)

Product module	Services tier	Engine tier	Client tier
IBM InfoSphere Glossary Anywhere (requires IBM InfoSphere Information Governance Catalog). The client is installed separately.	X		
IBM InfoSphere Metadata Integration Bridges			X
All others	X	X	X

Specific product considerations

InfoSphere DataStage

To deactivate an InfoSphere DataStage edition or feature pack (for example, if you are no longer entitled to an item), you do not use the IBM InfoSphere Information Server software removal program. Instead, deactivate them by using the **iisAdmin** command. The format of the command is as follows:

AIX Linux

```
<IS_install_path>/ASBServer/iisAdmin.sh -set -key
<license_key> -value 0
```

Windows

```
<IS_install_path>\ASBServer\iisAdmin.bat -set -key
<license_key> -value 0
```

Table 104. Value of the license key for each product edition or feature pack

Product edition or feature pack	<license_key> value
InfoSphere DataStage (parallel and server engine)	com.ibm.iis.datastage.license.option.parallel
Server engine only	com.ibm.iis.datastage.license.option.server
IBM InfoSphere QualityStage	com.ibm.iis.datastage.license.option.qualitystage
Balanced Optimization	com.ibm.iis.datastage.license.option.BAL_OPT
SAS Pack	com.ibm.iis.datastage.license.option.SAS_PACK

InfoSphere DataStage and InfoSphere QualityStage share many components between them. For this reason, if you remove one of these product modules from the system on which the other product module is installed, the removal process is extremely rapid.

If multiple versions of InfoSphere DataStage client programs are installed on a computer and you remove the current version of the InfoSphere DataStage client programs, older versions of the client programs might not function. To retain older InfoSphere DataStage client programs, complete the following tasks.

1. Remove the current version of the InfoSphere DataStage client programs.
2. Remove all older InfoSphere DataStage client programs.
3. Reinstall the older versions, starting from the oldest version to the most recent version that you want to retain.

InfoSphere Information Analyzer
Many of these components are also shared by InfoSphere Information Analyzer. The removal process for this product module is faster if InfoSphere DataStage or InfoSphere QualityStage remains installed.

Removing IBM InfoSphere Information Server by using the software removal program (graphical mode)

The software removal program leads you through the process of removing IBM InfoSphere Information Server. To run the software removal program in graphical mode, you can use a Web browser locally on the target computer. Alternatively, you can use a browser on any system that has network access to the computer where the program is running.

Before you begin

Complete the preparation steps that are described in Chapter 13, "Preparing to remove software," on page 369.

About this task

To provide security during the software removal process, access to the software removal application is password protected. When you first start the software removal program, a non-graphical interface is displayed. You are prompted to create a password to use for the software removal session. The software removal program displays a wizard that prompts you for parameters and other data.

On Microsoft Windows computers, the installation program automatically launches the default web browser on the local computer to display the graphical interface.

After you choose your settings in the graphical interface, the program stores your settings in a response file. If you run the software removal program again, you can use the response file to perform the same software removal process again.

Consider these conditions while the software removal program is running:

- When you run `uninstall.exe` or `java.exe`, you might see a security notification. You can safely click **Run** or **Yes** to allow the program to run and continue with the uninstallation.
- If you close your browser while the wizard is running, you can continue the same session by opening the browser again and reconnecting to the session.
- If your browser displays a "server busy" message, select the **Continue** option. You can also close the browser and then reconnect to same session.
- The wizard times out if there is no activity within two hours.

Procedure

1. On Microsoft Windows computers, verify that the user is a member of the local Administrators group. You must be a local administrator or a local user who is directly assigned to the local Administrators group to run the installation program.
   a. Open a command prompt and enter `compmgmt.msc`.
   b. In the navigation pane, click **Local Users and Groups** > **Groups**.
   c. In the right pane, right-click **Administrators** and click **Properties**.
   d. In the Administrators Properties window, locate the user name in the **Members** list. If the user name is in the list, the user is part of the Administrators group.

      If the user name is not in the list, the user is not part of the group. You must use a different user account to run the installation program, or ask a system administrator to add the current account into the local Administrators group.
2. Log in to the target computer.

Operating system	Description
AIX	Log in as the root user.
Linux	Log in as the root user.
Windows	Log in as a local administrator or as a local or domain user who is directly assigned to the local Administrators group.

3. Open a command prompt or shell window.

Operating system	Command
AIX	Open a shell window on the target computer, or use a graphical desktop sharing system. Keep the shell window open during the entire installation process. If you use telnet or ssh, make sure that the session does not shut down or timeout for the duration of the installation.
Linux	Open a shell window on the target computer, or use a graphical desktop sharing system. Keep the shell window open during the entire installation process. If you use telnet or ssh, make sure that the session does not shut down or timeout for the duration of the installation.
Windows	Start an administrator session by using one of the following methods: • Run the following command: `runas /user:Administrator cmd.exe` • Open an elevated command window. 1. Click **Start** > **Search**. 2. In the Search Results box, enter `cmd` and press Ctrl+Shift+Enter. 3. Click **Continue** to accept the prompt asking for your permission to continue.

4. In the newly created session, run the following command to clear any Java options that might be set.

Operating system	Command
AIX	In the sh or bash shell: `export IBM_JAVA_OPTIONS=` In the csh or tcsh shell: `unsetenv IBM_JAVA_OPTIONS`
Linux	In the sh or bash shell: `export IBM_JAVA_OPTIONS=` In the csh or tcsh shell: `unsetenv IBM_JAVA_OPTIONS`
Windows	`set IBM_JAVA_OPTIONS=`

5. **AIX Linux**: Configure the following settings for the session to help ensure that the installation is successful:
   - Set the file creation mask to 022:
     `umask 022`
   - Ensure that you have no less than 10240 file descriptor resources allocated.
   - Check that the current setting for the NOFILES kernel parameter is set to a value that is at least equal to the value that you specify for **`ulimit`**.
6. Run the software removal program. The following commands start the software removal process in graphical mode and include a response file.

Operating system	Procedure
AIX	1. Change to the `_uninstall` directory within the InfoSphere Information Server installation directory. By default, this directory is `/opt/IBM/InformationServer`. 2. Enter the following command to start the removal process.`./uninstall -uimode graphical -rsp` *`file_name`*
Linux	1. Change to the `_uninstall` directory within the InfoSphere Information Server installation directory. By default, this directory is `/opt/IBM/InformationServer`. 2. Enter the following command to start the removal process.`./uninstall -uimode graphical -rsp` *`file_name`*
Windows	1. Change to the `_uninstall` directory within the InfoSphere Information Server installation directory. By default, this directory is `C:\IBM\InformationServer`. 2. Enter the following command: `uninstall -uimode graphical -rsp` *`response_file_name`*

A welcome message is displayed. Then the program describes how to access the graphical interface:

`======> Installer available at: https://`*`hostname`*`:8443/ISInstall`

Specify options to run the software removal program

After you run the software removal program, you must open a web browser, enter the URL that was specified by the removal program, and configure options before proceeding with the removal process. You can then specify options for each of the tier computers by using the InfoSphere Information Server removal program.

Procedure

1. Open a web browser on the computer or on a remote computer, and navigate to one of the addresses that is listed in the messages from starting the installation program.

   The software removal program automatically launches the default Web browser on the local computer and navigates to the HTTP address.

   A Login Uninstallation window is displayed in the web browser.

   A message might appear in the browser that reads "The publisher cannot be verified". You can safely ignore the message and continue the removal process. This message does not indicate a security risk when it appears during the InfoSphere Information Server removal process.
2. To display installation messages in a language other than English, select the language in the **Select a language** menu, and click **Login**.
3. Accept the license agreement.

Saving a response file and removing the product

You can save a response file to tell the IBM InfoSphere Information Server removal program how to perform a software removal by using your saved configuration information. Saving a response file is important if you want to cancel the current removal session and use the response file in a later session.

Procedure

1. In the software removal wizard, specify a name and location for the response file in the **Response file location** field (at this point, you can choose to cancel the removal and use the saved response file in a later session).
2. Click **Remove IBM InfoSphere Information Server and save my settings in a response file**, and then click **Next** to begin the product removal. The program displays a summary of the choices that you made in the wizard. As the removal proceeds, the software removal program reports the progress in the window.

   Consider these conditions while the software removal program is running:

   - Leave the command prompt open until the software removal is complete.
   - After you have started the software removal process, keep the web browser open to monitor progress. If errors occur, the uninstallation continues. If you close the web browser, the software removal process continues to the end. You can reconnect to the session by opening a new web browser to the same URL. You can also monitor progress by reading the log files.
   - Microsoft Windows: Depending on your security settings, warning windows might display during the removal of the IBM DB2 database system or other components. Ignore the warnings, and click **OK** to continue the removal process.

Removing IBM InfoSphere Information Server by using a response file (silent mode)

Using silent mode, you can run an unattended removal of InfoSphere Information Server software and complete multiple similar software removals quickly. For example, if you want to remove identical client tier software from multiple workstations, you can create a single response file and use silent mode to remove the software from each workstation, rather than repeatedly using the graphical user interface.

Before you begin

Ensure that you have created a response file. To create a response file, run an uninstallation in graphical mode and save the response file.

Note: Passwords are not saved in the response file. To use the response file, you must first edit the response file with a text editor and add passwords.

Procedure

1. On Microsoft Windows computers, verify that the user is a member of the local Administrators group. You must be a local administrator or a local user who is directly assigned to the local Administrators group to run the installation program.
   a. Open a command prompt and enter `compmgmt.msc`.
   b. In the navigation pane, click **Local Users and Groups** > **Groups**.
   c. In the right pane, right-click **Administrators** and click **Properties**.
   d. In the Administrators Properties window, locate the user name in the **Members** list. If the user name is in the list, the user is part of the Administrators group.

      If the user name is not in the list, the user is not part of the group. You must use a different user account to run the installation program, or ask a system administrator to add the current account into the local Administrators group.
2. Log in to the target computer.

Operating system	Description
AIX	Log in as the root user.
Linux	Log in as the root user.
Windows	Log in as a local administrator or as a local or domain user who is directly assigned to the local Administrators group.

3. Open a command prompt or shell window.

Operating system	Command
AIX	Open a shell window on the target computer, or use a graphical desktop sharing system. Keep the shell window open during the entire installation process. If you use telnet or ssh, make sure that the session does not shut down or timeout for the duration of the installation.

Operating system	Command
Linux	Open a shell window on the target computer, or use a graphical desktop sharing system. Keep the shell window open during the entire installation process. If you use telnet or ssh, make sure that the session does not shut down or timeout for the duration of the installation.
Windows	Start an administrator session by using one of the following methods: • Run the following command: `runas /user:Administrator cmd.exe` • Open an elevated command window. 1. Click **Start** > **Search**. 2. In the Search Results box, enter `cmd` and press Ctrl+Shift+Enter. 3. Click **Continue** to accept the prompt asking for your permission to continue.

4. In the newly created session, run the following command to clear any Java options that might be set.

Operating system	Command
AIX	In the sh or bash shell: `export IBM_JAVA_OPTIONS=` In the csh or tcsh shell: `unsetenv IBM_JAVA_OPTIONS`
Linux	In the sh or bash shell: `export IBM_JAVA_OPTIONS=` In the csh or tcsh shell: `unsetenv IBM_JAVA_OPTIONS`
Windows	`set IBM_JAVA_OPTIONS=`

5. **AIX Linux**: Configure the following settings for the session to help ensure that the installation is successful:
   - Set the file creation mask to 022:
     `umask 022`
   - Ensure that you have no less than 10240 file descriptor resources allocated.
   - Check that the current setting for the NOFILES kernel parameter is set to a value that is at least equal to the value that you specify for **ulimit**.
6. Run the software removal program. The following commands starts the software removal process in silent mode by using a response file.

Operating system	Procedure
AIX	1. Change to the `_uninstall` directory within the InfoSphere Information Server installation directory. By default, this directory is `/opt/IBM/InformationServer`. 2. Enter the following command to start the removal process.`./uninstall -rsp response_file_name -verbose`
Linux	1. Change to the `_uninstall` directory within the InfoSphere Information Server installation directory. By default, this directory is `/opt/IBM/InformationServer`. 2. Enter the following command to start the removal process.`./uninstall -rsp response_file_name -verbose`
Windows	1. Change to the `_uninstall` directory within the InfoSphere Information Server installation directory. By default, this directory is `C:\IBM\InformationServer`. 2. Enter the following command: `uninstall -rsp response_file_name -verbose`

The `-verbose` option is optional. The option causes the installation program to display log messages to the console and also write them to the log file. If the `-verbose` option is not specified, log messages are written to the log file only.

The program runs a prerequisites check and displays the progress to the console. If a prerequisite check fails, the software removal program stops and displays an error. Correct the error before proceeding with the software removal process.

Consider these conditions while the software removal program is running:

- Leave the command prompt open until the software removal is complete.
- Microsoft Windows: Depending on your security settings, warning windows might display during the removal of the IBM DB2 database system or other components. Ignore the warnings, and click **OK** to continue the removal process.

When the installation is complete, the following message appears in the log file:

```
Installation Engine completed all steps successfully.
```

7. **Windows**: Restart the computer to complete the software removal process.

What to do next

If the software removal process fails, perform the following task to remove the software manually:

Operating system	Task
AIX	Manually remove IBM InfoSphere Information Server (AIX)
Linux	Manually remove IBM InfoSphere Information Server (Linux)
Windows	Manually remove IBM InfoSphere Information Server (Windows)

If the software removal process succeeds, perform the following task:

Operating system	Task
AIX	Complete software removal (AIX)
Linux	Complete software removal (Linux)
Windows	Complete software removal (Windows)

Software removal program encounters an error

During the software removal process, the removal program might encounter an error that requires your action to resolve.

Symptoms

The software removal program encounters an error and an Uninstallation Exception window appears.

Resolving the problem

Complete these steps in the Uninstallation Exception window:

1. Click **Retry failed action** and then click **OK**
2. If the problem recurs, review the last statements in the log file. See Chapter 17, "Log files," on page 407 for more information.
3. If the issue can be resolved, fix the issue, and then click **Retry**.

Chapter 15. Completing software removal

When the software removal process is complete, an Installation Summary for Current Computer page is displayed that contains the product, tiers, and product components that you installed, as well as information on how long the installation program ran and whether the installation was successful.

About this task

If the software removal process was successful, the following message is written to the log file:

```
Installation Engine completed all steps successfully.
```

Completing software removal (Windows)

After you remove IBM InfoSphere Information Server software from Microsoft Windows computers, perform the following procedure on all affected computers to complete the software removal.

Removing leftover files from the services tier (Windows)

When you remove suite software, the software removal program removes only the files that were created as part of the installation of InfoSphere Information Server. You must manually remove other files from the services tier computer if you plan to reinstall the suite.

Before you begin

Run the software removal program and ensure that it completed successfully.

Procedure

1. Remove unused databases.
   a. If you created the metadata repository database before you installed InfoSphere Information Server, manually drop the metadata repository database. The default database name is xmeta.
   b. If you created the IBM InfoSphere Information Analyzer analysis database before you installed InfoSphere Information Server, manually drop the database. The default database name is iadb.
2. Delete the application server directory.
   - If you installed IBM WebSphere Application Server Network Deployment by using the InfoSphere Information Server installation program, complete the additional application server removal steps that are listed in the WebSphere Application Server Network Deployment documentation:
     – For Version 8.5.5.1, go to the WebSphere Application Server information center and read “Cleaning your system after uninstalling the product”.
   - If the WebSphere Application Server installation folder is empty, delete the folder. The default directory for the folder is `C:\IBM\WebSphere\AppServer`.
3. If you previously installed older versions of IBM InfoSphere DataStage or InfoSphere Information Server clients, Windows registry keys might remain on the system after you uninstall the clients. If you have uninstalled all InfoSphere

© Copyright IBM Corp. 2005, 2014

DataStage or InfoSphere Information Server software, delete the following registry keys. Depending on the version of software you have installed, only some of the following keys might exist.

```
HKEY_CURRENT_USER\Software\Wow6432Node\Ascential Software
HKEY_CURRENT_USER\Software\Wow6432Node\IBM\InformationServer
HKLM\SOFTWARE\Ascential Software\ASBNode
HKLM\SOFTWARE\Ascential Software\DataStage Client
HKLM\SOFTWARE\Ascential Software\DataStage Multi-Client Manager
HKLM\SOFTWARE\Ascential Software\Documentation Tool
HKLM\SOFTWARE\IBM\InformationServer
```

4. Remove users if they are no longer needed.
   a. Run the following command to delete the users.

      ```
      net user username /delete
      ```

   b. If you used the local operating system user registry, delete the users that you created for InfoSphere Information Server.
5. Delete the groups that contained the deleted users, if you no longer need the groups.
6. If you used an existing installation of IBM WebSphere Application Server Network Deployment, remove the profile that you created for the suite.

   Note: If the suite installation program installed the application server, the profile is removed when you remove the suite from your computer.
   a. Click **Start** > **All Programs** > **IBM WebSphere** > **Application Server** > **Profiles** > ***profileName*** > **Stop the server**. *profileName* is the name of the profile to delete.
   b. Open a command prompt.
   c. Change to the *WASHome*\bin directory. *WASHome* is the directory where the application server is installed. The default directory is C:\IBM\WebSphere\AppServer.
   d. Run the following command:

      ```
      manageprofiles.bat -delete -profileName profile
      ```

      where *profile* is the name of the profile to be deleted.
   e. Delete the directory for the profile. The directory is in the *WASHome*\profiles*profileName* directory.
7. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.
8. Restart the computer.

Removing leftover files from the engine tier (Windows)

When you remove suite software, the software removal program removes only the files that were created as part of the installation of InfoSphere Information Server. You must manually remove other files from the engine tier computer if you plan to reinstall the suite.

Before you begin

Run the software removal program and ensure that it completed successfully.

Procedure

1. Remove unused databases.

a. If you created the metadata repository database before you installed InfoSphere Information Server, manually drop the metadata repository database. The default database name is xmeta.
b. If you created the IBM InfoSphere Information Analyzer analysis database before you installed InfoSphere Information Server, manually drop the database. The default database name is iadb.

2. Delete the application server directory. The default directory for the folder is `C:\IBM\WebSphere\AppServer`.
3. If you previously installed older versions of IBM InfoSphere DataStage or InfoSphere Information Server clients, Windows registry keys might remain on the system after you uninstall the clients. If you have uninstalled all InfoSphere DataStage or InfoSphere Information Server software, delete the following registry keys. Depending on the version of software you have installed, only some of the following keys might exist.

```
HKEY_CURRENT_USER\Software\Wow6432Node\Ascential Software
HKEY_CURRENT_USER\Software\Wow6432Node\IBM\InformationServer
HKLM\SOFTWARE\Ascential Software\ASBNode
HKLM\SOFTWARE\Ascential Software\DataStage Client
HKLM\SOFTWARE\Ascential Software\DataStage Multi-Client Manager
HKLM\SOFTWARE\Ascential Software\Documentation Tool
HKLM\SOFTWARE\IBM\InformationServer
```

4. Run the following command to delete users if they are no longer needed.

```
net user username /delete
```

5. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.
6. Restart the computer.

Removing leftover files from the repository tier (Windows)

When you remove suite software, the software removal program removes only the files that were created as part of the installation of InfoSphere Information Server. You must manually remove other files from the repository tier computer if you plan to reinstall the suite.

Before you begin

Run the software removal program and ensure that it completed successfully.

Procedure

1. Remove unused databases.
   a. If you created the metadata repository database before you installed InfoSphere Information Server, manually drop the metadata repository database. The default database name is xmeta.
   b. If you created the IBM InfoSphere Information Analyzer analysis database before you installed InfoSphere Information Server, manually drop the database. The default database name is iadb.
2. Delete the application server directory. The default directory for the folder is `C:\IBM\WebSphere\AppServer`.
3. Remove users if they are no longer needed.
   a. Run the following command to delete the users.

```
net user username /delete
```

   b. If you used the DB2 database system as the metadata repository, delete the following accounts:

- The DB2 administration server user. Delete this user only if you are also removing DB2. The default user is db2admin.
- The owner of the metadata repository. The default user is xmeta.
- The owner of the IBM InfoSphere Information Analyzer analysis database. The default user is iauser.

4. If you used the InfoSphere Information Server installation program to install the IBMDB2 database system, remove the following directories:

   DB2 installation directory
   The default directory is `C:\IBM\sqllib`.

   DB2 installation drive
   The directory is *x*`:\DB2`, where *x* is the drive where you installed the DB2 database system.

   DB2 users directory
   The default directories are `%ALLUSERSPROFILE%\Application Data\IBM\DB2` and `%ALLUSERSPROFILE%\Application Data\IBM\DB2History`.

5. If you previously installed older versions of IBM InfoSphere DataStage or InfoSphere Information Server clients, Windows registry keys might remain on the system after you uninstall the clients. If you have uninstalled all InfoSphere DataStage or InfoSphere Information Server software, delete the following registry keys. Depending on the version of software you have installed, only some of the following keys might exist.

   ```
   HKEY_CURRENT_USER\Software\Wow6432Node\Ascential Software
   HKEY_CURRENT_USER\Software\Wow6432Node\IBM\InformationServer
   HKLM\SOFTWARE\Ascential Software\ASBNode
   HKLM\SOFTWARE\Ascential Software\DataStage Client
   HKLM\SOFTWARE\Ascential Software\DataStage Multi-Client Manager
   HKLM\SOFTWARE\Ascential Software\Documentation Tool
   HKLM\SOFTWARE\IBM\InformationServer
   ```

6. Delete the groups that contained the deleted users, if you no longer need the groups.
7. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.
8. Restart the computer.

Completing software removal (AIX Linux)

After you remove IBM InfoSphere Information Server software from computers, perform the following steps to delete any remaining files and complete the software removal.

Removing leftover files from the services tier (AIX Linux)

When you remove suite software, the software removal program removes only the files that were created as part of the installation of InfoSphere Information Server. You must manually remove other files from the services tier computer if you plan to reinstall the suite.

Before you begin

Run the software removal program and ensure that it completed successfully.

Procedure

1. Remove unused databases.

a. If you created the metadata repository database before you installed InfoSphere Information Server, manually drop the metadata repository database. The default database name is xmeta.
b. If you created the IBM InfoSphere Information Analyzer analysis database before you installed InfoSphere Information Server, manually drop the database. The default database name is iadb.

2. Delete the application server directory.
   - If you installed IBM WebSphere Application Server Network Deployment by using the InfoSphere Information Server installation program, complete the additional application server removal steps that are listed in the IBM WebSphere Application Server Network Deployment documentation:
     – For Version 8.5.5.1, go to the WebSphere Application Server information center and read "Cleaning your system after uninstalling the product".
   - If the WebSphere Application Server installation folder is empty, delete the folder. The default directory for the folder is `/opt/IBM/WebSphere/AppServer`.
3. If you used the local operating system user registry, delete the users that you created for InfoSphere Information Server.
4. If you used an existing installation of IBM WebSphere Application Server, remove the profile that you created for the suite.

   Note: If the suite installation program installed the application server, the profile is removed when you remove the suite from your computer.
   a. From a shell prompt, change to the *WASHome*`/bin` directory, where *WASHome* is the directory where the application server is installed. The default directory is `/opt/IBM/WebSphere/AppServer`.
   b. Run the following command to delete the application serer profile:
      ```
      ./manageprofiles.sh -delete -profileName profile
      ```
      where *profile* is the name of the profile to be deleted.
   c. Delete the profile in the *WASHome*`/profiles` directory.
5. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.

Removing leftover files from the engine tier (AIX Linux)

When you remove suite software, the software removal program removes only the files that were created as part of the installation of InfoSphere Information Server. You must manually remove other files from the engine tier computer if you plan to reinstall the suite.

Before you begin

Run the software removal program and ensure that it completed successfully.

Procedure

1. Remove unused databases.
   a. If you created the metadata repository database before you installed InfoSphere Information Server, manually drop the metadata repository database. The default database name is xmeta.
   b. If you created the IBM InfoSphere Information Analyzer analysis database before you installed InfoSphere Information Server, manually drop the database. The default database name is iadb.
2. Delete the directory where you installed InfoSphere Information Server.

Operating System	Directory
AIX	/usr/IBM/InformationServer
Linux	/opt/IBM/InformationServer

3. If you no longer need the dsadm account, delete this user and the related group.
4. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.

Removing leftover files from the repository tier (AIX Linux)

When you remove suite software, the software removal program removes only the files that were created as part of the installation of InfoSphere Information Server. You must manually remove other files from the repository tier computer if you plan to reinstall the suite.

Before you begin

Run the software removal program and ensure that it completed successfully.

Procedure

1. Remove unused databases.
   a. If you created the metadata repository database before you installed InfoSphere Information Server, manually drop the metadata repository database. The default database name is xmeta.
   b. If you created the IBM InfoSphere Information Analyzer analysis database before you installed InfoSphere Information Server, manually drop the database. The default database name is iadb.
2. Delete the directory where you installed InfoSphere Information Server.

Operating System	Directory
AIX	/usr/IBM/InformationServer
Linux	/opt/IBM/InformationServer

3. If you installed the DB2 database system by using the InfoSphere Information Server installation program, delete the following directories:

   DB2 installation directory
   The default directory is /opt/IBM/db2/*version*.

   DB2 instance owner directory
   Remove the /~*db2inst1*/sqllib directory, where ~*db2inst1* is home directory of the instance owner. For example, the default home directory of the instance owner is typically /home/db2inst1.
4. If you used the DB2 database system as the metadata repository and it no longer needs a DB2 user account, delete the DB2 user accounts.
   a. Run the **id** command to identify the related groups for the following DB2 user accounts: db2inst1 and db2fenc1.

   ```
   bash$ id db2inst1
   uid=210(db2inst1) gid=202(db2iadm1) groups=1(staff)
   ```

   b. Run the **userdel** command to delete the user accounts.
   c. Run the **groupdel** command to delete the related groups.

5. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.

Chapter 16. Removing IBM InfoSphere Information Server software manually

If a previous installation or removal failed, you might have to manually remove the suite. You remove the software manually only after trying to use the software removal program.

Manually removing IBM InfoSphere Information Server (Windows)

If the software removal program fails to uninstall the suite, manually uninstall the suite.

Before you begin

Before you manually remove software, try to remove the suite on each computer in your installation by using the software removal program. If the software removal program fails to operate properly, follow this procedure to remove the entire InfoSphere Information Server instance. You can also use this procedure to delete any files that the software removal program does not automatically remove.

Attention: Removing an InfoSphere Information Server installation also deletes any projects, jobs, models, and other user-created items that are associated with that installation. Before you remove the software, ensure that you have backed up all computers in the installation. See "Backing up and restoring IBM InfoSphere Information Server" in the *IBM InfoSphere Information Server Administration Guide* for more information on how to back up your current IBM InfoSphere Information Server installation.

All scheduled tasks must be shut down, and IBM InfoSphere Information Services Director applications must be disabled. See "Shutting down tasks and applications before removing software (service tier and engine tier)" on page 370.

About this task

You cannot manually remove only a single tier or product module. A manual uninstallation must remove the entire instance. If multiple client versions are installed, you cannot remove only the latest version. You must remove all client versions.

Manually removing InfoSphere Information Server from the services tier (Windows)

If you cannot remove suite software by using the software removal program, then you must manually remove files and directories from the services tier computer to ensure that the product is completely removed. Completely removing this files and directories is important if you plan to reinstall the suite.

Procedure

1. Back up the computer.
2. Open an elevated command window:
   a. Click **Start** > **Search**.
   b. In the Search Results box, enter `cmd` and press Ctrl+Shift+Enter.

© Copyright IBM Corp. 2005, 2014

c. Click **Continue** to accept the prompt asking for your permission to continue.

3. If you previously installed older versions of IBM InfoSphere DataStage or InfoSphere Information Server clients, Windows registry keys might remain on the system after you uninstall the clients. If you have uninstalled all InfoSphere DataStage or InfoSphere Information Server software, delete the following registry keys: Depending on the version of software you have installed, only some of the following keys might exist.

```
HKEY_CURRENT_USER\Software\Wow6432Node\Ascential Software
HKEY_CURRENT_USER\Software\Wow6432Node\IBM\InformationServer
HKLM\SOFTWARE\Ascential Software\ASBNode
HKLM\SOFTWARE\Ascential Software\DataStage Client
HKLM\SOFTWARE\Ascential Software\DataStage Multi-Client Manager
HKLM\SOFTWARE\Ascential Software\Documentation Tool
HKLM\SOFTWARE\IBM\InformationServer
```

4. Delete the application server directory.
   - If you installed IBM WebSphere Application Server Network Deployment by using the InfoSphere Information Server installation program, complete the additional application server removal steps that are listed in the IBM WebSphere Application Server Network Deployment documentation:
     – For Version 8.5.5.1, go to the WebSphere Application Server information center and read "Cleaning your system after uninstalling the product".
   - If the WebSphere Application Server installation folder is empty, delete the folder. The default directory for the folder is `C:\IBM\WebSphere\AppServer`.
5. If you used an existing installation of WebSphere Application Server Network Deployment, remove the profile that you created for the suite.

   Note: If the suite installation program installed the application server, the profile is removed when you remove the suite from your computer.

   a. Click **Start** > **All Programs** > **IBM WebSphere** > **Application Server** > **Profiles** > ***profileName*** > **Stop the server**. *profileName* is the name of the profile to delete.
   b. Open a command prompt.
   c. Change to the *WASHome*`\bin` directory. *WASHome* is the directory where the application server is installed. The default directory is `C:\IBM\WebSphere\AppServer`.
   d. Run the following command:

      manageprofiles.bat -delete -profileName *profile*

      where *profile* is the name of the profile to be deleted.
   e. Delete the directory for the profile. The directory is in the *WASHome*`\profiles\`*profileName* directory.
6. Delete the InfoSphere Information Server installation directory. The default directory is `C:\IBM\InformationServer\`.
7. If only the most recent version of the clients are installed on this computer, back up your registry and then delete the `\HKEY_LOCAL_MACHINE\SOFTWARE\Ascential Software` and `\HKEY_LOCAL_MACHINE\SOFTWARE\IBM\InformationServer` registry keys.
8. Delete the following temporary directories that contain log files for the installation program.

```
%TEMP%\ibm_is_logs
%TEMP%\ISInstall*
```

Note: If you chose a different directory for the temporary files, delete that directory instead of the default directories.

9. Delete the Windows Start menu shortcuts from the following directories:

```
C:\Documents and Settings\All Users\Start Menu\Programs\
  IBM InfoSphere Information Server\
C:\ProgramData\Microsoft\Windows\Start Menu
%ALLUSERSPROFILE%\Start Menu\Programs\IBM InfoSphere Information Server\
```

10. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.
11. Restart each computer after you remove the suite. Microsoft Windows might not be able to remove some files and services until the computer is restarted.

    Attention: Do not perform a new product installation before restarting each computer. If you reinstall the product before restarting your computer, the new installation might fail.

Manually removing InfoSphere Information Server from the engine tier (Windows)

If you cannot remove suite software by using the software removal program, then you must manually remove files and directories from the engine tier computer to ensure that the product is completely removed. Completely removing these files and directories is important if you plan to reinstall the suite.

Procedure

1. Back up the computer.
2. Open an elevated command window:
   a. Click **Start** > **Search**.
   b. In the Search Results box, enter `cmd` and press Ctrl+Shift+Enter.
   c. Click **Continue** to accept the prompt asking for your permission to continue.
3. If you previously installed older versions of IBM InfoSphere DataStage or InfoSphere Information Server clients, Windows registry keys might remain on the system after you uninstall the clients. If you have uninstalled all InfoSphere DataStage or InfoSphere Information Server software, delete the following registry keys: Depending on the version of software you have installed, only some of the following keys might exist.

```
HKEY_CURRENT_USER\Software\Wow6432Node\Ascential Software
HKEY_CURRENT_USER\Software\Wow6432Node\IBM\InformationServer
HKLM\SOFTWARE\Ascential Software\ASBNode
HKLM\SOFTWARE\Ascential Software\DataStage Client
HKLM\SOFTWARE\Ascential Software\DataStage Multi-Client Manager
HKLM\SOFTWARE\Ascential Software\Documentation Tool
HKLM\SOFTWARE\IBM\InformationServer
```

4. Run the following commands to stop InfoSphere Information Server services:

```
sc stop ASBAgent
sc stop DSEngine
sc stop dsrpc
sc stop dstelnet
```

5. Run the following commands to remove InfoSphere Information Server services:

```
sc delete ASBAgent
sc delete DSEngine
sc delete dsrpc
sc delete dstelnet
```

6. The installation program uses a particular version of the MKS Toolkit on the client and engine tiers. If you did not first install your own supported version of MKS Toolkit, use Add or Remove Programs to remove the MKS Toolkit that was installed by the InfoSphere Information Server installation program.
7. Delete the InfoSphere Information Server installation directory. The default directory is `C:\IBM\InformationServer\`.
8. If only the most recent version of the clients are installed on this computer, back up your registry and then delete the `\HKEY_LOCAL_MACHINE\SOFTWARE\Ascential Software` and `\HKEY_LOCAL_MACHINE\SOFTWARE\IBM\InformationServer` registry keys.
9. Delete the following temporary directories that contain log files for the installation program.

   ```
   %TEMP%\ibm_is_logs
   %TEMP%\ISInstall*
   ```

   Note: If you chose a different directory for the temporary files, delete that directory instead of the default directories.
10. Delete the Windows Start menu shortcuts from the following directories:

    ```
    C:\Documents and Settings\All Users\Start Menu\Programs\
      IBM InfoSphere Information Server\
    C:\ProgramData\Microsoft\Windows\Start Menu
    %ALLUSERSPROFILE%\Start Menu\Programs\IBM InfoSphere Information Server\
    ```
11. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.
12. Restart each computer after you remove the suite. Microsoft Windows might not be able to remove some files and services until the computer is restarted.

    Attention: Do not perform a new product installation before restarting each computer. If you reinstall the product before restarting your computer, the new installation might fail.

Manually removing InfoSphere Information Server from the repository tier (Windows)

If you cannot remove suite software by using the software removal program, then you must manually remove files and directories from the repository tier computer to ensure that the product is completely removed. Completely removing these files and directories is important if you plan to reinstall the suite.

Procedure

1. Back up the computer.
2. Open an elevated command window:
   a. Click **Start** > **Search**.
   b. In the Search Results box, enter `cmd` and press Ctrl+Shift+Enter.
   c. Click **Continue** to accept the prompt asking for your permission to continue.
3. If you previously installed older versions of IBM InfoSphere DataStage or InfoSphere Information Server clients, Windows registry keys might remain on the system after you uninstall the clients. If you have uninstalled all InfoSphere DataStage or InfoSphere Information Server software, delete the following registry keys: Depending on the version of software you have installed, only some of the following keys might exist.

```
HKEY_CURRENT_USER\Software\Wow6432Node\Ascential Software
HKEY_CURRENT_USER\Software\Wow6432Node\IBM\InformationServer
HKLM\SOFTWARE\Ascential Software\ASBNode
HKLM\SOFTWARE\Ascential Software\DataStage Client
HKLM\SOFTWARE\Ascential Software\DataStage Multi-Client Manager
HKLM\SOFTWARE\Ascential Software\Documentation Tool
HKLM\SOFTWARE\IBM\InformationServer
```

4. If you used the DB2 database system as the metadata repository, delete the following accounts:
   - The DB2 administration server user. Delete this user only if you are also removing DB2. The default user is db2admin.
   - The owner of the metadata repository. The default user is xmeta.
   - The owner of the IBM InfoSphere Information Analyzer analysis database. The default user is iauser.
5. If you used the InfoSphere Information Server installation program to install the IBMDB2 database system:
   a. Remove the DB2 database system by using the Add or Remove Programs utility in the Windows Control Panel. If DB2 cannot be removed by using the Add or Remove Programs utility, remove the DB2 database system by running the command `db2unins.bat /f` in the `bin` directory of your DB2 database system installation location.
   b. Remove the GSK7 entries from the PATH environment variable.
   c. Delete the following directories:

      DB2 installation directory
      The default directory is `C:\IBM\sqllib`.

      DB2 installation drive
      The directory is *x*`:\DB2`, where *x* is the drive where you installed the DB2 database system.

      DB2 users directory
      The default directories are `%ALLUSERSPROFILE%\Application Data\IBM\DB2` and `%ALLUSERSPROFILE%\Application Data\IBM\DB2History`.
6. Delete the InfoSphere Information Server installation directory. The default directory is `C:\IBM\InformationServer\`.
7. If only the most recent version of the clients are installed on this computer, back up your registry and then delete the `\HKEY_LOCAL_MACHINE\SOFTWARE\Ascential Software` and `\HKEY_LOCAL_MACHINE\SOFTWARE\IBM\InformationServer` registry keys.
8. Delete the following temporary directories that contain log files for the installation program.

   ```
   %TEMP%\ibm_is_logs
   %TEMP%\ISInstall*
   ```

   Note: If you chose a different directory for the temporary files, delete that directory instead of the default directories.
9. Delete the Windows Start menu shortcuts from the following directories:

   ```
   C:\Documents and Settings\All Users\Start Menu\Programs\
     IBM InfoSphere Information Server\
   C:\ProgramData\Microsoft\Windows\Start Menu
   %ALLUSERSPROFILE%\Start Menu\Programs\IBM InfoSphere Information Server\
   ```
10. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.

11. Restart each computer after you remove the suite. Microsoft Windows might not be able to remove some files and services until the computer is restarted.

    Attention: Do not perform a new product installation before restarting each computer. If you reinstall the product before restarting your computer, the new installation might fail.

Manually removing InfoSphere Information Server from the client tier (Windows)

If you cannot remove suite software by using the software removal program, then you must manually remove files and directories from the client tier computer to ensure that the product is completely removed. Completely removing these files and directories is important if you plan to reinstall the suite.

Procedure

1. Back up the computer.
2. Open an elevated command window:
   a. Click **Start** > **Search**.
   b. In the Search Results box, enter `cmd` and press Ctrl+Shift+Enter.
   c. Click **Continue** to accept the prompt asking for your permission to continue.
3. If you previously installed older versions of IBM InfoSphere DataStage or InfoSphere Information Server clients, Windows registry keys might remain on the system after you uninstall the clients. If you have uninstalled all InfoSphere DataStage or InfoSphere Information Server software, delete the following registry keys: Depending on the version of software you have installed, only some of the following keys might exist.

   ```
   HKEY_CURRENT_USER\Software\Wow6432Node\Ascential Software
   HKEY_CURRENT_USER\Software\Wow6432Node\IBM\InformationServer
   HKLM\SOFTWARE\Ascential Software\ASBNode
   HKLM\SOFTWARE\Ascential Software\DataStage Client
   HKLM\SOFTWARE\Ascential Software\DataStage Multi-Client Manager
   HKLM\SOFTWARE\Ascential Software\Documentation Tool
   HKLM\SOFTWARE\IBM\InformationServer
   ```

4. The installation program uses a particular version of the MKS Toolkit on the client and engine tiers. If you did not first install your own supported version of MKS Toolkit, use Add or Remove Programs to remove the MKS Toolkit that was installed by the InfoSphere Information Server installation program.
5. Remove the following desktop shortcuts:
   - `Administrator Client.lnk`
   - `Designer Client.lnk`
   - `Console for IBM Information Server.lnk`
   - `Web Console for IBM Information Server.lnk`
6. If the Multi-Client Manager was installed, run the following command to remove the Multi-Client Manager service entry:

   ```
   "C:\IBM\InformationServer\MCM\ClientSwitcherService.exe" /uninstall
   ```

   Note: If multiple client versions are installed on the computer, all clients are removed. To use a previous client version, you must reinstall the client.
7. Delete the InfoSphere Information Server installation directory. The default directory is `C:\IBM\InformationServer\`.

8. If only the most recent version of the clients are installed on this computer, back up your registry and then delete the `\HKEY_LOCAL_MACHINE\SOFTWARE\Ascential Software` and `\HKEY_LOCAL_MACHINE\SOFTWARE\IBM\InformationServer` registry keys.
9. Delete the following temporary directories that contain log files for the installation program.

   ```
   %TEMP%\ibm_is_logs
   %TEMP%\ISInstall*
   ```

   Note: If you chose a different directory for the temporary files, delete that directory instead of the default directories.
10. Delete the Windows Start menu shortcuts from the following directories:

    ```
    C:\Documents and Settings\All Users\Start Menu\Programs\
      IBM InfoSphere Information Server\
    C:\ProgramData\Microsoft\Windows\Start Menu
    %ALLUSERSPROFILE%\Start Menu\Programs\IBM InfoSphere Information Server\
    ```
11. Remove any remaining InfoSphere Information Server users and groups if you are not using them for other purposes.
12. Restart each computer after you remove the suite. Microsoft Windows might not be able to remove some files and services until the computer is restarted.

    Attention: Do not perform a new product installation before restarting each computer. If you reinstall the product before restarting your computer, the new installation might fail.

Manually removing IBM InfoSphere Information Server (AIX Linux)

If a previous installation or removal failed, you might have to manually remove the suite. You remove the software manually only after trying to use the software removal program.

Before you begin

Before you manually remove software, try to remove the suite on each computer in your installation by using the software removal program. If the software removal program fails to operate properly, follow this procedure to remove the entire InfoSphere Information Server instance. You can also use this procedure to delete any files that the software removal program does not automatically remove.

Attention: Removing an InfoSphere Information Server installation also deletes any projects, jobs, models, and other user-created items that are associated with that installation. Before you remove the software, ensure that you have backed up all computers in the installation. See "Backing up and restoring IBM InfoSphere Information Server" in the *IBM InfoSphere Information Server Administration Guide* for more information on how to back up your current IBM InfoSphere Information Server installation.

All scheduled tasks must be shut down, and IBM InfoSphere Information Services Director applications must be disabled. See "Shutting down tasks and applications before removing software (service tier and engine tier)" on page 370.

About this task

You cannot manually remove a single tier or product module only. A manual removal must remove the entire instance.

Manually removing InfoSphere Information Server from the services tier (AIX Linux)

If you cannot remove suite software by using the software removal program, then you must manually remove files and directories from the services tier computer to ensure that the product is completely removed. Completely removing this files and directories is important if you plan to reinstall the suite.

Procedure

1. Back up the computer.
2. Log in as the root user and complete the following steps to stop and remove services:
   a. Change to the InfoSphere Information Server installation directory (default is `/opt/IBM/InformationServer`).
   b. Stop IBM WebSphere Application Server.
   c. Run the following command to determine which startup entries to remove. The `MetadataServer.sh` files are started automatically by symbolic links in the `rc` directory. **AIX Linux**: The directory is `/etc/rc.d`.

      ```
      find /etc -name "*" -print | xargs grep -i installation_dir
      ```

      where *installation_dir* is the directory where you installed InfoSphere Information Server. The default directory is `/opt/IBM/InformationServer`.
   d. Remove all entries that contain `99ISF`.
3. Delete the application server directory.
   - If you installed IBM WebSphere Application Server Network Deployment by using the InfoSphere Information Server installation program, complete the additional application server removal steps that are listed in the WebSphere Application Server Network Deployment documentation:
     – For Version 8.5.5.1, go to the WebSphere Application Server information center and read “Cleaning your system after uninstalling the product”.
   - If you pre-installed WebSphere Application Server Network Deployment manually, instead of reinstalling the product, delete the profile and then recreate it.
4. Delete the InfoSphere Information Server installation directory. The default directory is `/opt/IBM/InformationServer`.
5. Delete the temporary log files for the installation program by running the following command:

   ```
   rm -rf /tmp/ibm_is_logs
   ```

   Note: If you chose a different directory for the temporary files, delete that directory instead of the default directories.
6. Delete the InfoSphere DataStage installation location file by running the following command:

   ```
   rm -rf /.dshome
   ```
7. Delete the InfoSphere DataStage log files by running the following command:

Operating system	Command
AIX	`rm -rf /ds_logs`
Linux	`rm -rf /root/ds_logs`

8. Remove any InfoSphere Information Server users and groups if you are not using them for other purposes.

9. Restart the computer to ensure that all processes have been terminated.

Manually removing InfoSphere Information Server from the engine tier (AIX Linux)

If you cannot remove suite software by using the software removal program, then you must manually remove files and directories from the engine tier computer to ensure that the product is completely removed. Completely removing this files and directories is important if you plan to reinstall the suite.

Procedure

1. Back up the computer.
2. Log in as the root user and complete the following steps to stop and remove services:
   a. Change directories to the installation directory for the InfoSphere Information Server engine. The default location is `/opt/IBM/InformationServer/Server/DSEngine`.
   b. Run the following command to source the `dsenv` file to configure the required environment variables:
      ```
      . ./dsenv
      ```
   c. Run the following command to stop the engine services:
      ```
      ./bin/uv -admin -stop
      ```
   d. Change to the InfoSphere Information Server installation directory. The default directory is `/opt/IBM/InformationServer`.
   e. Run the following command to stop the agents:
      ```
      ./ASBNode/bin/NodeAgents.sh stop
      ```
   f. Run the following command to determine which startup entries to remove. The `NodeAgents.sh` file is started automatically by symbolic links in the `rc` directory. **AIX Linux**: The directory is `/etc/rc.d`.
      ```
      find /etc -name "*" -print | xargs grep -i installation_dir
      ```
      installation_dir is the directory where you installed InfoSphere Information Server. The default directory is `/opt/IBM/InformationServer`.
   g. Remove all entries that contain `99ISF`, `DSEngine`, and `ISFAgent`.
   h. Open the `/etc/services` file, and search for and remove the entry that starts with `dsrpc` entry. The following example shows what this entry might look like:
      ```
      dsrpc 31538/tcp # RPCdaemon DSEngine@/opt/IBM/InformationServer/Server/
         DSEngine
      ```
3. Delete the InfoSphere Information Server installation directory. The default directory is `/opt/IBM/InformationServer`.
4. Delete the temporary log files for the installation program by running the following command:
   ```
   rm -rf /tmp/ibm_is_logs
   ```
   Note: If you chose a different directory for the temporary files, delete that directory instead of the default directories.
5. Delete the InfoSphere DataStage installation location file by running the following command:
   ```
   rm -rf /.dshome
   ```
6. Delete the InfoSphere DataStage log files by running the following command:

Operating system	Command
AIX	`rm -rf /ds_logs`
Linux	`rm -rf /root/ds_logs`

7. Remove any InfoSphere Information Server users and groups if you are not using them for other purposes.
8. Restart the computer to ensure that all processes have been terminated.

Manually removing InfoSphere Information Server from the repository tier (AIX Linux)

If you cannot remove suite software by using the software removal program, then you must manually remove files and directories from the repository tier computer to ensure that the product is completely removed. Completely removing this files and directories is important if you plan to reinstall the suite.

Procedure

1. Back up the computer.
2. If you installed IBMDB2 database system as part of the InfoSphere Information Server installation, remove the DB2 database system from the computer on which you installed the repository tier. Removing the DB2 database system automatically removes the metadata repository database and the analysis database.
   a. Log in as the DB2 instance owner. This account is db2inst1 by default.
   b. Drop the databases by running the following commands:

      ```
      db2 drop database xmeta
      db2 drop database iadb
      ```

      This example uses the default database names: `xmeta` and `iadb`. If you used different names, use your database names instead.
   c. If you installed the DB2 database system as part of the installation of InfoSphere Information Server and it is the only DB2 installation, remove the DB2 database system. See the DB2 Information Center for more information on removing the DB2 database system.
   d. Use the following commands to remove any remaining files:

      ```
      rm -rf /tmp/db2*
      rm -rf /var/db2
      rm -rf /var/opt/db2
      rm -rf ~db2inst1/db2inst1
      rm -rf ~db2inst1/sqllib
      rm -rf /opt/IBM/db2
      ```

      This example uses the default paths and user accounts. If you used different paths and user accounts, use your paths and user accounts instead.
   e. Remove any entries in the `/etc/services` file that start with `DB2_db2inst1` or `db2c_db2inst1` where db2inst1 is the name of the instance owner that you used during the suite installation.
   f. If the `/etc/inittab` file exists and this is the only DB2 database system installation, open the file and search for and remove the `fmc` entry. The following example shows what this entry might look like:

      ```
      #DB2 Fault Monitor Coordinator
      fmc:number:respawn:/opt/IBM/db2/V10.5/bin/db2fmcd
      ```

3. If you used an existing installation of DB2 database system, use the DB2 command prompt to drop the metadata repository database and the analysis database. The default database names are xmeta and iadb.
4. Delete the InfoSphere Information Server installation directory. The default directory is `/opt/IBM/InformationServer`.
5. Delete the temporary log files for the installation program by running the following command:

   ```
   rm -rf /tmp/ibm_is_logs
   ```

   Note: If you chose a different directory for the temporary files, delete that directory instead of the default directories.
6. Delete the InfoSphere DataStage log files by running the following command:

Operating system	Command
AIX	`rm -rf /ds_logs`
Linux	`rm -rf /root/ds_logs`

7. Remove any InfoSphere Information Server users and groups if you are not using them for other purposes.
8. Restart the computer to ensure that all processes have been terminated.

Part 6. Troubleshooting InfoSphere Information Server software installation

Resolve installation problems with this installation troubleshooting information. You can find more troubleshooting information in the information centers and customer support technotes for various product modules and components.

© Copyright IBM Corp. 2005, 2014

Chapter 17. Log files

To troubleshoot issues with IBM InfoSphere Information Server, you can consult log files and temporary files.

Installation log files

Log files are created during the installation and uninstallation process. You can use the log files to troubleshoot installation errors.

During the installation and uninstallation process, the log files are saved to the following directory:

Table 105. Default directory where log files are saved during the installation and uninstallation process

Operating system	Directory
AIX Linux	`/tmp/ibm_is_logs`
Windows	`%TEMP%\ibm_is_logs`

Note: If you redirected the installation program's temporary files to a different directory, the installation log files are in that location, instead of in the default temporary directory that is specified here.

After a successful installation, a .zip file of the log files that are created during the installation process is saved to the IBM InfoSphere Information Server installation directory. This .zip file is named `isdump-`*`os-timestamp`*`.zip`. After a failed installation, the `ISInstall-`*`date-timestamp`*`.log` log file is retained in the temporary directory.

After a successful uninstallation, the `ISInstall-`*`date-timestamp`*`.log` log file is retained in the logs directory under the installation directory. After a failed uninstallation, the `ISinstall-`*`date-timestamp`*`.log` log file is retained in the `ibm-is-logs` directory within the temporary directory.

Note: Do not delete the installation log files or the log file directory during the installation process. When installation is complete, save the log files for troubleshooting and verification purposes.

Application server system log files

The application server log files contain information that you can use to monitor the startup of the application server and diagnose errors.

WebSphere Application Server logs

These files are located in the following directories on each node in your WebSphere Application Server installation:

Operating system	Directory
AIX	*`path`*`/profiles/`*`profile`*`/logs/`*`server1`*

© Copyright IBM Corp. 2005, 2014

Operating system	Directory
Linux	*path*/profiles/*profile*/logs/*server1*
Windows	*path*\profiles*profile*\logs*server1*

path
The WebSphere Application Server installation path. By default, *path* is one of the following paths:

Operating system	Directory
AIX	opt/IBM/WebSphere/AppServer
Linux	opt/IBM/WebSphere/AppServer
Windows	C:\IBM\WebSphere\AppServer

profile
The profile name where IBM InfoSphere Information Server is installed. For a stand-alone installation, the default value is InfoSphere. For a clustered installation, the default value for a custom profile is Custom*xx*, where *xx* is the number of the profile.

server1
The name of the application server. For a stand-alone installation, the default value is server1. For cluster installations, there might be multiple application server directories under the custom profile. The typical value is server*x*, where *x* is the number of the application server instance. For a Deployment Manager profile, the default value is dmgr. For a node agent under the custom profile, the default value is nodeagent.

The following log files can be found in these directories:

SystemOut.log
WebSphere Application Server messages to STDOUT are redirected to this file.

SystemErr.log
WebSphere Application Server messages to STDERR are redirected to this file.

For more information about WebSphere Application Server log files, see the WebSphere Application Server documentation:

- For Version 8.5.5.1, go to the WebSphere Application Server information center and read "Using basic or traditional message logs to troubleshoot applications"

WebSphere Application Server Liberty Profile logs

The log files are located in the following directory:

Operating system	Directory
AIX	*IS_install_path*/wlp/usr/servers/iis/logs
Linux	*IS_install_path*/wlp/usr/servers/iis/logs
Windows	*IS_install_path*\wlp\usr\servers\iis\logs

The following log files can be found in the directory:

messages.log
Current log.

messages_*YY.MM.DD_HH.MM.SS***.log**
Previous logs.

InfoSphere Information Server console client log file

Client applications report debug, error, and information messages in the IBM InfoSphere Information Server console client log file.

The location of the InfoSphere Information Server console client log file is on the client computer in the `%USERPROFILE%\AppData\Local\IBM Information Server console\isc.log` file.

Chapter 18. General installation problems

These problems and resolutions are related to general issues that might occur across the product suite.

ExecuteScriptAction failed: File Names in the following package are too long

During a Microsoft Windows installation, if you receive a message that reads "File Names in the following package are too long" and the log entry indicates a CDIN5106I message that the error occurred during the IAServices install action, you can continue the installation and fix the problem afterward.

Symptoms

During installation, you receive the following messages:

- com.ibm.is.install.exception.ActionException: CDIIN5125E: ExecuteScriptAction failed: the error output matched the standard error regular expression RegularExpression (_pattern = SubstitutableValue: Token=".+", Value=".+").: Error Encountered the following service exception: [File Names in the following package are too long. Trim the depth of folder or the folder name length: [.]]

  CDIIN5106I: Details: InstallAction (_id = ia-report-template-baseline-analysis, _installUnitModelId = IAServices)
- com.ibm.is.install.exception.ActionException: CDIIN5123E: ExecuteScriptAction failed: the process returned error code -1.:

  CDIIN5106I: Details: InstallAction (_id = https.deploy.predefined.general.rules, _installUnitModelId = IAServices)

Causes

The path to the package in the temporary folder is too long for Microsoft Windows to execute the script.

Environment

Microsoft Windows

Resolving the problem

1. Continue the installation.
2. After installation, temporarily set the TEMP environment variable to a shorter path, such as `C:\temp` and deploy the analysis report template (the ⇒ symbol indicates a line continuation):

   ```
   set TEMP=C:\temp
   cd IS_install_path\ASBServer\bin
   ReportingAdmin.bat -url services_url -user is_admin_user ⇒
   -password is_admin_password -deploy -templates -file ⇒
   is_suite_path\payloads\InformationAnalyzer\IAServices\resources⇒
   \templates\reports\ia-report-template-baseline-analysis.jar
   ```

   Where:

   - *IS_install_path* is the location where IBM InfoSphere Information Server is installed.
   - *is_suite_path* is the extracted location of the IBM InfoSphere Information Server installation image, its `is-suite` folder.

© Copyright IBM Corp. 2005, 2014

Recovering from a UVREGEN installation exception

When you run an installation after a previous attempt that has not been cleaned up properly, you might encounter the error described here. You should not encounter this error if you rebooted the engine computer prior to your current installation process.

Symptoms

During installation, you encounter the following error message:

```
com.ibm.is.install.exception.InstallException: CDIIN5063E:
ExecuteDataStgeProcessAction UVREGEN failed with exit code 255. Error
output is "INSTANCETAG tunable cannot be "dcd" or "ace". uvregen: error
during regen, process aborted".
```

In addition, if you run the following command, you see output:

```
ipcs -a | grep -0xtag
```

where *tag* is the ITAG for the engine tier instance. For the IBM InfoSphere Data Click release, the tag is ade.

Causes

A recent removal or manual cleanup process did not successfully complete before you attempted the reinstallation.

Resolving the problem

1. If you have not already done so, do not close the current installation process. The problem can be resolved and you will be able to resume installation when complete. Whether you closed the installation program or not, complete the following steps:
   a. Open a command shell and become the root user, either directly with the **su** command or through the **sudo** command if the current user is in the sudoers list.
      ```
      sudo su - root
      ```
   b. Find the shmids of the shared memory segments that are still open.
      ```
      ipcs -m | grep 0xtag | awk '{print $2}'
      ```
   c. For each shared memory segment that is listed, run the remove command:
      ```
      ipcrm -m shmid
      ```
   d. Find the semids of the semaphore arrays:
      ```
      ipcs -s | grep 0xtag | awk '{print $2}'
      ```
   e. For each semaphore array, run the remove command:
      ```
      ipcrm -s semid
      ```
   f. Ensure that all shared memory segments and semaphore arrays have been removed. You should see no output when you run the following command:
      ```
      ipcs -a | grep 0xtag
      ```
   g. Run the following action, which had failed during installation:
      ```
      cd IS_install_path/Server/DSEngine
      bin/uv -admin -settag -itag tag -force -noautostart
      ```
2. Resume the installation process. Or, if you did close the installation program, restart it by using the response file so that you do not need to re-enter your values:

a. The only values that the response file does not save are the passwords. Edit the response file and add the values for the passwords.
b. If you closed the command shell from which you started the installation program originally, open a new shell as root and prepare the shell environment:

```
export IBM_JAVA_OPTIONS=
umask 022
```

c. Run the installation program with the response file:

```
setup -rsp response_file_name -uiMode graphical
```

"User credentials are not valid" error appears unexpectedly during installation

If you receive a "user credentials are not valid" error when you create a new operating system user account from the installation program, specify the correct password or remove the user account.

Symptoms

When you specify a new operating system user account from the IBM InfoSphere Information Server installation program, you receive a "user credentials are not valid" error and you cannot continue the installation. See “Setting up administrator and owner accounts” on page 107 for more information about operating system user accounts.

Causes

When you specify a new operating system user account from the installation program, the program checks whether the user account already exists. If the account exists, the program checks whether the password is correct. If the account does not exist, the program immediately creates the user account on the computer. If you exit the program before you complete the installation, the user account remains on the computer. If you subsequently restart the installation and specify the same user name but a different password, you receive a "user credentials are not valid" error.

Environment

All supported operating systems.

Resolving the problem

Use operating system utilities to determine whether the user account exists on the computer. If necessary, change the password. Then continue the installation.

IWAV0003E warning in installation log file

IWAV0003E warnings in the installation log file can be ignored.

Symptoms

A warning similar to the following text is written to the installation log file during the IBM WebSphere Application Server enterprise archive (EAR) deployment phase:

```
[exec] +++ Warning +++: Mon Sep 18 19:36:18 BST 2010 IWAV0003E
Could not reflect methods for com.ascential.xmeta.shared.
ILocalTransactionalBusiness because one of the methods references
a type that could not be loaded. Exception: java.lang.
NoClassDefFoundError: com.ascential.xmeta.exception.ServiceException
```

Environment
All environments.

Resolving the problem
This error can be ignored. The EAR installs correctly.

This warning is issued because the classes inside the EAR are not on the classpath for the EAR installation process. After the EAR is installed, the classes are resolved correctly.

Warnings occur during installation of DB2

While installing IBM InfoSphere Information Server, you might see a message indicating that warnings occurred during the installation of IBM DB2 Database for Linux, UNIX, and Windows, or "Open File - Security Warning" messages might be displayed.

Symptoms
During the DB2 installation, warnings or other messages are displayed.

Resolving the problem
If the warnings are “Open File - Security Warning” messages, ignore the messages. Otherwise, do these steps:

1. Look in the installation log file for a message that is similar to the following line:

   ```
   WARNING: com.ascential.acs.installer.asb.InstallDB2Action
   WRepositoryInstaller65507: Warnings occurred installing DB2.
   For details see the log file:
   C:\IBM\InformationServer\logs\db2_install.log
   ```

2. View the log file that is specified in the warning message. You can ignore warnings such as SQL1362W or SQL20189W, which warn that the updates do not take effect until the next database startup. These warnings are not important because the installation program stops and restarts DB2 during the installation process.

For more information about specific warning messages in the DB2 log file, refer to the DB2 documentation:

- DB2 V10.5: publib.boulder.ibm.com/infocenter/db2luw/v10r5/

Installation problems with VMware ESX

If you experience problems when installing IBM InfoSphere Information Server on a VMware ESX virtual computer, make sure that your system is not overloaded.

Symptoms
When you attempt the installation, one or more of the following problems occurs:

- The system performs slowly, and the installation takes a much longer time than expected.
- IBM DB2 installation fails because not all ports are available.
- IBM WebSphere Application Server installation fails because not all ports are available.

Causes
The target system is overloaded.

Environment

- Linux (SUSE, Red Hat)
- Microsoft Windows

Resolving the problem

Reduce the load on the system, and try the installation again.

"The publisher could not be verified" message

You can ignore messages such as “The publisher could not be verified. Are you sure you want to run this software?” during IBM InfoSphere Information Server installation.

Symptoms

A window is displayed in the browser window. “The message in the window reads The publisher could not be verified. Are you sure you want to run this software?”

Causes

The installation program and related utilities are not signed.

Environment

Microsoft Windows 7 and Microsoft Windows Vista

Resolving the problem

Ignore the message and continue the installation. This message does not indicate a security risk when it is displayed during InfoSphere Information Server installation.

"There is no script engine for file extension '.js'" message appears during installation

During a Microsoft Windows installation, if you receive a message that reads “There is no script engine for file extension '.js'”, repair or reinstall the Windows Script Engine.

Symptoms

During installation, you receive a message that reads “There is no script engine for file extension '.js'.”

Causes

The Microsoft Windows Script Engine is not installed or is not operational.

Environment

Microsoft Windows

Diagnosing the problem

Open a command prompt window and enter the following command:

```
cscript myfile1.js
```

where *myfile1* is the name of a new file.

If you receive the following message, there is a problem with the Windows Script Engine installation:

```
Input Error: There is no script engine for file extension ".js".
```

Resolving the problem

Repair or reinstall the Windows Script Engine. Refer to Microsoft documentation for more information. There are known problems with the Windows Script Engine that cause this error.

Installation process hangs when a web browser session is closed

If an exception occurs after a web browser or browser tab has been closed, or if the browser session has timed out, the installation process will be stuck waiting for user response.

Symptoms

The installation process is stuck waiting for a response.

Causes

When the web browser window or tab that is running an installation process is closed during the installation session, if the installation runs into exception, the installation process gets stuck waiting for a user response.

Note: To avoid this problem, keep the browser or browser tab open throughout the installation session.

Environment

Graphical mode installation

Resolving the problem

1. Uninstall your current IBM InfoSphere Information Server installation by running the following commands:

   ```
   cd /opt/IBM/InformationServer/_uninstall
   ./uninstall -verbose
   ```

2. Complete the manual cleanup:

Operating system	Procedure
AIX	Completing software removal (AIX)
Linux	Completing software removal (Linux)
Windows	Completing software removal (Windows)

3. Start a new InfoSphere Information Server installation by running the following commands:

   ```
   cd <image>/is-suite
   ./setup -rsp response.txt -force -verbose
   ```

Fix pack version errors during IBM InfoSphere Information Server installation

When you pre-install IBM WebSphere Application Server separately before you install InfoSphere Information Server, the installation program halts with an error if fix packs later than 7.0.0.11 are installed. The problem also occurs if you have installed Fix Pack 7.0.0.11 and iFix IFPM20596. You can ignore these errors and continue the installation.

Symptoms

The Prerequisites Check page displays messages like the following lines. The messages also appear in the installation log file:

```
WARNING: com.ibm.is.install.check.WebSphereVersionChecker No iFixes found
in the WebSphere Application Server installation.

INFO: FAILED: CDIPR2118I: Ensure that the WebSphere Application Server
version has the fix packs 6.1.0.29:IFPK96868, 6.1.0.29:MultiOS-IFPK99266,
6.1.0.29:IFPK99560, 7.0.0.11:IFPK96868, 7.0.0.11:IFPM20596,
7.0.0.11:IFPM21207. Detected fix packs for version 7.0.0.11. Resolution:
Follow WebSphere Application Server Fix Pack installation instructions to
install the Java SDK and Fix Packs to the required level.
```

Causes

The Java software development kit (SDK) that is installed with these fix packs is a higher version than the installation program expects. However, InfoSphere Information Server supports this version.

Environment

All environments

Resolving the problem

Restart the installation program with the `-force` option on the command line. This option enables a check box within the Prerequisites Check page that allows you to continue even if the prerequisites check returns warnings or errors.

Use the `-force` option with caution. Be sure to resolve any other warnings or errors that appear during the prerequisites check. Continuing an installation with unresolved warnings or errors might result in a non-operational InfoSphere Information Server instance.

Chapter 19. Authentication problems

These problems and resolutions are related to authentication.

Confirming user privileges

If the installation fails, confirm that the users who install IBM InfoSphere Information Server have the necessary privileges.

Symptoms

The installation fails to complete successfully.

Environment

All supported operating systems.

Resolving the problem

Log in to the computer with the proper privileges.

Operating system	Procedure
AIX	Log in as the root user.
Linux	Log in as the root user.
Windows	Log in to the target computer as a local administrator or as a local or domain user who is directly assigned to the local Administrators group.

© Copyright IBM Corp. 2005, 2014

Chapter 20. WebSphere Application Server Network Deployment server problems

These problems and resolutions are related to the services tier and IBM WebSphere Application Server Network Deployment.

Resolving IBM WebSphere Application Server installation problems

Use this troubleshooting information when you install IBM WebSphere Application Server with IBM InfoSphere Information Server.

Symptoms

When you install WebSphere Application Server with IBM InfoSphere Information Server, the WebSphere Application Server installation fails.

Environment

All supported environments.

Diagnosing the problem

If the WebSphere Application Server installation fails, review the following log files:

- WebSphere Application Server log files. For more information, refer to the following topics in the WebSphere Application Server documentation:
  - Version 8.5.5.1: Verifying the installation
- InfoSphere Information Server installation log file:

Operating system	Directory
AIX	`/opt/IBM/InformationServer/logs/WebSphereInstall.log`
Linux	`/opt/IBM/InformationServer/logs/WebSphereInstall.log`
Windows	`C:\IBM\InformationServer\logs\WebSphereInstall.log`

Search for following error:

```
INSTCONFFAILED: Cannot create profile:
The profile does not exist
```

Resolving the problem

Any of the following issues can cause problems with the WebSphere Application Server installation process:

Microsoft Windows

The `DEBUG` variable is set in your environment variables.

Verify that the **`DEBUG`** variable is not set in your environment variables. If the **`DEBUG`** variable is set:

1. Remove the **`DEBUG`** variable from your environment variables.
2. Restart the computer.
3. Reinstall InfoSphere Information Server.

© Copyright IBM Corp. 2005, 2014

4. After the installation is complete, reset the **DEBUG** variable.

The installation user cannot write to the Microsoft Windows directory.

The installation user must have write permissions to the Windows directory. This directory is the directory specified by the `%WinDir%` environment variable. By default, the directory is `C:\WINDOWS`.

In Microsoft Windows 2008, this problem might occur if you run as the true administrator or if User Account Control (UAC) is enabled.

AIX

The `/etc/filesystems` file contains a duplicate entry.

If the `/etc/filesystems` file contains a duplicate entry, the WebSphere Application Server installation reports 0 MB of available disk space. Edit the `/etc/filesystems` file, and comment out the duplicate entry.

Installation fails when the IBM WebSphere Application Server Network Deployment port is changed after installation

After you install IBM InfoSphere Information Server, if you change the WebSphere Application Server HTTPS port (WC_defaulthost_secure) by using the IBM WebSphere Application Server administrative console, future installations, product additions, and updates fail.

Symptoms

A "Connection refused" error message is displayed on the user interface or in the log file.

Causes

The original installation settings are stored in several configuration files are used for subsequent installation activities. If the WebSphere Application Server port number changes after installation, the installation settings in these files are incorrect.

Resolving the problem

Follow the procedure in Changing the WebSphere Application Server Network Deployment HTTP port to change the HTTP port in the WebSphere Application Server administrative console and in the configuration files.

Resolving connection problems

An installation might fail because the Domain Name System (DNS) server is unable to resolve system names or because the application server is not running.

Symptoms

IBM InfoSphere Information Server installation fails when the Domain Name System (DNS) is unable to resolve system names. The installation program returns an error message that is similar to the following message:

```
CDIIN2107E: Unable to validate the IBM InfoSphere Information Server user name name
and password for host hostname at port port_number. The
returned message is:
java.net.UnknownHostException: hostname.

Confirm that 1) WebSphere Application Server is running; 2) you entered the correct
user name and password; 3) the port number is correct.
```

Causes

The communication between the installation computer and the application server on the services tier is not established. The following issues might cause this problem:

- The connection to the application server uses the short host name.
- The connection to the application server uses a host name that does not match the name that the computer had when the application server was installed.
- The short host name does not map to the same IP address as the long host name.
- The application server is not running.

Diagnosing the problem

When the services tier is installed on *hostname.domain*, use the following ping commands to test the connection to the short host name and to the long host name:

```
ping hostname
ping hostname.domain
```

Resolving the problem

1. Ensure that the application server is running.
2. Ensure that the `etc/hosts` file is correctly configured with the short host name and long host name of the services tier computer.
3. Ensure that all necessary ports are open among the client tier, engine tier, and services tier computers.
4. Ensure that the host name that is used to access the services tier computer is the same name that the computer had when the application server was installed.
5. Ensure that localhost and the computer host name are not both mapped to the loopback IP address. The IPv4 version of this address is 127.0.0.1 and the IPv6 version is ::1/128. Map only localhost to IP address 127.0.0.1. The following example shows an `etc/hosts` file where these entries are specified correctly:

   ```
   127.0.0.1 | ::1/128 localhost.localdomain localhost
   xx.xx.xx.xx | xx:xx::xx:xx machinelonghostname
   machineshorthostname
   ```

   xx.xx.xx.xx | *xx:xx::xx:xx* is the unique IP address of the computer.
6. If the engine tier or services tier was installed on a computer that has incorrectly configured host names, reinstall IBM InfoSphere Information Server.

Unable to start IBM WebSphere Application Server after the user name changes or expires

If the user name or password that you created for WebSphere Application Server or for the metadata repository changes or expires, the application server does not start.

Symptom

The application server does not start.

Cause

The WebSphere Application Server user name or password changed or expired, or the metadata repository user name or password changed or expired. Starting and

stopping WebSphere Application Server requires user authentication. These start and stop commands fail when user authentication is unsuccessful because the user names or passwords changed or expired.

Resolving the problem

The IBM InfoSphere Information Server installation directory is shown as *install_home* in the following commands.

Note: The AppServerAdmin commands do not validate user names or passwords.

AIX Linux

Follow these steps to resolve the problem:

1. Log in to the services tier computer as root.
2. To change the password for the WebSphere Application Server administrator user, enter the following command:

   ```
   install_home/ASBServer/bin/AppServerAdmin.sh -was -user user_name
        -password new_password
   ```

3. To change the password for the metadata repository user, do the following steps:
   a. Confirm that you can use the new password to connect to the database.
   b. Enter the following command to register the new password information:

      ```
      install_home/ASBServer/bin/AppServerAdmin.sh -db -user user_name
           -password new_password
      ```

   c. Restart all WebSphere Application Server processes. See the *IBM InfoSphere Information Server Administration Guide*.

Microsoft Windows

Follow these steps to resolve the problem:

1. Log in to the services tier computer as a user with administrator privileges.
2. To change the password for the WebSphere Application Server administrator user, enter the following command:

   ```
   install_home\ASBServer\bin\AppServerAdmin.bat -was -user user_name
        -password new_password
   ```

3. To change the password for the metadata repository user, do the following steps:
   a. Confirm that you can use the new password to connect to the database.
   b. Enter the following command to register the new password information:

      ```
      install_home\ASBServer\bin\AppServerAdmin.bat -db -user user_name
           -password new_password
      ```

   c. Restart all WebSphere Application Server processes. See the *IBM InfoSphere Information Server Administration Guide*.

For details on adding a new WebSphere Application Server user when using the internal registry, see "Adding an IBM WebSphere Application Server administrator user to an internal user registry" on page 425.

Adding an IBM WebSphere Application Server administrator user to an internal user registry

If you do not know the WebSphere Application Server administrator user and password, you can create a user in the internal user registry.

Symptom

You are using the internal user registry and you do not know the WebSphere Application Server administrator user and password.

Cause

The password changed, or the user was deleted.

Resolving the problem

The IBM InfoSphere Information Server installation directory is shown as *install_home* in the following commands.

AIX Linux

Follow these steps to resolve the problem:

1. Log in to the services tier computer as root.
2. Enter the following command to add a user with administrator privileges to the InfoSphere Information Server internal user registry:

   ```
   install_home/ASBServer/bin/DirectoryAdmin.sh -user -admin
        -userid new_user_name -password new_password
   ```

3. Enter the following **AppServerAdmin** command to update the user name or password within the WebSphere Application Server configuration:

   ```
   /opt/IBM/InformationServer/ASBServer/bin/AppServerAdmin.sh
   -was -user name -password password
   ```

Microsoft Windows

Follow these steps to resolve the problem:

1. Log in to the services tier computer as a user with administrator privileges.
2. Enter the following command to add a user with administration privileges to the InfoSphere Information Server internal user registry:

   ```
   install_home\ASBServer\bin\DirectoryAdmin.bat -user -admin
        -userid new_user_name -password new_password
   ```

3. Enter the following **AppServerAdmin** command to update the user name or password within the WebSphere Application Server configuration:

   ```
   C:\IBM\InformationServer\ASBServer\bin\AppServerAdmin.bat
   -was -user name -password password
   ```

Repairing the WebSphere Application Server Network Deployment registry after switching to LDAP

If you have switched from an internal registry to a Lightweight Directory Access Protocol (LDAP) user registry without using the documented procedure, "Switching to an LDAP user registry," you might be unable to stop IBM WebSphere Application Server processes.

Symptoms

You can not stop WebSphere Application Server Network Deployment when using either the **MetadataServer.sh** stop command or the **stopServer.sh** command.

Environment
Optional. Describe any environmental details that are not already in the title or short description.

Resolving the problem
Repair the WebSphere Application Server registry:

1. Kill the WebSphere Application Server processes.
2. Use **sudo vi** to edit the following settings in the file, `/opt/IBM/WebSphere/AppServer/profiles/InfoSphere/config/cells/hostNameNode01Cell/security.xml`:
   a. Change the first instance of **enabled**="**true**" to "**false**"
   b. Change the first instance of **appEnabled**="**true**" to "**false**"
3. Start WebSphere Application Server.
4. Log into the WebSphere Application Server administrative console by using wasadmin.

   Note: Because of your modifications to `security.xml` in step 2, the console does not prompt you for a password.
5. In **Security** > **Global Security**, click **Configure** and enter the following information:
   a. Set **Primary administrator user name** to wasadmin.
   b. Select **Server identity that is stored in the repository** and set the user ID to wasadmin and type the password.
6. Click **Apply** and then **OK**. This returns you to the Global Security page.
7. Select the check boxes: **Enable administrative security** and **Enable application security**
8. Clear the **Java 2 security** check box.
9. Click **Apply**. A message appears on the top of the form. Click the **Save** link inside the message.
10. Log out of the WebSphere Application Server administrative console.
11. Test that you can now stop WebSphere Application Server smoothly:
    a. Stop WebSphere Application Server.

       Note: WebSphere Application Server should be able to stop without errors, but you must still do the following steps to ensure the that the problem is fixed.
    b. Start WebSphere Application Server.
    c. Stop WebSphere Application Server. It should be able to stop without errors.

Services tier performance

If the services tier is slow, ensure that you have not turned on IBM WebSphere Application Server tracing.

Symptoms
The services tier is running slowly. The default setting for IBM WebSphere Application Server tracing is off.

Resolving the problem

To avoid the problem, ensure that any changes made to IBM WebSphere Application Server are restored to the default settings before installing and using IBM InfoSphere Information Server.

To resolve the problem, ensure that you have not turned on the IBM WebSphere Application Server tracing.

Chapter 21. IBM InfoSphere Information Server reporting problems

These problems and resolutions are related to reporting.

Cannot view report result from InfoSphere Information Server console or InfoSphere Information Server Web console

In a cluster environment, if a "Page Not Found" error message is displayed while you are trying to view a report result, a security process might be running. You can edit the `httpd.conf` configuration file to allow encoded slash characters, and then restart the Web server.

Symptoms

After you run a process in a cluster environment, a "Page Not Found" error message is displayed, and you cannot view the report results.

Causes

As a security measure, the web server encodes some URLs. Each slash in a URL is encoded as special character that the `httpd.conf` configuration file reads and then rejects by default.

Environment

This advice applies to cluster environments that run on IBM HTTP Server or Apache Web Server.

Resolving the problem

1. Open the `httpd.conf` configuration file. The file is in the configuration directory: *install-home*`\IBM\HTTPServer\conf\httpd.conf`.
2. Set the **`AllowEncodedSlashes`** value to `ON`.
3. Save this file, and restart the web server. If the web server is installed as a managed server, use theIBM WebSphere Application Server administrative console to restart it. Otherwise, log into the computer where the web server is installed, and restart the web server.

Enabling and disabling trace logging of the reporting engine

When you encounter a problem or an error that requires attention by IBM Software Support, you can set up trace logging to report detailed internal processes. You can enable trace logging in the reporting engine to generate logs that help IBM Software Support diagnose errors.

Symptoms

Not applicable.

Resolving the problem

1. Open a command prompt or a terminal window on the computer.
2. Enter the following command to enable logging in the reporting engine:

© Copyright IBM Corp. 2005, 2014

Operating system	Command
AIX	*installdir*/ASBServer/bin/ ReportingAdmin.sh -user *username* -password *password* -el
Linux	*installdir*/ASBServer/bin/ ReportingAdmin.sh -user *username* -password *password* -el
Windows	*installdir*\ASBServer\bin\ ReportingAdmin.bat -user *username* -password *password* -el

If you specify the `-authfile` parameter, the `-user` and `-password` parameters are optional. Use the `-authfile` parameter to specify the location of your credentials file.

3. Rerun the process that resulted in the error that you received earlier.
4. After the error occurs, create a `.zip` file that contains the log files. The log files are located in the following directories:
   - If IBM WebSphere Application Server clustering is enabled, the log files are created in the reporting workspace on each server in the cluster. On each server, the reporting workspace is located in the temp folder that is defined by the operating system. The path has the following format:

     tempfolder\informationServer*servername*\engine\JREPORT

     For example: `C:\WINDOWS\Temp\informationServer\myservernode1\engine\JREPORT`
   - If clustering is not enabled, the log files are located in the following directory:

Operating system	Directory
AIX	*installdir*/bin/InformationServer/ ASBServer/apps/Reporting/engine/JReport/ logs
Linux	*installdir*/bin/InformationServer/ ASBServer/apps/Reporting/engine/JReport/ logs
Windows	*installdir*\bin\InformationServer\ ASBServer\apps\Reporting\engine\JReport\ logs

5. Send this `.zip` file to IBM Software Support.
6. When you finish collecting detailed logs, roll back the trace level settings to the original state. At the command prompt, enter the following command to disable logging:

Operating system	Command
AIX	*installdir*/ASBServer/bin/ ReportingAdmin.sh -user *username* -password *password* -dl
Linux	*installdir*/ASBServer/bin/ ReportingAdmin.sh -user *username* -password *password* -dl

Operating system	Command
Windows	*installdir*\ASBServer\bin\ ReportingAdmin.bat -user *username* -password *password* -dl

Chapter 22. Client and console problems

These problems and resolutions are related to client errors and login problems.

Java Virtual Machine fails to load

During startup, client applications might fail to load the Java Virtual Machine (JVM) if the operating system is unable to allocate a sufficient amount of heap memory.

Symptoms

The client application fails with an error message that is similar to the following message:

```
Internal error - Failed to load the Java Virtual Machine
```

Causes

The JVM is unable to allocate the memory that it has been configured to use. This error typically occurs when too many client processes are running concurrently.

Resolving the problem

Decrease the value that is specified for MaximumHeapSize in the application-specific sections of the following file: `\ASBNode\conf\proxy.xml`. This file is located in the IBM InfoSphere Information Server installation directory.

Continue to decrease this value until clients start normally. If multiple client applications are running, you might need to decrease this value for all clients. Comments in the `proxy.xml` file identify the configuration section for each client application. Common values for MaximumHeapSize include 64, 128, 256, 348, 512, and so on.

After making changes, you must restart the clients to reallocate the memory.

Rich client login failure

If the rich client login fails because the services tier cannot be reached, you can troubleshoot several possible causes.

Symptoms

The client login failure can be identified with the following information:

- The stack trace from the client does not include the message `Trace from server`.
- The stack trace shows that the root cause of the exception is a socket connection error.
- The host name and port number are included in the error message.
- A RegistrationHelper exception is included in the error message.

Stack traces and error messages vary depending on the situation. A connection failure might be identified with an error message about network connection, such as "Could not connect to server" or "port out of range." A connection failure might also be revealed in a stack trace that is similar to the following example:

```
Exception type: javax/security/auth/login/LoginException
Exception message: Could not connect to server [green3] on port [9081].
Exception stack trace:
```

© Copyright IBM Corp. 2005, 2014

```
javax.security.auth.login.LoginException:
  Could not connect to server [green3] on port [9081].
 at com.ascential.acs.security.auth.client.AuthenticationService.getLoginException
  (AuthenticationService.java:978)
 at com.ascential.acs.security.auth.client.AuthenticationService.doLogin
  (AuthenticationService.java:355)
Caused by: com.ascential.acs.registration.client.RegistrationContextManagerException:
  Caught an unexpected exception.
 at com.ascential.acs.registration.client.RegistrationContextManager.setContext
  (RegistrationContextManager.java:76)
 at com.ascential.acs.security.auth.client.AuthenticationService.doLogin
  (AuthenticationService.java:349)
Caused by: com.ascential.acs.registration.client.RegistrationHelperException:
  Caught an unexpected exception.
 at com.ascential.acs.registration.client.RegistrationHelper.getBindingProperties
  (RegistrationHelper.java:694)
 at com.ascential.acs.registration.client.RegistrationHelper.getBindingProperties
  (RegistrationHelper.java:587)
 at com.ascential.acs.registration.client.RegistrationHelper.getBindingConfigProperties
  (RegistrationHelper.java:566)
 at com.ascential.acs.registration.client.RegistrationContextManager.setContext
  (RegistrationContextManager.java:173)
 at com.ascential.acs.registration.client.RegistrationContextManager.setContext
  (RegistrationContextManager.java:73)
 ... 1 more
Caused by: java.net.ConnectException: Connection refused: connect
 at java.net.PlainSocketImpl.socketConnect(Native Method)
 at java.net.PlainSocketImpl.doConnect(PlainSocketImpl.java:352)
 at java.net.PlainSocketImpl.connectToAddress(PlainSocketImpl.java:214)
 at java.net.PlainSocketImpl.connect(PlainSocketImpl.java:201)
 at java.net.SocksSocketImpl.connect(SocksSocketImpl.java:368)
 at java.net.Socket.connect(Socket.java:526)
 at java.net.Socket.connect(Socket.java:476)
 at sun.net.NetworkClient.doConnect(NetworkClient.java:175)
 at sun.net.www.http.HttpClient.openServer(HttpClient.java:407)
 at sun.net.www.http.HttpClient.openServer(HttpClient.java:542)
 at sun.net.www.http.HttpClient.<init>(HttpClient.java:246)
 at sun.net.www.http.HttpClient.New(HttpClient.java:319)
 at sun.net.www.http.HttpClient.New(HttpClient.java:336)
 at sun.net.www.protocol.http.HttpURLConnection.getNewHttpClient
  (HttpURLConnection.java:858)
 at sun.net.www.protocol.http.HttpURLConnection.plainConnect
  (HttpURLConnection.java:799)
 at sun.net.www.protocol.http.HttpURLConnection.connect
  (HttpURLConnection.java:724)
 at sun.net.www.protocol.http.HttpURLConnection.getInputStream
  (HttpURLConnection.java:1047)
 at com.ascential.acs.registration.client.RegistrationHelper.getBindingProperties
  (RegistrationHelper.java:677)
 ... 5 more
```

Resolving the problem

To resolve the problem, apply one of the following solutions, depending on the cause of the failure:

Cause	Solution
The client computer host file has an invalid host name or IP address entry.	Update the client computer hosts file so that the server host name can be resolved from client.
The server listening port might be blocked by a firewall.	Ensure that the application server TCP/IP ports are not blocked by the firewall.
The application server is not running.	Ensure that the application server is running.

Cause	Solution
In a cluster environment, you did not specify the host name and port of the front-end HTTP dispatcher.	If IBM WebSphere Application Server clustering is implemented within your installation, direct the client to the host name and port of the front-end HTTP dispatcher. (The default port is 80.) The client login will fail if you specify the Deployment Manager host name and port or if you specify the host name and port of a single cluster member.
You changed a WebSphere Application Server port or modified the host file on the system, but you did not restart WebSphere processes after making the change.	In a stand-alone (non-cluster) environment, restart IBM WebSphere Application Server. In a cluster environment, restart the cluster.

Resolving problems logging in to the WebSphere Application Server administrative console

Use this troubleshooting information when you are running a stand-alone installation of IBM WebSphere Application Server Network Deployment with IBM InfoSphere Information Server and you cannot log in to the WebSphere Application Server administrative console.

Symptoms

You cannot log in to the WebSphere Application Server administrative console of an WebSphere Application Server installation that is running InfoSphere Information Server.

Causes

There are multiple reasons why the login fails.

Environment

This advice applies to all environments that are supported by InfoSphere Information Server.

Diagnosing the problem

Ensure that you are using the default IBM WebSphere Application Server administrator user name and password. If you are not sure of the WebSphere Application Server administrator user name, you can view the WebSphere Application Server `security.xml` file.

Attention: Do not modify this file.

The `security.xml` file is in the following directory:

Operating system	Directory
AIX	`/opt/IBM/WebSphere/AppServer/profiles/`*`profile_name`*`/config/cells/`*`cell`*
Linux	`/opt/IBM/WebSphere/AppServer/profiles/`*`profile_name`*`/config/cells/`*`cell`*
Windows	`\IBM\WebSphere\AppServer\profiles\`*`profile_name`*`\config\cells\`*`cell`*

profile_name is the name of the WebSphere Application Server profile being used (the default name is InfoSphere). *cell* is the name of the WebSphere Application Server cell.

The user name is specified by the serverId attribute in the userRegistries element of the active user registry. Ensure that you are looking at the serverId of the active user registry. To find the active user registry in the `security.xml` file, look for the **`activeUserRegistry`** string in the file. The file contains an entry that is similar to **`activeUserRegistry=`***UserRegistry*, where the *UserRegistry* value depends on your configuration. Search the remainder of the `security.xml` file for this value, which determines the current active user registry. Ensure that you can access the registry. For example, if you are using an LDAP user registry, ensure that the LDAP server is running and that you can connect to it. If you are using the InfoSphere Information Server internal registry, ensure that the database server that is used by InfoSphere Information Server is running and that you can access it.

Resolving the problem

If the user name or password is invalid or unknown:

1. Change the password.
   - If your installation is configured to use the InfoSphere Information Server internal user registry, enter the following **`DirectoryAdmin`** command to change the password:

Operating system	Command
AIX	*install_home*/ASBServer/bin/DirectoryAdmin.sh -user -userid *wasadmin_username* -password *wasadmin_password*
Linux	*install_home*/ASBServer/bin/DirectoryAdmin.sh -user -userid *wasadmin_username* -password *wasadmin_password*
Windows	*install_home*\ASBServer\bin\DirectoryAdmin -user -userid *wasadmin_username* -password *wasadmin_password*

   - If your installation is configured to use an external user registry (such as a local operating system user registry or an LDAP user registry), use operating system or LDAP utilities to change the password.
2. Enter the following **`AppServerAdmin`** command to update the WebSphere Application Server and InfoSphere Information Server configurations:

Operating system	Command
AIX	*install_home*/ASBServer/bin/AppServerAdmin.sh -was -user *username* -password *password*
Linux	*install_home*/ASBServer/bin/AppServerAdmin.sh -was -user *username* -password *password*
Windows	*install_home*\ASBServer\bin\AppServerAdmin -was -user *username* -password *password*

Connection problems after IBM WebSphere Application Server cluster members are removed

You might experience connection problems with the WebSphere Application Server administrative console or IBM InfoSphere DataStage clients if you remove all members of a WebSphere Application Server cluster.

Symptoms

You cannot access the WebSphere Application Server administrative console, or log in to InfoSphere DataStage clients. You recently removed all members of the WebSphere Application Server cluster and then added members again.

Causes

When all cluster members are removed from the cluster, the cluster template is removed. New cluster members are set up with the default cluster template. The default template is not compatible with IBM InfoSphere Information Server.

Environment

All environments.

Resolving the problem

1. Rebuild the cluster as described in "Adding a new cluster member" in the *IBM InfoSphere Information Server Administration Guide*. Do not start the cluster yet.
2. On the computer that hosts the Deployment Manager, run the following command:

Operating system	Command
AIX	*installdir*/ASBServer/bin/ reconfigure_was_cluster.sh -user *wasadmin_user* -password *wasadmin_password*
Linux	*installdir*/ASBServer/bin/ reconfigure_was_cluster.sh -user *wasadmin_user* -password *wasadmin_password*
Windows	*installdir*\ASBServer\bin\ reconfigure_was_cluster.bat -user *wasadmin_user* -password *wasadmin_password*

   wasadmin_user and *wasadmin_password* are the user name and password of the WebSphere Application Server administrator.
3. Start the cluster as described in "Restarting application server processes" in the *IBM InfoSphere Information Server Administration Guide*.

Chapter 23. Removal problems

These problems and resolutions are related to removing IBM InfoSphere Information Server products.

IBM InfoSphere Information Server installation fails repeatedly

If an installation fails repeatedly on a computer that has been used for a previous IBM InfoSphere Information Server installation, verify that the previous installation was completely removed.

Symptoms

Installation fails on a computer that has been used for a previous IBM InfoSphere Information Server installation.

Causes

Previous installations are not completely removed, or the computer was not restarted after the software removal process completed.

Environment

This advice applies to all operating systems.

Resolving the problem

Before you install InfoSphere Information Server, you must completely remove previous installations, including failed installations, of InfoSphere Information Server. You must remove previous versions of InfoSphere Information Server from the installation location where you intend to install the new instance.

You must also remove DB2 and WebSphere from the target installation locations before the InfoSphere Information Server installation program can reinstall DB2 and WebSphere.

To ensure that the previous installation is completely removed, follow the manual software removal steps for the previous installation.

Note: On Microsoft Windows, you can have more than one version of InfoSphere Information Server only if the previously installed versions were client-only installations. You must completely remove any non-client InfoSphere Information Server installations before you install InfoSphere Information Server.

Ignoring runtime error during uninstallation

A known problem causes an unimportant error message to be displayed during the uninstallation process.

Symptoms

During the uninstallation process, the following error message is displayed when the ASBAgent is being removed:

```
This application has requested the Runtime to terminate it in an unusual way
```

Causes

© Copyright IBM Corp. 2005, 2014

This is a known problem with the Visual C++ Runtime .dll file (msvcrt.dll), which is documented in the Microsoft Knowledge Base.

Environment
This advice applies to the Microsoft Windows XP operating system.

Resolving the problem
Ignore this error message. The uninstallation process completes successfully with no errors encountered in the process.

MKS Toolkit errors occur when removing clients

Only one version of the MKS Toolkit can exist on a Microsoft Windows computer at a time. The MKS Toolkit is installed with the IBM InfoSphere Information Server clients. If you install a newer version of the clients, the version of the MKS Toolkit that is packaged with the clients overwrites any older version of the MKS Toolkit on the system, resulting in version incompatibility issues between the MKS Toolkit and older versions of IBM InfoSphere Information Server clients. Version incompatibility might also cause problems when you remove older InfoSphere Information Server clients.

Symptoms
When you try to remove InfoSphere Information Server clients, you might see several possible symptoms that point to version incompatibility issues:

- When you attempt to remove an older InfoSphere Information Server client, error messages are displayed when the software removal program tries to remove the MKS Toolkit. The messages are displayed because the older version of the InfoSphere Information Server client is not compatible with the new version of the MKS Toolkit.

  The following error message is displayed:

  ```
  ERROR: The product bean "installMksClientAction" could not be loaded...
  ```
- At the end of the software removal of the older InfoSphere Information Server client, the following error message is displayed:

  ```
  Could not uninstall assembly: uid=mksinstaller : version = 8.0.0.0: instance = 1
  ```
- After you remove IBM InfoSphere Information Server, Version 8.0.1 clients, the product shortcuts are not removed from the Windows **Start** menu.

Causes
A newer version of the MKS Toolkit was installed and overwrote the previously existing version that was installed by the InfoSphere Information Server client that you are trying to remove. If you install multiple versions of the InfoSphere Information Server client, the only version of the MKS Toolkit that remains on your computer is the version that is packaged with the newest version of the InfoSphere Information Server client that you install.

If you remove the version of the InfoSphere Information Server client that installed the current version of the MKS Toolkit without removing the MKS Toolkit, the remaining clients continue to run. However, if you later try to remove the remaining clients, you might receive error messages when you try to remove the MKS Toolkit because the versions are incompatible.

Environment
This advice applies to all Microsoft Windows operating systems.

Resolving the problem

If you removed all IBM InfoSphere DataStage and InfoSphere Information Server clients but the MKS Toolkit remains, you can remove it by using the add or remove program function that is available with your version of Windows. Reboot the computer after removing the MKS Toolkit.

Part 7. Appendixes

© Copyright IBM Corp. 2005, 2014

Appendix A. Product accessibility

You can get information about the accessibility status of IBM products.

The IBM InfoSphere Information Server product modules and user interfaces are not fully accessible.

For information about the accessibility status of IBM products, see the IBM product accessibility information at http://www.ibm.com/able/product_accessibility/index.html.

Accessible documentation

Accessible documentation for InfoSphere Information Server products is provided in IBM Knowledge Center. IBM Knowledge Center presents the documentation in XHTML 1.0 format, which is viewable in most web browsers. Because IBM Knowledge Center uses XHTML, you can set display preferences in your browser. This also allows you to use screen readers and other assistive technologies to access the documentation.

The documentation that is in IBM Knowledge Center is also provided in PDF files, which are not fully accessible.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

© Copyright IBM Corp. 2005, 2014

Appendix B. Reading command-line syntax

This documentation uses special characters to define the command-line syntax.

The following special characters define the command-line syntax:

[ ] Identifies an optional argument. Arguments that are not enclosed in brackets are required.

... Indicates that you can specify multiple values for the previous argument.

| Indicates mutually exclusive information. You can use the argument to the left of the separator or the argument to the right of the separator. You cannot use both arguments in a single use of the command.

{ } Delimits a set of mutually exclusive arguments when one of the arguments is required. If the arguments are optional, they are enclosed in brackets ([ ]).

Note:

- The maximum number of characters in an argument is 256.
- Enclose argument values that have embedded spaces with either single or double quotation marks.

For example:

wsetsrc*[-S server] [-l label] [-n name] source*

The *source* argument is the only required argument for the **wsetsrc** command. The brackets around the other arguments indicate that these arguments are optional.

wlsac *[-l | -f format] [key...] profile*

In this example, the -l and -f format arguments are mutually exclusive and optional. The *profile* argument is required. The *key* argument is optional. The ellipsis (...) that follows the *key* argument indicates that you can specify multiple key names.

wrb -import *{rule_pack | rule_set}...*

In this example, the rule_pack and rule_set arguments are mutually exclusive, but one of the arguments must be specified. Also, the ellipsis marks (...) indicate that you can specify multiple rule packs or rule sets.

© Copyright IBM Corp. 2005, 2014

Appendix C. Contacting IBM

You can contact IBM for customer support, software services, product information, and general information. You also can provide feedback to IBM about products and documentation.

The following table lists resources for customer support, software services, training, and product and solutions information.

Table 106. IBM resources

Resource	Description and location
IBM Support Portal	You can customize support information by choosing the products and the topics that interest you at www.ibm.com/support/entry/portal/Software/Information_Management/InfoSphere_Information_Server
Software services	You can find information about software, IT, and business consulting services, on the solutions site at www.ibm.com/businesssolutions/
My IBM	You can manage links to IBM Web sites and information that meet your specific technical support needs by creating an account on the My IBM site at www.ibm.com/account/
Training and certification	You can learn about technical training and education services designed for individuals, companies, and public organizations to acquire, maintain, and optimize their IT skills at http://www.ibm.com/training
IBM representatives	You can contact an IBM representative to learn about solutions at www.ibm.com/connect/ibm/us/en/

© Copyright IBM Corp. 2005, 2014

Appendix D. Accessing the product documentation

Documentation is provided in a variety of formats: in the online IBM Knowledge Center, in an optional locally installed information center, and as PDF books. You can access the online or locally installed help directly from the product client interfaces.

IBM Knowledge Center is the best place to find the most up-to-date information for InfoSphere Information Server. IBM Knowledge Center contains help for most of the product interfaces, as well as complete documentation for all the product modules in the suite. You can open IBM Knowledge Center from the installed product or from a web browser.

Accessing IBM Knowledge Center

There are various ways to access the online documentation:

- Click the **Help** link in the upper right of the client interface.
- Press the F1 key. The F1 key typically opens the topic that describes the current context of the client interface.

  Note: The F1 key does not work in web clients.
- Type the address in a web browser, for example, when you are not logged in to the product.

  Enter the following address to access all versions of InfoSphere Information Server documentation:

  `http://www.ibm.com/support/knowledgecenter/SSZJPZ/`

  If you want to access a particular topic, specify the version number with the product identifier, the documentation plug-in name, and the topic path in the URL. For example, the URL for the 11.3 version of this topic is as follows. (The ⇒ symbol indicates a line continuation):

  ```
  http://www.ibm.com/support/knowledgecenter/SSZJPZ_11.3.0/⇒
  com.ibm.swg.im.iis.common.doc/common/accessingiidoc.html
  ```

  Tip:

  The knowledge center has a short URL as well:

  `http://ibm.biz/knowctr`

  To specify a short URL to a specific product page, version, or topic, use a hash character (#) between the short URL and the product identifier. For example, the short URL to all the InfoSphere Information Server documentation is the following URL:

  `http://ibm.biz/knowctr#SSZJPZ/`

  And, the short URL to the topic above to create a slightly shorter URL is the following URL (The ⇒ symbol indicates a line continuation):

  ```
  http://ibm.biz/knowctr#SSZJPZ_11.3.0/com.ibm.swg.im.iis.common.doc/⇒
  common/accessingiidoc.html
  ```

© Copyright IBM Corp. 2005, 2014

Changing help links to refer to locally installed documentation

IBM Knowledge Center contains the most up-to-date version of the documentation. However, you can install a local version of the documentation as an information center and configure your help links to point to it. A local information center is useful if your enterprise does not provide access to the internet.

Use the installation instructions that come with the information center installation package to install it on the computer of your choice. After you install and start the information center, you can use the **iisAdmin** command on the services tier computer to change the documentation location that the product F1 and help links refer to. (The ⇒ symbol indicates a line continuation):

Windows

```
IS_install_path\ASBServer\bin\iisAdmin.bat -set -key ⇒
com.ibm.iis.infocenter.url -value http://<host>:<port>/help/topic/
```

AIX Linux

```
IS_install_path/ASBServer/bin/iisAdmin.sh -set -key ⇒
com.ibm.iis.infocenter.url -value http://<host>:<port>/help/topic/
```

Where `<host>` is the name of the computer where the information center is installed and `<port>` is the port number for the information center. The default port number is 8888. For example, on a computer named `server1.example.com` that uses the default port, the URL value would be `http://server1.example.com:8888/help/topic/`.

Obtaining PDF and hardcopy documentation

- The PDF file books are available online and can be accessed from this support document: https://www.ibm.com/support/docview.wss?uid=swg27008803&wv=1.
- You can also order IBM publications in hardcopy format online or through your local IBM representative. To order publications online, go to the IBM Publications Center at http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

Notices and trademarks

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

Notices

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 character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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.

© Copyright IBM Corp. 2005, 2014

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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © 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.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session or persistent cookies. If a product or component is not listed, that product or component does not use cookies.

Table 107. Use of cookies by InfoSphere Information Server products and components

Product module	Component or feature	Type of cookie that is used	Collect this data	Purpose of data	Disabling the cookies
Any (part of InfoSphere Information Server installation)	InfoSphere Information Server web console	• Session • Persistent	User name	• Session management • Authentication	Cannot be disabled
Any (part of InfoSphere Information Server installation)	InfoSphere Metadata Asset Manager	• Session • Persistent	No personally identifiable information	• Session management • Authentication • Enhanced user usability • Single sign-on configuration	Cannot be disabled

Table 107. Use of cookies by InfoSphere Information Server products and components (continued)

Product module	Component or feature	Type of cookie that is used	Collect this data	Purpose of data	Disabling the cookies
InfoSphere DataStage	Big Data File stage	• Session • Persistent	• User name • Digital signature • Session ID	• Session management • Authentication • Single sign-on configuration	Cannot be disabled
InfoSphere DataStage	XML stage	Session	Internal identifiers	• Session management • Authentication	Cannot be disabled
InfoSphere DataStage	IBM InfoSphere DataStage and QualityStage Operations Console	Session	No personally identifiable information	• Session management • Authentication	Cannot be disabled
InfoSphere Data Click	InfoSphere Information Server web console	• Session • Persistent	User name	• Session management • Authentication	Cannot be disabled
InfoSphere Data Quality Console		Session	No personally identifiable information	• Session management • Authentication • Single sign-on configuration	Cannot be disabled
InfoSphere QualityStage Standardization Rules Designer	InfoSphere Information Server web console	• Session • Persistent	User name	• Session management • Authentication	Cannot be disabled
InfoSphere Information Governance Catalog		• Session • Persistent	• User name • Internal identifiers • State of the tree	• Session management • Authentication • Single sign-on configuration	Cannot be disabled
InfoSphere Information Analyzer	Data Rules stage in the InfoSphere DataStage and QualityStage Designer client	Session	Session ID	Session management	Cannot be disabled

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies:

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows and Windows NT 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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The United States Postal Service owns the following trademarks: CASS, CASS Certified, DPV, LACS^{Link}, ZIP, ZIP + 4, ZIP Code, Post Office, Postal Service, USPS and United States Postal Service. IBM Corporation is a non-exclusive DPV and LACS^{Link} licensee of the United States Postal Service.

Other company, product or service names may be trademarks or service marks of others.

Index

Special characters

/3GB switch
  removing from boot.ini file 335
/etc/magic file
  editing 329
/PAE switch
  adding to boot.ini file 335
.ds files
  data set descriptor files 329
.zip files
  installation log files 407

Numerics

32-bit systems 14
64-bit systems 14

A

accounts, user 107
active metadata repository 287
active-passive configuration
  configuration example (MSCS) 170
  configuration overview (AIX) 167
  configuration overview (Linux) 167
  configuration overview (Windows) 183
  creating users and user groups (Microsoft example) 188
  creating users and user groups (System Automation for MultiPlatforms example) 178
  defining and starting resource group (System Automation for Multiplatforms example) 175
  high availability control scripts 169
  installing high availability software (System Automation for Multiplatforms example) 173
  installing InfoSphere Information Server in cluster (Microsoft example) 188
  installing InfoSphere Information Server in cluster (System Automation for Multiplatforms example) 178
  Microsoft example 1, services and engine 184
  preparing server hardware (Microsoft) 186
  preparing server hardware (MSCS example) 172
  road map for installation 76
  setting up high availability for the engine and services tier (Microsoft example) 188
  setting up network (Microsoft example) 186
  setting up network (MSCS example) 172
active-passive configuration *(continued)*
  setting up resource agent scripts (System Automation for Multiplatforms example) 179
  setting up the cluster (Microsoft example) 186
  setting up virtual IP address (Microsoft example) 187
  setting up virtual IP address (System Automation for Multiplatforms example) 174
  testing (Microsoft example) 189
  testing (System Automation for Multiplatforms example) 181
analysis database
  setting up in existing DB2 instance 114
analysis database, changing 287
analysis databases
  sizing 56
  topologies 26
application server
  changing HTTP port 304
  planning 14
  removing profiles 385, 388
    engine tier (AIX) 389
    engine tier (Linux) 389
    engine tier (Windows) 386
    metadata repository tier (AIX) 390
    metadata repository tier (Linux) 390
    repository tier (Windows) 387
    services tier (AIX) 388
    services tier (Linux) 388
    services tier (Windows) 385
  troubleshooting 422, 423, 425
applying fix packs 359
architecture
  engine 6
  installation 17
  services 4
ASB agents 6
authentication
  configuring PAM 338
automatic client reroute
  configuring for a DB2 database cluster 124
  configuring for HADR setup 137
  for a DB2 clustered or HADR configuration 293

B

backing up
  planning 54
  primary to standby in HADR setup 138
bind
  DB2 package 344
boot.ini file
  adding /PAE switch 335
  removing /3GB switch 335
buildop components
  compilers 322

C

C++ compiler
  GCC for Linux on 64-bit Intel systems 323
  IBM AIX 322
  Microsoft Visual Studio .NET 323
C++ compilers
  installing 322
capacity planning 54
changing
  database high availability 297
Changing configuration 287
character encodings
  supported 164
client
  firewall exceptions 306
client settings
  for InfoSphere Glossary Anywhere 271
client tier
  described 3
cluster members
  resolving connection problems after removing 437
clustering
  application servers 39
clusters
  planning 17
code pages 164
command prompt
  installing from 215
  running software removal program from 380
command-line syntax
  conventions 447
commands
  syntax 447
component installer
  described 6
components
  described 3
compute node
  AIX 309
  Linux 309
  Windows 314
compute nodes 51
  distributing transform libraries to 321
  Windows 313, 315
conductor nodes 51
config.ini file
  InfoSphere Glossary Anywhere 258
configuration file
  InfoSphere Glossary Anywhere 258

© Copyright IBM Corp. 2005, 2014

configuring
InfoSphere Information Server 279
Configuring
repository high availability 297
configuring the browser
removing software 371
connections
maximum number 334
connectivity
planning 96
testing ODBC drivers 345
Connector access services agent 6
customer support
contacting 449

D

data collection and analysis 251
data source names
creating 346
specifying 343
data sources
modifying connection properties 293, 299
data stores, changing 287
database
Match Designer database 143
Database client libraries 337
database clustering
overview 47
database connection properties
modifying 299
QSSRDDB 299
QualityStage Standardization Rules Designer 299
database connection propertiesmetadata repository data source
modifying 293
Database connectivity
configuring 337
database system
determining how to install 13
database.properties file
modifying 290
specifying alternate host for cluster 127, 135
DB2 cluster
Liberty profile configuration 127
specifying alternate database address and port 125
DB2 clustered configuration
configuring automatic client reroute 293
DB2 database clustering
configuration overview 119
configuring automatic client reroute 124
creating a cluster 122
creating the metadata repository database 122
creating users and groups 123
heartbeat 120
installing DB2 database system 122
installing HA software 121
Liberty Profile: specifying alternate database address and port 127
DB2 database clustering *(continued)*
Network Deployment: specifying alternate database address and port 125
overview 47
preparing servers, file system, and storage 120
running installation program 123
updating database properties file 127, 135
verifying 129
DB2 database system
using existing instance for repository tier databases 114
DB2 HADR
Liberty profile configuration 134
specifying alternate database address and port 132
updating database properties file 127, 135
DB2 installation
warnings 414
deployments
installation topologies 17
diagnostic tools 251
directories
required for installation 95
domain connections
troubleshooting 422
DS_CONNECT 345
dsenv file
sample entries 341
DSNs
creating 346
DSODB
changing for high availability 297

E

EAR deployment
warnings during 413
encoded slashes
allowing 429
encodings
supported 164
engine tier
example scenario with isolated engine 59
overview 6
topology example 19
environment variables
parallel engine 325
setting 324
setting for InfoSphere DataStage 341
ESDB
changing for high availability 297
exceptions database
capacity planning 58
changing for high availability 297
exceptions database, changing 287

F

failover 166
firewall
clustering 46
firewall *(continued)*
exceptions for client programs 306
Fix Central 193
fix packs 193
applying 359
front-end dispatcher
modifying 283
front-end load balancer
installing 158
front-end web server
alternative, front-end load balancer 157
configuring 280
installing 157
front-end Web server
clustering 42

G

genkey 216
globalization (NLS)
configuring 340
supported character encodings 164
GLTABSZ parameter 345
graphical mode
configuring the browser 194
removing software
configuring the browser 371
starting software removal program in 376
grids
configuring 50, 53
planning 17

H

HA software
installing 121
HADR configuration
configuring automatic client reroute 293
hardware requirements 94
health checkers 251
heap size
increasing 332
high availability
active-passive configuration example 170
active-passive configuration example 1, services and engine 184
active-passive configuration scripts (System Automation for Multiplatforms example) 179
active-passive configuration testing (Microsoft example) 189
active-passive configuration testing (System Automation for Multiplatforms example) 181
configuration overview 30
DB2 options 47
design considerations 30
failover 166
installing management software (System Automation for Multiplatforms example) 173
required configuration 283

high availability *(continued)*
setting up high availability for the engine and services tiers (Microsoft example) 188
setting up network (Microsoft example) 186
setting up network (MSCS example) 172
high availability clustered configuration
road map for installation 83
high availability control scripts
overview 169
using (System Automation for Multiplatforms example) 179
high availability disaster recovery (HADR)
configuring automatic client reroute 137
configuring primary server 137
configuring standby server 138
creating the database 131
creating users and groups 132
installing DB2 database system 47, 131
preparing servers and storage 130
running installation program 132
setting up metadata repository in 129
starting 140
validation 141
HTTP port 299, 304
HTTPS
example scenario 59

I

IBM DB2 database system
configuring existing instance for metadata repository (Windows) 115, 146
determining how to install 13
preinstallation 113
using existing instance for metadata repository database (AIX) 117
using existing instance for metadata repository database (Linux) 117
individual modules
removing 374
information analysis
planning 56
InfoSphere DataStage
configuring 337
configuring globalization (NLS) 340
testing the installation 252
InfoSphere DataStage plug-ins
monitoring installation 250
InfoSphere Glossary Anywhere 268
installing and configuring a client 263, 266
preparing for installation 257
InfoSphere Information Analyzer
agent 6
database, analysis 56
InfoSphere Information Governance Catalog
configuring 349
InfoSphere Information Governance Catalog for Eclipse
installation 271
InfoSphere Information Server
in Windows 393
client tier 398
engine tier 395
repository tier 396
services tier 393
installing 361, 365
on AIX 399
engine tier 401
repository tier 402
services tier 400
on Linux 399
engine tier 401
repository tier 402
services tier 400
preparing to install 197
removing 369, 374, 385, 388, 393, 395, 396, 398, 399, 400, 401, 402
engine tier (Windows) 386
repository tier (Windows) 387
services tier (Windows) 385
shutting down tasks and applications before removing 370
InfoSphere Information Server engine
tuning 345
InfoSphere Information Server installation program
running 160
InfoSphere Information Server Web console
described 4
InfoSphere Information Services Director
agent 6
testing installation 254
InfoSphere Metadata Interchange Bridges
planning 3
InfoSphere Metadata Server
troubleshooting startup 423, 425
InfoSphere QualityStage
configuring 337
configuring globalization (NLS) 340
testing the installation 252
installation
CDs 193
kernel parameter configuration 101
log files 407
monitoring overview 246
order to install tiers 193
planning 3, 13
process overview 3, 13
sample checklist: basic installation with no high availability 63
sample checklist: isolated engine tier and parallel engine configuration 69
sample road map: high availability active-passive configuration 76
sample road map: high availability clustered configuration 83
scenario with basic client-server topology 59
scenario with clustered metadata repository tier and services tier 62
scenario with isolated engine tier 59
installation *(continued)*
scenario with parallel engine configuration 59
scenario with two-server active-passive high availability topology 60
topologies 17
troubleshooting permissions 103
user privileges 419
verification 251
installation logs
monitoring 246
installation program 201
running for DB2 database clustered configuration 123
running for HADR configuration 132
installation records
updating 295
installation software
obtaining 193
installing
C++ compilers 322
command prompt (silent) 215
completing the installation process 214
fix packs 359
InfoSphere Information Server 361, 365
patches 359
preparation steps 197
running the installation wizard 211
saving a response file 211
specifying options 206
investigation reports
WebSphere Application Server 347
IP sprayer topology 46, 158
ISALite 251
ISInstall*.log 247, 250
ITAG installations 6, 21
IWAV0003E 413

J

Java 1.7 162
Java Virtual Machine (JVM)
troubleshooting 433

K

kernel parameters
modifying 101
setting 326

L

language 163
legal notices 453
Library path environment
configuring 337
library path environment variables
InfoSphere DataStage libraries 341
load balancers
introduction 44
locale 163
log files
installation 407

log files *(continued)*
log files
application server 407
WebSphere Application Server 407
log.txt 248
logical tiers 3
login
troubleshooting 433
LOGTO 345
lssam 129

M

maintenance
planning 54
massively parallel processing
configurations 51
configuring compute node (AIX) 309
configuring compute node (Linux) 309
configuring compute nodes (Windows) 313
copying components (Windows) 314
Match Designer database 287
configuring 346
database requirements 56, 143
locations 26
planning 143
MAXRLOCK parameter 345
maxuproc parameter
setting 326
memory
setting de-commit threshold 332
setting for parallel jobs 327
messages.log 249
metadata
importing with Metadata Asset Manager 195
installing importing applications 195
metadata repository 287
changing host name and port 290
Metadata repository
registering schemas within 295
metadata repository database
configuring in existing DB2 instance by using scripts (Windows) 115, 146
creating for HADR configuration 131
creating in existing DB2 instance by using scripts (AIX) 117
creating in existing DB2 instance by using scripts (Linux) 117
creating within DB2 database cluster 122
determining how to create 13
supported topologies 26
updating properties 290
updating properties for cluster 127, 135
metadata server
troubleshooting startup 423, 425
MFILES parameter 345
Microsoft Cluster service
creating users and user groups 188
installing InfoSphere Information Server in cluster 188
setting up virtual IP address 187
Microsoft failover clustering service
configuration example 1, services and engine 184
setting up high availability for engine and services tier 188
setting up network 186
setting up the cluster 186
testing the configuration 189
Microsoft Windows Cluster service (MSCS)
setting up network 172
MKS Toolkit 246
configuring for Windows parallel processing config 317
installation errors 440
monitoring
InfoSphere DataStage plug-in installation 250
installation logs 246
PX operators installation 250
WebSphere Application Server application deployment 249
WebSphere Application Server installation 248
WebSphere Application Server Liberty Profile application deployment 249
monitoring installation
overview 246
MPP 309
multi-client manager 3

N

naming conventions
projects 112
NAS (network attached storage) 53
national language support
supported character encodings 164
network attached storage (NAS) 53
networks
planning 96

O

ODBC
data source bind 344
vendor drivers 344
ODBC data sources
configuring access (AIX 341
configuring access (Linux) 341
configuring access (Windows) 345
ODBC drivers
testing connectivity 345
operations database
changing for high availability 297
installation planning 26
operations database, changing 287
order in which to install tiers 193

P

Page Not Found errors
troubleshooting 429
page tables
auto-tuning 331
paged pools
auto-tuning 330
paging space
on IBM AIX 328
on Linux 329
setting 328
PAM
configuring for InfoSphere Information Server engine 338
parallel engine
adding computers 309
C++ compilers 322
configuring for Windows 330
configuring heap size 332
configuring maximum connections 334
configuring maximum port number 334
copying components 309
data set descriptor files 329
editing boot.ini file 335
enabling auto-tuning for paged pool size 330
enabling auto-tuning for system page size 331
environment variables 324, 325
example scenario 59
monitoring system resources 6
running rsh without a password 311
setting paging space 328
setting paging space (IBM AIX) 328
setting paging space (Linux) 329
setting threshold for de-committing memory 332
using ssh 310
parallel jobs
configuring 326
setting the maximum memory 327
user privileges 312
parallel processing
configuring compute nodes (AIX) 309
configuring compute nodes (Linux) 309
configuring compute nodes (Windows) 313, 315
configuring MKS Rexecd and Rshd daemons 317
configuring remote shell script for parallel engine (Windows) 317
configuring RSH (Windows) 316
installing components (Windows) 314
making system environment changes (Windows) 314
massively parallel processing (MPP) systems 51
overview 50
strengthening remote shell permissions (Windows) 320
symmetric multiprocessing (SMP) systems 51
testing parallel engine configuration using InfoSphere DataStage Designer (Windows) 319

parallel processing *(continued)*
testing parallel engine configuration using MKS Korn shell (Windows) 318
transform library distribution methods 321
Passport Advantage 193
passwords
in response files 216, 217
installation 107
naming restrictions 110
patches
applying 359
permissions
directory and user 103
planning
capacity 54
client tier 3
clusters 17
connectivity 96
engine tier 6
grids 17
IDs 107
installation 3, 13
Match Designer database 143
naming, projects 112
networks 96
ports 96
repository tier 9
services tier 4
topologies 17
users 107
WebSphere Application Server 14
ports
changing 302
configuring for HADR 137
maximum number 334
planning 96
preinstallation requirements
directories 95
processes, used by InfoSphere Information Server 353
product accessibility
accessibility 445
product documentation
accessing 451
profiles
removing (AIX) 388
engine tier 389
metadata repository tier 390
services tier 388
removing (Linux) 388
engine tier 389
metadata repository tier 390
services tier 388
removing (Windows) 385
engine tier 386
repository tier 387
services tier 385
project naming 112
PX operators
monitoring installation 250

R

RAID 54
Reconfiguring 287
relocating a database 287
remote shell (RSH)
configuring for Windows parallel processing config 316
remote shell command
configuring 311
remove 268
removing
command prompt (silent) 380
InfoSphere Information Server 369
removing software
manual removal 393
repositories
configuring for high availability 297
Repositories
in the metadata repository database 295
repository tier
overview 9
setting up in existing DB2 instance 114
requirements
system 94
Resource Tracker
monitoring system resources 6
response files 215
AIX parameters 229
creating 215
encrypting information in 216
Linux parameters 229
passwords 217
samples 217
silent mode (removing software) 380
syntax 217
Windows parameters 218
RLTABSZ parameter 345
rsh command
configuring 311
program location 312
runtime errors
troubleshooting 439

S

SAN (storage area networks) 53, 54
saving settings
for InfoSphere Glossary Anywhere 271
secure shell program
configuring 310
security
configuring 279
configuring PAM 338
encrypting information in response files 216
server engine
starting and stopping 351
server parameters
configuring all clients 258
global configuration 258
InfoSphere Glossary Anywhere configuration 258
services tier
troubleshooting 426
services, used by InfoSphere Information Server 353
setup command syntax 244
silent installation 215
AIX parameters 229
by using a response file 240
InfoSphere Glossary Anywhere 266
Linux parameters 229
Windows parameters 218
silent mode
starting installation program in 215
starting software removal program in 380
software removal
Microsoft Windows security notifications 376
overview 373
silent mode 380
starting in graphical mode 376
troubleshooting 439
verifying 439
software requirements 94
software services
contacting 449
special characters
in command-line syntax 447
SQA reports
WebSphere Application Server 347
SRDDB
changing for high availability 297
ssh program
configuring 310
staging area, changing 287
standardization rules designer database
capacity planning 58
Standardization Rules Designer database
changing for high availability 297
locations 26
Standardization Rules Designer database, changing 287
starting and stopping
InfoSphere Information Server engine 351
starting the installation program
by using a response file 240
storage area networks (SAN) 53, 54
suite installation
monitoring 247
suite removal
in Windows 385
engine tier 386
repository tier 387
services tier 385
manually in Windows 393
client tier 398
engine tier 395
repository tier 396
services tier 393
manually on AIX 399
engine tier 401
repository tier 402
services tier 400
manually on Linux 399
engine tier 401
repository tier 402
services tier 400
on AIX 388
on Linux 388
removing remaining files 385
engine tier (AIX) 389

suite removal *(continued)*
removing remaining files *(continued)*
engine tier (Linux) 389
engine tier (Windows) 386
metadata repository tier (AIX) 390
metadata repository tier (Linux) 390
repository tier (Windows) 387
services tier (AIX) 388
services tier (Linux) 388
services tier (Windows) 385
shutting down tasks and applications 370
support
customer 449
symmetric multiprocessing (SMP) systems 51
syntax
command-line 447
System Automation for Multiplatforms
configuration example 170
system health checker 251
system pages
auto-tuning 331
system requirements 94
SystemErr.log 407
SystemOut.log 249, 407

T

T30FILES parameter 345
testing the installation
running the IBM Support Assistant Lite tool 251
tier computers
required directories 95
tiers
client 3
described 3
engine 6
installation options 206
product components contained in 3
relationships among 11
services 4
Tivoli System Automation for Multiplatforms
creating users and user groups 178
defining and starting resource group 175
installing 173
installing InfoSphere Information Server in cluster 178
setting up high availability control scripts 179
setting up virtual IP address 174
testing the configuration 181
Tivoli System Automation for Multiplatforms (SA MP)
installing in DB2 database cluster 121
topologies
active-passive 33
analysis databases and 26
client/server 18
computer for each tier 20
dedicated engine tier 19
topologies *(continued)*
fully clustered 37
installation 17
instance with multiple server engines 25
ITAG installations 21
Match Designer database and 26
multiple InfoSphere Information Server instances 21
single computer 17
trace logs
disabling 429
enabling 429
trademarks
list of 453
troubleshooting
application server 422
DB2 installation warnings 414
domain connections 422
enabling trace logs 429
failed installations 416, 439
file name in the package too long 411
ISALite tools 251
JVM failure 433
log files 407
login failure 433
messages during installation 415
Page Not Found errors 429
report results, viewing 429
runtime errors, uninstallation 439
services tier 426
software removal 439
user accounts 413
user privileges for installation 419
VMware ESX 414
WebSphere Application Server administrative console 435
WebSphere Application Server installation with InfoSphere Information Server 421
WebSphere Application Server registry 425

U

uinstallation script 268
uninstall 268
updating software 359
upgrading
monitoring overall suite installation 247
user IDs
naming restrictions 110
user limits
modifying 101
user privileges
overview 419
parallel jobs 312
users
planning 107
uvconfig file
GLTABSZ parameter 345
MAXRLOCK parameter 345
MFILES parameter 345
RLTABSZ parameter 345
T30FILES parameter 345
uvodbc.config file
configuring 343

V

verifying the installation
running the IBM Support Assistant Lite tool 251
virtual IP address
setting up for active-passive configuration (Microsoft example) 187
setting up for active-passive configuration (System Automation for Multiplatforms example) 174
VMware ESX 414

W

WebSphere Application Server
changing HTTP port 299
cluster members, creating 155
cluster, creating 155
configurations, required 148
front-end load balancer, for clustering 157
front-end web server, for clustering 157, 280
front-end Web server, for clustering 42
HTTP session persistence 155
installation troubleshooting 421
installing 148
Java 1.7, switching to 162
non-clustered installation 150
planning 14
planning for clusters 151
removing profiles 385, 388
reports 347
troubleshooting startup (Windows) 423, 425
WebSphere Application Server application deployment
monitoring 249
WebSphere Application Server cluster
administrative security 153
custom profile, creating 154
deployment manager profile, creating 153
Profile Management Tool 153, 154
WebSphere Application Server clustering
configurations, required 148, 152
WebSphere Application Server clusters
load balancing 44
running as a service 159
WebSphere Application Server deployment manager
node, federating 154
WebSphere Application Server installation
monitoring 248
WebSphere Application Server Liberty Profile
alternate address for cluster 127, 134
WebSphere Application Server Liberty Profile application deployment
monitoring 249

WebSphere Application Server Network Deployment
changing ports 302
cluster environments, installing 148, 152
cluster install options 41
installation 148, 152
replication domains 156
Windows Registry
editing for parallel engine 330
increasing heap size 332
MaxUserPort 334
PagedPoolSize 330
SystemPages 331
TcpNumConnections 334



Printed in USA

GC19-4316-00



Spine information:

IBM InfoSphere Information Server **Version 11 Release 3**

Planning, Installation, and Configuration Guide

