IBM FileNet P8
Version 4.5.1

Plan and Prepare Your Environment for IBM FileNet P8 for installation on Microsoft Windows with Microsoft SQL Server, Oracle WebLogic Server, and Windows Active Directory

IBM®

IBM FileNet P8
Version 4.5.1

Plan and Prepare Your Environment for IBM FileNet P8 for installation on Microsoft Windows with Microsoft SQL Server, Oracle WebLogic Server, and Windows Active Directory

Contents

ibm.com and related resources

Product support and documentation are available from ibm.com.

Support and assistance

Product support is available on the Web. Click Support from the product Web site at:

FileNet Content Manager Support
> http://www.ibm.com/software/data/content-management/filenet-content-manager/support.html

Information center

You can view the product documentation in an Eclipse-based information center that you can install when you install the product. By default, the information center runs in a Web server mode that other Web browsers can access. You can also run it locally on your workstation. See the information center at http://publib.boulder.ibm.com/infocenter/p8docs/v4r5m1/index.jsp.

PDF publications

You can view the PDF files online using the Adobe Acrobat Reader for your operating system. If you do not have the Acrobat Reader installed, you can download it from the Adobe Web site at http://www.adobe.com.

See the following PDF publications Web sites:

Product	Web site
Product Documentation for FileNet P8 Platform	http://www.ibm.com/support/docview.wss?rs=3278&uid=swg27010422

"How to send your comments"
Your feedback is important in helping to provide the most accurate and highest quality information.

"Contacting IBM" on page vi
To contact IBM customer service in the United States or Canada, call 1-800-IBM-SERV (1-800-426-7378).

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

Send your comments by using the online reader comment form at https://www14.software.ibm.com/webapp/iwm/web/signup.do?lang=en_US &source=swg-rcf.

Consumability survey

You are invited to tell IBM how to improve the consumability of software products. If you want to help IBM make IBM® FileNet® P8 easier to use, take the Consumability Survey at http://www.ibm.com/software/data/info/consumability-survey/.

Contacting IBM

To contact IBM customer service in the United States or Canada, call 1-800-IBM-SERV (1-800-426-7378).

To learn about available service options, call one of the following numbers:
- In the United States: 1-888-426-4343
- In Canada: 1-800-465-9600

For more information about how to contact IBM, see the Contact IBM Web site at http://www.ibm.com/contact/us/.

Planning and preparing for IBM FileNet P8 installation

To prepare to install IBM FileNet P8 Platform, you must review the planning information before you begin. You must also complete the prerequisite tasks assigned to the various roles.

"Planning the installation"
You must review the installation planning information before your IBM FileNet P8 Platform installation so that you know what kind of deployments are supported, understand how the tasks in the installation tasks are organized by role, and know how to use the Installation and Upgrade Worksheet.

"Performing the required installation preparation tasks" on page 9
To efficiently carry out the required installation preparation tasks, you must assign your staff to carry out the tasks that are organized by administrative role.

Planning the installation

You must review the installation planning information before your IBM FileNet P8 Platform installation so that you know what kind of deployments are supported, understand how the tasks in the installation tasks are organized by role, and know how to use the Installation and Upgrade Worksheet.

"FileNet P8 Platform sample architecture"
You can distribute IBM FileNet P8 Platform components and expansion products across a variety of machines.

"Planning your deployment" on page 3
Content Engine, Process Engine, and Application Engine, or Workplace XT, can all be installed as standalone, single instance applications, or as multiple instances with different names on the same server, or in a cluster.

"Definition of installation roles" on page 5
The tasks in this guide and the rows in the Installation and Upgrade Worksheet are organized by administrative roles. Your organization might have different roles, and some of the responsibilities of listed roles will vary from those assigned by default in this documentation.

"Using the installation and upgrade worksheet" on page 7
The Installation and Upgrade Worksheet is a Microsoft® Excel spreadsheet (p8_worksheet.xls). The worksheet describes the properties and parameters required to complete IBM FileNet P8 installation, upgrade, and configuration programs, and provides a way to record the values you assign to these properties and parameters.

FileNet P8 Platform sample architecture

You can distribute IBM FileNet P8 Platform components and expansion products across a variety of machines.

"Baseline configuration with expansion components" on page 2
A production environment includes required components plus one or more optional expansion components.

"Sample development architecture" on page 2
To provide an application development configuration, you can install the components that the application is based on and also the required API toolkits on developer workstations.

1

You can install required and any expansion components on a single server in order to support demonstrations, proof-of-concepts, and single-server development. Because Enterprise Manager runs only on Windows, a single-server installation that includes this administration tool must be on a Windows server.

Baseline configuration with expansion components

A production environment includes required components plus one or more optional expansion components.
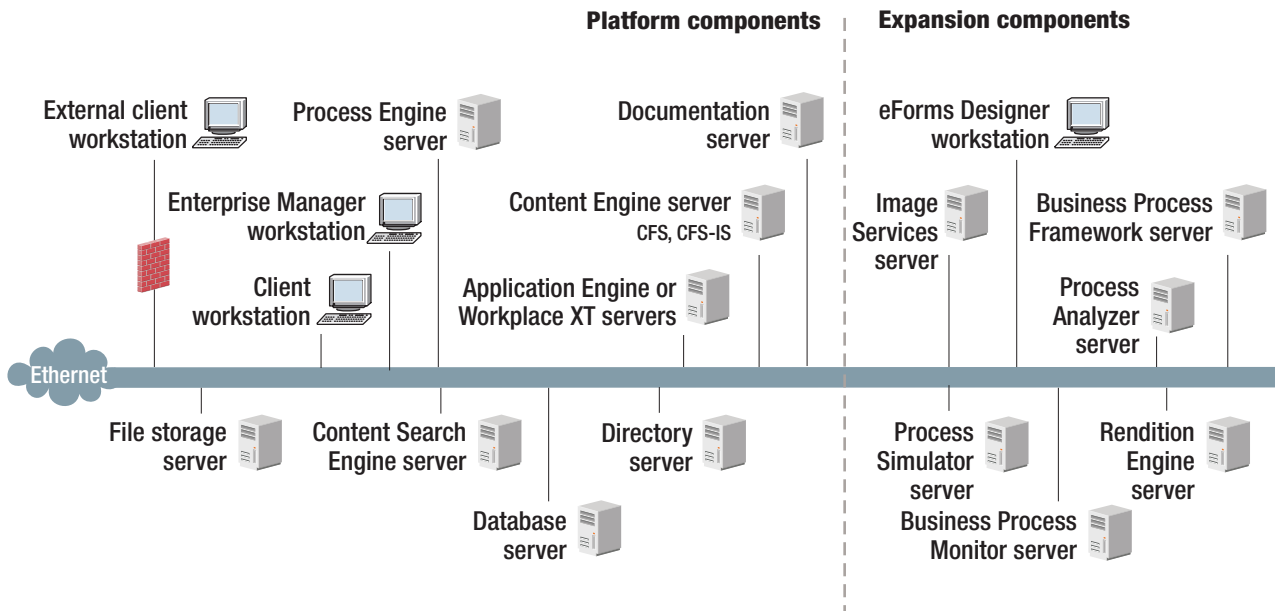
Figure 1. Baseline configuration with expansion components

To understand this graphic, keep in mind the following details:

- You must set up Content Engine, Application Engine, and the documentation server on application servers.
- IBM FileNet P8 Platform supports several different operating systems. However, Enterprise Manager runs only on Windows.
- The graphic does not show components in a high availability or clustered configuration.
- For information about collocating components on a single server, see *IBM FileNet P8 Hardware and Software Requirements*.
- Not all expansion components are shown, for example IBM FileNet P8 Portlets and Image Services Resource Adapter. Check with your service representative for availability of other expansion products.

Sample development architecture

To provide an application development configuration, you can install the components that the application is based on and also the required API toolkits on developer workstations.

Refer to the IBM FileNet P8 Developer Help topic **IBM FileNet P8 Documentation** → **Developer Help** → **Developer Roadmap** → **Introduction** for information on setting up your development environment and installing the IBM FileNet P8 API toolkits.

Sample demonstration architecture

You can install required and any expansion components on a single server in order to support demonstrations, proof-of-concepts, and single-server development. Because Enterprise Manager runs only on Windows, a single-server installation that includes this administration tool must be on a Windows server.

Installing many FileNet P8 components on the same server is likely to affect system performance. For more information, see the *IBM FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Performance Tuning Guide*.

Planning your deployment

Content Engine, Process Engine, and Application Engine, or Workplace XT, can all be installed as standalone, single instance applications, or as multiple instances with different names on the same server, or in a cluster.

"Planning Content Engine deployment"
You can deploy Content Engine in either a stand-alone, multi-instance single-server, or cluster environment.

"Application server planning considerations" on page 4
To prepare your application servers, you must carry out certain tasks including creating administrative accounts, specifying environmental variables, and making sure your JDBC drivers are at a supported level.

Planning Content Engine deployment

You can deploy Content Engine in either a stand-alone, multi-instance single-server, or cluster environment.

Standalone deployment

When you deploy Content Engine as a stand-alone application, you configure a single application server. You must configure your Content Engine instances and deploy those Content Engine instances on a single server, using a single directory for the configuration files.

Multi-instance single-server deployment

When you deploy multiple Content Engine instances on a single server, you configure a single application server.

Each Content Engine instance is isolated from the others, and there is no exchange of information between the instances. For example, you can dedicate a Content Engine instance for use by the Research and Development department, and you can dedicate a Content Engine instance for use by the Human Resources department. Each instance has its own set of object stores.

You must perform all the procedures for each instance on a single server, using a different directory for the configuration files for each instance. It is a best practice to deploy only a single instance into a cluster.

Managed deployment

When you deploy Content Engine in an Oracle WebLogic Server managed environment, you must install and configure Content Engine on the administrative node to avoid cross-network configuration issues. Install and configure Content Engine on the administration server. Then use the administration server tools to deploy the Content Engine EAR file to the managed servers. The managed servers can be in a cluster.

Cluster deployment

Cluster deployments can be used for load balancing or to provide high availability. Refer to the *IBM FileNet P8 High Availability Technical Notice* for details on how to set up your FileNet P8 system using clusters, farms, and other high availability software and hardware.

Application server planning considerations

To prepare your application servers, you must carry out certain tasks including creating administrative accounts, specifying environmental variables, and making sure your JDBC drivers are at a supported level.

Content Engine and Application Engine are J2EE application server-based applications. (Process Engine is not.) You must install Content Engine and Application Engine in a homogeneous J2EE environment in which all of your application servers and their version numbers are identical for both components.

Assuming that a user application is required for your system and that you have not built or customized one using the IBM FileNet P8 API toolkits, you can install either of the following general user interfaces:
- Workplace XT
- Application Engine /Workplace

Even if you have your own customized application, it is a best practice to install one of these applications for testing and support purposes. The applications must use Enterprise Java Bean (EJB) transport.

If the application server where Content Engine will be deployed is running on a 32-bit JVM, it is a best practice to create no more than 75 Content Engine object stores. On a 64-bit JVM, it is a best practice to create no more than 150 Content Engine object stores.

Content Engine is a resource-intensive enterprise application. Running Content Engine and other J2EE applications on the same machine is possible but not a best practice. Other J2EE applications will compete with Content Engine for the same CPU, memory, and disk I/O resources, and increase the complexity of the installation and the risk of the deployment, because configurations will not match what has been qualified by IBM FileNet P8 Engineering.

Although you might need to host Content Engine and other applications on the same machine, it is preferable to host Content Engine on its own machine or logical partition. If an architecture requires Content Engine and a non-P8 J2EE application to be on the same machine, be sure to thoroughly test the configuration in your integration environment before deploying them into production.

See the IBM FileNet P8 help topic **System Administration → Enterprise-wide Administration → FileNet P8 Security → Authentication** for reference information about support for EJB and Web Services transports.

"WebSphere Application Server and WebLogic Server - Content Engine multi-server deployment"
In an environment of load-balanced or highly-available farmed (or clustered) application servers, you will initially install Content Engine on the Deployment Manager node (IBM WebSphere® Application Server) or the Administrator node (WebLogic Server).

WebSphere Application Server and WebLogic Server - Content Engine multi-server deployment:

In an environment of load-balanced or highly-available farmed (or clustered) application servers, you will initially install Content Engine on the Deployment Manager node (IBM WebSphere Application Server) or the Administrator node (WebLogic Server).

To install Content Engine in centrally managed servers in farms or clusters, see the *IBM FileNet P8 High Availability Technical Notice*.

Important: The base version of WebSphere Application Server that is bundled with Content Engine does not include cluster support. To configure clusters, you must have Content Engine installed on servers that are running the IBM WebSphere Application Server Network Deployment version.

Note the impact of deploying centrally managed servers not in farms or clusters. In an environment in which multiple Content Engine instances are geographically dispersed, and where each instance might have its own local Application Engine or Workplace XT server, you will install Content Engine on the Deployment Manager node (WebSphere Application Server) or the Administrator node (WebLogic Server). After installing Content Engine, you will use Configuration Manager to deploy it to other servers.

Definition of installation roles

The tasks in this guide and the rows in the Installation and Upgrade Worksheet are organized by administrative roles. Your organization might have different roles, and some of the responsibilities of listed roles will vary from those assigned by default in this documentation.

Installation administrator
- Runs IBM FileNet P8 installers during initial setup.
- Runs the Configuration Manager tool during initial setup, followed by starting IBM FileNet Enterprise Manager.
- Runs IBM FileNet P8 Upgrade programs during upgrades.
- Abbreviated as IA. Responsible for coordinating the information described in this worksheet. The information itself will require the input from the other roles.

The role of IA is usually filled by an IBM FileNet Certified Professional (FCP).

Information technology administrator
- Responsible for the networking and operating systems settings required by IBM FileNet P8.
- Responsible for performing certain security configurations.

- Abbreviated as ITA. Responsible for providing the information in the rows in the *Installation and Upgrade Worksheet* with a value of ITA in the Role column.

Security administrator

- Responsible for configuring the directory servers required by IBM FileNet P8 components, including Content Engine and Application Engine.
- Creates and maintains directory server user and group accounts.
- Abbreviated as SA. Responsible for providing the information in the rows in the *Installation and Upgrade Worksheet* with a value of SA in the Role column.

Database administrator

- Creates, configures, maintains database installations and database or table spaces.
- Responsible for creating database accounts needed by IBM FileNet P8.
- For purposes of this documentation, the database administrator is expected to have responsibilities regarding the JDBC data sources.
- Abbreviated as DBA. Responsible for providing the information in the rows in the *Installation and Upgrade Worksheet* with a value of DBA in the Role column.

Application server administrator

- Responsible for providing the application servers required by IBM FileNet P8.
- Responsible for application server administrative accounts.
- Abbreviated as ASA. Responsible for providing the information in the rows in the *Installation and Upgrade Worksheet* with a value of ASA in the Role column.

IBM FileNet P8 administrator

- This role designation actually refers to the administrator or administrators who perform regular maintenance of Content Engine, Process Engine Application Engine, Workplace or Workplace XT.
- The administrator who logs on to Enterprise Manager by using the gcd_admin account or an object_store_admin account is considered an IBM FileNet P8 administrator.
- Abbreviated as P8A. Responsible for providing the information in the rows of the *Installation and Upgrade Worksheet* with a value of P8A in the Role column.

E-mail Administrator

- Creates an e-mail account that will be used to configure the **Notification** tab of the Process Task Manager so that Process Engine can send e-mail notifications to end users. (Required only if you use this feature.)
- Abbreviated as EA.

Related concepts

"Using the installation and upgrade worksheet"
The Installation and Upgrade Worksheet is a Microsoft Excel spreadsheet
(`p8_worksheet.xls`). The worksheet describes the properties and parameters
required to complete IBM FileNet P8 installation, upgrade, and configuration
programs, and provides a way to record the values you assign to these properties
and parameters.

The Information Technology administrator must prepare the network and
operating systems, and carry out certain security configurations to prepare your
environment for IBM FileNet P8 Platform.

The Security administrator must prepare the security environment for IBM FileNet
P8 Platform, including planning the security environment, configuring the
directory server, and creating accounts.

The Application Server Administrator must prepare the application servers for IBM
FileNet P8 Platform, including planning deployment, creating administrative
accounts, and configuring JDBC drivers for both Content Engine and Application
Engine.

Related tasks

The Database administrator must prepare the databases required for IBM FileNet
P8, including gathering information about data sources, creating databases and
database accounts, and installing client software.

Using the installation and upgrade worksheet

The Installation and Upgrade Worksheet is a Microsoft Excel spreadsheet
(`p8_worksheet.xls`). The worksheet describes the properties and parameters
required to complete IBM FileNet P8 installation, upgrade, and configuration
programs, and provides a way to record the values you assign to these properties
and parameters.

Administrators who are preparing the environment for installation or upgrade of
IBM FileNet P8 components must use the worksheet during their preparation tasks
to record the appropriate values and provide them to the Installation
Administrator who runs the installation or upgrade programs.

Some of the features of the Installation and Upgrade Worksheet are:
- **Instructions**: describes the worksheet and includes a button that runs the
  Customize Worksheet macro.
- The two highlighted columns, **Property or Parameter** and **ENTER YOUR
  VALUE HERE**, provide the simplest view of the requirement. The others add
  identifying information and help you sort and filter the rows usefully.
- The **Role** column assigns each row to an administrator and uses the following
  acronyms:
  – ITA: Information Technology Administrator
  – ASA: Application Server Administrator
  – DBA: Database Administrator
  – SA: Security Administrator
  – P8A: IBM FileNet P8 Administrator
- Property definitions are contained in the column titled **Description**.

- Some rows, though not all, contain a hyperlink in the **IC help link** column. Click this hyperlink to run a query against the IBM Information Center, which opens with the Search Results pane showing the topics that contain the words in the query phrase. Browse the search results until you have enough information to be able to enter a value in the Worksheet row.

  "Running the Customize Worksheet macro"
  The Customize Worksheet macro lets you extract only those rows that describe your environment.

  "Autofiltering and sorting the Worksheet"
  There are several ways to organize the Worksheet to make finding properties and entering values easier.

Running the Customize Worksheet macro

The Customize Worksheet macro lets you extract only those rows that describe your environment.

Important: For support of the full range of built-in filter and macro features, use Microsoft Excel to view the Installation and Upgrade Worksheet file. You can use other spreadsheet programs to view the file; however, filter and macro support can vary. For example, in Calc from OpenOffice.Org, the column filters work as expected, but the Customize Worksheet button does not.
To run the Customize Worksheet macro:

1. Open the Installation and Upgrade Worksheet (`p8_worksheet.xls`) and click the **Instructions** worksheet (also called a tab).
2. Scroll down until you see the button representing the Customize Worksheet macro. Click the button.
3. Select the components and options that describe the environment you are preparing for IBM FileNet P8.
   - Installation or Upgrade
   - FileNet P8 Components
   - Application Server type
   - Operating system
   - Database type
   - Directory Server type
   - Number of object stores (adds new sets of rows for creating additional data sources)
   - Name of customized sheet
4. Click **OK**. The macro copies the rows that fulfill your selection criteria into a new worksheet with the name you entered. Enter the values for your environment into this new worksheet.
5. Click the name of the new worksheet at the bottom of the Excel window. Add your preparation values into this new worksheet.
6. Notice that the new worksheet has buttons at the top titled **Show Installer View** and **Show Full View**, depending on its state. The **Show Installer View** displays only those columns that you need while running installation or configuration programs.

Autofiltering and sorting the Worksheet

There are several ways to organize the Worksheet to make finding properties and entering values easier.

AutoFiltering is a quick way to display only those rows that meet a certain criteria. To use AutoFilter:

1. Make sure AutoFiltering is enabled. (Select the entire row with the column headers, then click **Data** → **Filter** → **Autofilter**.) AutoFilter arrows will appear to the right of the column labels.

2. Click the **AutoFilter** arrow in the **Installation or Configuration Program** column header and select the program you are interested in (for example, PE installer).

3. Click the **AutoFilter** arrow in the **Setup Type** column header, select **Custom**, and specify **Setup Type contains Installation**.

4. For a custom AutoFilter, click the **AutoFilter** in any column header, select **Custom**, and specify Setup Type contains "Installation".

5. To turn off AutoFiltering in a column, click the column **AutoFilter** arrow and select **(All)**.

6. To reorder rows alphabetically, do a Sort:

   a. Click anywhere in a column, for example, Column A Role.

      The only possible values in the Role column are ASA, SA, DBA, ITA, and P8A. Sorting on Role therefore groups the rows by this attribute, in alphabetic order. Several other columns also have a limited number of possible values which means they can be usefully sorted.

   b. Click the **Sort Ascending** icon in the Excel toolbar, or use the **Data** → **Sort** menu command. The rows sort on Role.

      Sorting the Worksheet reassigns row numbers. If you refer to rows by number, be aware that row numbers change if you change the sort order.

Performing the required installation preparation tasks

To efficiently carry out the required installation preparation tasks, you must assign your staff to carry out the tasks that are organized by administrative role.

Some tasks require input that results from other preparation tasks performed by other administrator roles. While performing the tasks, record results in the *Installation and Upgrade Worksheet*. See the "Using the installation and upgrade worksheet" on page 7 topic for details.

To prepare the IBM FileNet P8 environment, perform the tasks assigned to the following roles.

"IT administrator installation tasks" on page 10
The Information Technology administrator must prepare the network and operating systems, and carry out certain security configurations to prepare your environment for IBM FileNet P8 Platform.

"Security administrator installation tasks" on page 26
The Security administrator must prepare the security environment for IBM FileNet P8 Platform, including planning the security environment, configuring the directory server, and creating accounts.

"Database administrator installation tasks" on page 38
The Database administrator must prepare the databases required for IBM FileNet P8, including gathering information about data sources, creating databases and database accounts, and installing client software.

"Application Server administrator installation tasks" on page 51
The Application Server Administrator must prepare the application servers for

IBM FileNet P8 Platform, including planning deployment, creating administrative accounts, and configuring JDBC drivers for both Content Engine and Application Engine.

IT administrator installation tasks

The Information Technology administrator must prepare the network and operating systems, and carry out certain security configurations to prepare your environment for IBM FileNet P8 Platform.

- Review all rows assigned to the IT administrator (ITA) in the "Using the installation and upgrade worksheet" on page 7. While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation.

   Tip: With the **Data** → **Filter** → **AutoFilter** command enabled, as it is by default in the shipping worksheet file (`p8_worksheet.xls`), perform the following actions to quickly see only the properties assigned to a particular role:
   – Click the **AutoFilter** drop-down arrow in the **Role** column header and select ITA.
   – Further filter the result set by clicking the **AutoFilter** drop-down arrow in any of the other columns and selecting a value or clear a filter by selecting **All**.

- If you are installing in a non-English environment, review the considerations and procedures in *IBM FileNet P8 Non-English Support Guide* before you begin your preparation tasks.

   "Creating Content Engine operating system accounts"
   The IT Administrator must create several operating system accounts required by FileNet P8 components during installation.

   "Creating Process Engine operating system accounts" on page 13
   Create new or designate existing operating system accounts for Process Engine, as shown in the following table.

   "Configuring Microsoft Windows" on page 16
   You must perform certain operating system procedures on all Windows-based servers where you will installFileNet P8, including planning for adequate disk and temp space and port requirements, and installing required levels of Microsoft .NET Framework, Web Services Enhancements (WSE), or Windows® Communication Foundation (WCF).

   "Configuring the network" on page 20
   You must perform certain configurations on the network before installing IBM FileNet P8 Platform.

   "Preparing storage areas for object stores" on page 21
   To prepare for file storage, you must prepare the locations where file stores will go, configure remote access protocol, and, if you plan to implement content-based retrieval, perform security configurations.

Creating Content Engine operating system accounts

The IT Administrator must create several operating system accounts required by FileNet P8 components during installation.

Create the following users and groups:

User or group name	Description
Application Server Installation Administrator: (WebSphere Application Server, WebLogic Server, JBoss Application Server): *ce_appserver_install_user*	An operating system user account you used to install your application server. Use your local machine's administrative tools to grant *ce_appserver_install_user* at least the following permissions: • For Windows, *ce_appserver_install_user* must be a member of the local administrators group. Later in the installation process, the *ce_appserver_install_user* performs the following tasks: • Create and configure the application server/domain/profile for Content Engine. • Start or stop the application server when needed. • Modify the application server files or directories as needed for deploying the Content Engine application using the Configuration Manager tool. • Provide create, read and write permissions for directories on devices or drives that are used for external Content Engine file storage. *ce_appserver_install_user* must be a member of the *ce_appserver_install_group*.
Application Server Installation Group: *ce_appserver_install_group*	Create an operating system group account and add to it the following accounts: • *ce_appserver_install_user* • *ce_install_user* (Windows) or *ce_install_user* (UNIX). • *config_mgr_user* Use the user accounts in *ce_appserver_install_group* to do the following tasks: • Give operating system privileges to the directories used for the Content Engine Installation and for the application server's instance/domain/profile. • Configure and deploy the Content Engine EAR files which require access to the application server's `instance/domain/profile` directories. • Have permissions on devices/drives to read and write that are designated for external Content Engine file storage.

User or group name	Description
Content Engine Installer (Windows): *ce_install_user*	An operating system account you will use to install Content Engine. *ce_install_user* can be the same user as the *ce_appserver_admin*. Use Windows administrative tools to add *ce_install_user* to the Local Administrators group and to the *ce_appserver_install_group*.

Configuration Manager user (WebSphere, WebLogic, JBoss): *config_mgr_user*	The operating system account you will use to run Configuration Manager. At several points in *Plan and Prepare Your Environment for IBM FileNet P8* you will be instructed to grant additional permissions to *config_mgr_user*, including the following permissions:
	
	Add *config_mgr_user* to the *ce_appserver_install_group*.
	
	(Windows only) Using Active Directory tools, add *config_mgr_user* to either the Power Users group or the Local Administrators group.
	
	At several points in the *IBM FileNet P8 Platform Installation and Upgrade Guide* you will be instructed to grant additional permissions to *config_mgr_user*, including the following permissions:
	
	• Execute permission to the Configuration Manager tool executable file, `configmgr.exe` (Windows) or `configmgr.sh` (UNIX).
	
	• Write permission to the directory where the Configuration Manager tool will create the configuration XML files. For example:
	  – the directory you specify using the optional **-path** parameter when you run the tool
	  – the default directory, `ce_install_path/tools/configurationmanager/tasks`, if you don't specify a path parameter
	
	• Read and write permission on the contents of the directory specified by **-path** parameter, or the default directory.

User or group name	Description
Content Engine Operating System user: *ce_os_user*	An operating system account you must log on as to create and configure the shared root directory of a file storage area or content cache area. See "Preparing storage areas for object stores" on page 21. **Windows** For Windows-based Content Engine and file storage areas, *ce_os_user* must reside in the same Windows domain or in trusted Windows domains as the servers that host Content Engine and the file storage area.

Creating Process Engine operating system accounts

Create new or designate existing operating system accounts for Process Engine, as shown in the following table.

Create several operating system users and groups for Process Engine.

Create the following users and groups:

User or group name	Description
The operating system user Process Engine uses to start and stop software and services (Windows): *fnsw*	This user is an operating system user that is used to start and stop Process Engine software and services. It can also optionally be used to install Process Engine. On Windows the IMS ControlService and Process Engine Services Manager service are configured to start with the *fnsw* user and the *fnsw_password* that is entered during Process Engine installation. Permissions on the Process Engine software directory are also set for this local user.
	Record this value in the Installation and Upgrade Worksheet.
	The Process Engine installation program asks if default users will be used or if aliases for the users will be defined. If a user name other than the default FNSW is defined, indicate to the Process Engine installation program that aliases will be defined. When prompted, provide this *fnsw* local user name to the installation program. Assign the *fnsw* local user to the following groups:
	• *fnusr*
	• *fnadmin*
	• *fnop*
	• Administrators (if it exists)
	The Process Engine installation program sets the following policies for this user automatically:
	• Act as part of the operating system
	• Adjust memory quotas for a process
	• Log on as a service
	• Replace a process level token
	If you choose to install Process Engine while logged on as this user, and run Process Engine SQL scripts from the installation program, ensure that this user has privileges in the database to run the scripts.
The password for the *fnsw* user (Windows): *fnsw_password*	This password is for the *fnsw* local user.
	Record this value in the Installation and Upgrade Worksheet.
	Process Engine installation program on Windows prompts for the *fnsw_password*. Leaving the field for the *fnsw_password* blank during Process Engine installation causes the installation program to set the default password, BPMtemp1pzwd.
	If this password does not match the operating system password, these services will not start.
	Set this password to never expire. Process Engine activity ceases if the *fnsw_password* expires.

User or group name	Description
An operating system group for Process Engine: *fnop*	This operating system group is a group account for Process Engine whose members have operator non-administrator privileges on Image Services used by Process Engine.

Record this value in the Installation and Upgrade Worksheet.

Add the *fnsw* user to this group.

If you will install Process Engine or run Process Engine Windows services as either the *fnsw* domain user or *pe_install_user_domain*, you must add that domain user to this local group.

The Process Engine installation asks if default users and groups will be used or if aliases will be defined. If a group name other than the default *fnop* is assigned, indicate to the Process Engine installation program that aliases will be defined and provide this name to the installation program as an alias for the default group. |
| An operating system group for Process Engine (Windows): *fnusr* | This operating system group is a group account for Process Engine whose members have non-administrator privileges on Process Engine files and databases.

Record this value in the Installation and Upgrade Worksheet.

The Process Engine installation asks if default users and groups will be used or if aliases will be defined. fnusr is the default group name. If a group name other than the default is assigned, indicate to the Process Engine installation program that aliases will be defined and provide this name to the installation program as an alias for the fnusr group.Assign the following users to the *fnusr* group:

- *fnsw*

- Administrator

If you will install Process Engine or run Process Engine Windows services as either the *fnsw* domain user or *pe_install_user_domain*, you must add that domain user to this local group. |

User or group name	Description
An operating system group for Process Engine (Windows): *fnadmin*	This operating system group is a group account for Process Engine whose members have all privileges on Process Engine files and databases. Record this value in the Installation and Upgrade Worksheet. The Process Engine installation asks if default users and groups will be used or if aliases will be defined. fnadmin is the default group name. If a group name other than the default is assigned, indicate to the Process Engine installation program that aliases will be defined and provide this name to the installation program as an alias for the fmadmin group.Add the following users to this group: • *fnsw* • Administrator If you will install Process Engine or run Process Engine Windows services as either the *fnsw* domain user or *pe_install_user_domain*, you must add that domain user to this local group.
An operating system user for Process Engine (Windows): *pe_install_user_local*	This is a Windows operating system account that can be used to install Process Engine software who is a Windows administrator for the local machine. If you choose to install Process Engine while logged on as this user, and run Process Engine SQL scripts from the installation program, ensure that this user has privileges in the database to run the scripts. You can log on as the *fnsw* user or create this user or the *pe_install_user_domain* user (with the specified permissions) expressly for the purpose of running the Process Engine installation program.

Configuring Microsoft Windows

You must perform certain operating system procedures on all Windows-based servers where you will installFileNet P8, including planning for adequate disk and temp space and port requirements, and installing required levels of Microsoft .NET Framework, Web Services Enhancements (WSE), or Windows Communication Foundation (WCF).

"Configuring Windows for IBM FileNet P8 servers" on page 17
To configure Windows for IBM FileNet P8 servers, ensure the minimum disk and temporary space and determine the port requirements.

"Configuring Windows for Enterprise Manager, .NET and COM compatibility clients" on page 17
You must first install Microsoft .NET Framework and either Web Services Enhancements (WSE) or Windows Communication Foundation (WCF) before you can install .NET API Clients, COM Compatibility clients, or Enterprise Manager.

"Configuring Windows for Content Engine on Active Directory" on page 17
If Windows Active Directory is your directory service, set the primary DNS server IP address on your Content Engine machine to the IP address of the machine where DNS is installed.

"Configuring Windows for Process Engine"
You must perform certain operating system procedures on all Windows-based
servers where you will install Process Engine, including verifying TCP/IP
settings and configuring the hosts file.

"Configuring Windows for Content Search Engine" on page 19
To configure Content Search Engine servers:

Configuring Windows for IBM FileNet P8 servers:

To configure Windows for IBM FileNet P8 servers, ensure the minimum disk and
temporary space and determine the port requirements.

To configure Windows for IBM FileNet P8 servers:
- Ensure minimum required disk space and temporary space for installation. See
  IBM FileNet P8 4.5 Hardware and Software Requirements.
- Determine port requirements. Consult with the application server, database, and
  FileNet P8 administrators to determine port requirements for all the servers in
  your installation environment.

**Configuring Windows for Enterprise Manager, .NET and COM compatibility
clients:**

You must first install Microsoft .NET Framework and either Web Services
Enhancements (WSE) or Windows Communication Foundation (WCF) before you
can install .NET API Clients, COM Compatibility clients, or Enterprise Manager.

To configure Windows for Enterprise Manager, .NET, and COM compatibility
clients:

If you have not already done so, install the required Windows components.

Option	Description
Microsoft Windows XP, Microsoft Windows Server 2003	Microsoft .NET Framework 2.0 and Web Services Enhancements (WSE) 3.0
Content Engine version 4.5.1 Fix Pack 002 and higher with Microsoft Windows Vista, Microsoft Windows Server 2008, Microsoft Windows 7	Microsoft .NET Framework 3.0 and Windows Communication Foundation (WCF)

Configuring Windows for Content Engine on Active Directory:

If Windows Active Directory is your directory service, set the primary DNS server
IP address on your Content Engine machine to the IP address of the machine
where DNS is installed.

Configuring Windows for Process Engine:

You must perform certain operating system procedures on all Windows-based
servers where you will install Process Engine, including verifying TCP/IP settings
and configuring the hosts file.

"Verifying TCP/IP parameter settings (Windows)" on page 18
To verify TCP/IP parameter settings, log on as the administrator and run
`regedit`.

"Adding inbound rules to the Windows 2008 firewall"
Configure inbound rules in the Windows 2008 firewall to allow the following
Process Engine ports access.

"Configuring the /etc/hosts file"
Information related to Process Engine IP address, server name and optional
load balancer must be entered into either the server's DNS table or the hosts
file.

"Configuring Process Engine clients for ORB (Windows)" on page 19
Process Engine clients require either the IBM or the Sun Object Request Broker
(ORB).

Verifying TCP/IP parameter settings (Windows):

To verify TCP/IP parameter settings, log on as the administrator and run `regedit`.

1. Log on as the Administrator user and run `regedit` to verify the following
   registry key values. These values are decimal. The default in Registry Editor is
   hexadecimal.

   ```
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
   Tcpip\Parameters\MaxUserPort => 65534 (default = 5000)

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
   Tcpip\Parameters\TcpTimedWaitDelay => 90 (default = 240, or 4 min)
   ```

2. If necessary, add or modify a new DWORD value with the values as described
   above and save the changes.

Adding inbound rules to the Windows 2008 firewall:

Configure inbound rules in the Windows 2008 firewall to allow the following
Process Engine ports access.

Port	Protocol
32768	TCP
32769	TCP
32770	UDP
32776	TCP
32777	TCP

Configuring the /etc/hosts file:

Information related to Process Engine IP address, server name and optional load
balancer must be entered into either the server's DNS table or the hosts file.

For non-farmed configurations, the information can be in either the DNS table or
the hosts file on the server. For farmed configurations, the information must be
entered into the host file. In a farmed environment, entries must exist for every
Process Engine server in the farm.

Process Engine can be configured to run with IPv4, IPv6, or a dual stack network.

Entries must use the following format for each Process Engine server. The load
balancer name must also be associated with the appropriate server in a farmed
configuration.

```
IP_addr hostname load_balancer_name
```

IP_addr
> The IP address of the Process Engine server.

hostname
> The corresponding host name, also referred to as the network name. This name must be < 32 characters and cannot contain any dots. Record this value in the Installation and Upgrade Worksheet as the network name and provide it to the Process Engine installation program.

load_balancer_name
> The name of the load balancer in a farmed configuration.

In the following example there are IPv4 and IPv6 addresses and cmepwin11 is the hostname (network name). If the server is configured for a dual stack network, both IPv4 and IPv6 entries must exist for the same hostname.

```
127.0.0.1       localhost
::1             localhost
2007:0:0:0:9:181:124:192  cmepwin11
#fe80::fda2:48f4:7106:2f19%10 cmepwin11
123.45.6.78  cmepwin11
```

Configuring Process Engine clients for ORB (Windows):

Process Engine clients require either the IBM or the Sun Object Request Broker (ORB).

This applies to the following configurations:
- J2EE application server clients such as Workplace or Workplace XT
- Content Engine when using the workflow subscription processor to launch workflows
- Non-J2EE or custom applications

Therefore, if you have Windows-based non-J2EE or custom applications, configure the Java installation on those servers with either the IBM or the Sun ORB.

Application server	Changes
Oracle WebLogic Server with the Sun JVM	No changes are required.
WebLogic Server with the JRockit JVM	Replace the default ORB with the Sun ORB by adding the following code to your application server startup file. Enter without carriage returns:
	`set JAVA_OPTIONS=%JAVA_OPTIONS% -Dorg.omg.CORBA.ORBClass =com.sun.corba.se.impl.orb.ORBImpl set JAVA_OPTIONS=%JAVA_OPTIONS% -Dorg.omg.CORBA.ORBSingletonClass =com.sun.corba.se.impl.orb.ORBSingleton`

Configuring Windows for Content Search Engine:

To configure Content Search Engine servers:
1. Set the *k2_os_user* user as an administrator on each Windows Content Search Engine server. See "Specifying IBM FileNet P8 accounts" on page 29 for details on creating required users and groups.

2. On each Windows machine where you intend to install an additional locale (other than English), complete the following substeps. If you do not intend to install additional locales, skip this step.

   a. Install the Java™ 1.4 SDK.

   b. Set the PATH environment variable on the Windows machine to *Java (SDK) 1.4 install_path*\bin, where *Java (SDK) 1.4 install_path* is the install path to Java 1.4 SDK. For example, the install path might be C:\Program Files\Java\jdk1.4.2_17.

3. Install Java 1.5 SDK and set the JAVA_HOME environment variable on each machine where you will install Content Search Engine software, as follows:

   a. Install the Java 1.5 SDK.

   b. Open the Windows Control Panel.

   c. Choose **System**.

   d. Click the **Advanced** tab.

   e. Click **Environment Variables** .

   f. Click **New** under **System Variables**.

   g. Set the variable information in the following way:

   Variable name: JAVA_HOME

   Variable value: *Java (JDK) 1.5 install path*. For example, the install path might be C:\Program Files\Java\jdk1.5.0_11.

   Important: The Content Search Engine installation program will not allow you proceed until the JAVA_HOME environment variable is set.

4. Set the PATH environment variable in the following way:

   a. Open the System control panel.

   b. Click the **Advanced** tab.

   c. Click **Environment Variables**.

   d. Click **PATH** under **System Variables**.

   e. Append the following variable information: ;%JAVA_HOME%\bin

Configuring the network

You must perform certain configurations on the network before installing IBM FileNet P8 Platform.

"Prerequisites to configuring your network"
Perform the following prerequisite tasks in any order.

"Synchronizing time and date" on page 21
System users will experience a variety of problems if one or more servers are not synchronized with the rest of the system.

Prerequisites to configuring your network:

Perform the following prerequisite tasks in any order.

- Assign all IBM FileNet P8 servers a static IP address.
- Ensure TCP/IP settings. Verify TCP/IP configuration settings on all UNIX and Windows servers and Enterprise Manager clients intended for IBM FileNet P8 so that they can all communicate with one another.
- Ensure NetBIOS over TCP/IP is enabled on Windows.
- Ensure availability of required port numbers. Several port numbers are required by the various IBM FileNet P8 components.

Synchronizing time and date:

System users will experience a variety of problems if one or more servers are not synchronized with the rest of the system.

The Process Engine database server (the machine that hosts the database used by Process Engine) is considered the master time keeper; the UTC time of that machine is considered the correct time. The server hosting the Process Engine API and the server hosting Content Engine must have the UTC time set.

- To change the time on the machine hosting Process Engine, you must stop the server. In a farmed Process Engine system, if you want to change the time of one of the servers in the farm, you must stop only that server.

- To change the time in the machine hosting the Process Engine API, be sure it is not connected to any Process Engine system. If the API is connected to a Process Engine server, and you change the time, you will experience authentication errors, and you might need to log on again.

- If your Content Engine server is being used with a Process Engine server, and you change the time on the Content Engine server, you will experience authentication errors in Process Engine and you might need to log on again.

Preparing storage areas for object stores

To prepare for file storage, you must prepare the locations where file stores will go, configure remote access protocol, and, if you plan to implement content-based retrieval, perform security configurations.

An object store can have up to three types of storage areas for the content of documents and business objects:

file storage area
> A file storage area stores content in a network-accessible directory. The path name to this directory specifies the location of the file storage area.
>
> For information about file storage areas, see the IBM FileNet P8 help topic **System Administration** → **Content Engine Administration** → **Content storage** → **File storage areas**.

fixed storage area
> A fixed storage area is a file storage area on a large-capacity, (possibly) write-once, fixed content device.
>
> For information about fixed storage areas, see the IBM FileNet P8 help topic **System Administration** → **Content Engine Administration** → **Content storage** → **File storage areas**.

database storage area
> A database storage area stores content as binary large objects (BLOBs) in a database.

A *file storage area* refers only to a network-accessible directory that is not on a fixed content device. The names of file storage areas and database storage areas must be unique within an object store. File storage areas on encrypted NTFS devices are not supported. However, the Decru, Vormetrics, and IBM Encryption Expert hardware-based encryption solutions are supported.

By default, the Create Object Store wizard creates a database storage area. If your object stores will use database storage areas only, you can skip the rest of this task, provided that one of the following conditions is met:

- Your database type is non-DB2.

- Your database type is DB2 and your database storage areas will not contain large content elements (larger than 300 MB, for example).

Besides creating a database storage area, the Create Object Store wizard allows you to create an initial file storage area or initial fixed storage area. But the wizard requires that you first do at least one of the following, depending on the type of storage areas you want for your object stores:

- For fixed storage areas, create at least one fixed content device (typically via Enterprise Manager). Multiple fixed storage areas can share the same fixed content device, or a fixed storage area can have its own fixed content device.

  If the content of all your object stores will be in fixed storage areas only, create your fixed content devices now, and skip the rest of this topic.

  To create a fixed content device, refer to the procedures in IBM FileNet P8 help topic **System Administration** → **Content Engine Administration** → **Content storage** → **File storage areas**.

- For file storage areas, prepare locations on one or more file servers (which usually are not a machine where you installed Content Engine), as shown in the remainder of this task.

  "Configuring file servers for file storage areas"
  You must configure file servers for the initial file storage areas of the object stores to be created, and for additional file storage areas of existing object stores.

  "Users and groups" on page 23
  The following table shows the operating system users and groups on the machine where Content Engine is to be deployed that are involved in securing file storage areas. These users and groups must be defined in the directory service that the operating system uses to authenticate users, which is not necessarily the same directory service that Content Engine Server uses.

  "Configuring the remote access protocol on the client machine" on page 25
  When configuring the remote file access protocol (NFS or CIFS), the client machine is the one where Content Engine Server or Content Search Engine are running. Configuring the remote access protocol (NFS or CIFS) means designating a directory (where content is be stored) so that it appears to be on the a local file system of the client machine.

  "Preparing storage areas to support content search functionality" on page 25
  To prepare your system for index area creation, each file storage area that will be full-text indexed must be accessible by both the Content Engine and the Autonomy K2 server that will perform the full-text indexing.

Configuring file servers for file storage areas:

You must configure file servers for the initial file storage areas of the object stores to be created, and for additional file storage areas of existing object stores.

Refer to the *IBM FileNet P8 Hardware and Software Requirements* for currently supported operating systems for file servers.

Configuring a file server for file storage areas involves the following general steps, which are described in more detail in the procedures later in this task.

To configure file servers for file storage areas:

1. Create or designate an existing top-level directory on the file server where file storage areas will reside.

2. Secure the directory so only Content Engine Server and Content Search Engine can access it.
3. Expose the directory via the remote file access protocol that applies to the operating system of the file server.
4. (Best practice) Under the top-level directory, create a subdirectory for each file storage area you intend to create. If you decide to put a file storage area directly within a top-level directory, rather than in a subdirectory, and you later decide to create an additional file storage area on this file server, you will have to create another top-level directory for it, because you will not be able to use the previously created top-level directory.

   "Remote file access protocols"
   The supported remote file access protocols between Content Engine and a file server are: Common Internet File System (CIFS), Network File System (NFS), and Distributed File System (DFS). DFS is supported if you are using it to manage a file storage area; however, the replication feature of DFS is not supported.

Remote file access protocols:

The supported remote file access protocols between Content Engine and a file server are: Common Internet File System (CIFS), Network File System (NFS), and Distributed File System (DFS). DFS is supported if you are using it to manage a file storage area; however, the replication feature of DFS is not supported.

The communication method between the Content Engine machine and the file server depends on the operating systems running on the two machines, as shown in the following table:

Content Engine Operating System	File Server Operating System	File Access Protocol
Windows 2003	Windows 2003	CIFS
UNIX	UNIX	NFS
UNIX	Windows 2003	NFS

Install a UPS power supply backup system on each file server to enable graceful shutdown. Loss or corruption of data will result if a file server does not shut down gracefully.

Users and groups:

The following table shows the operating system users and groups on the machine where Content Engine is to be deployed that are involved in securing file storage areas. These users and groups must be defined in the directory service that the operating system uses to authenticate users, which is not necessarily the same directory service that Content Engine Server uses.

The user and group account variables in this table are placeholders for the actual account names that you designate.

Users and Groups	Role
Content Engine operating system user (ce_os_user)	The user under which Content Engine Server runs (typically, the user that starts Content Engine Server).

Users and Groups	Role
K2 operating system user (k2_os_user)	The user under which Content Search Engine runs (typically, the user that starts Content Search Engine).
Content Engine operating system group (ce_os_group)	The group that contains: • Content Engine operating system user • K2 operating system user

For details on file storage area security, see the IBM FileNet P8 help topic **System Administration → Enterprise-wide Administration → FileNet P8 Security → Authorization → Storage area security**.

"Configuring a Windows-based file server for a Windows client using CIFS" You must configure security permissions on the directories where file storage areas are going to be located.

"Configuring a Windows-based file server for a UNIX client using NFS" To configure Windows Services for NFS, use the procedures in Microsoft documentation.

Configuring a Windows-based file server for a Windows client using CIFS:

You must configure security permissions on the directories where file storage areas are going to be located.

To configure a Windows-based file server:
1. Log on to the Windows file server as *ce_os_user*.
2. Create (or designate) a directory `fsa1` where content will be stored. For example:`C:\> md c:\filenet\file_stores\fsa1`
3. Navigate in Windows Explorer to `fsa1`, right-click the file icon, and choose **Properties**.
4. In the Security tab, click **Advanced**.
5. In the **Advanced Security Settings** dialog box:
   a. Grant Full Control to *ce_os_user* and *ce_os_group*, and select **This Folder, subfolders, and files** from the **Apply onto** drop-down list.
   b. Remove all other users and groups in the **Permission entries** table.
   c. Click **OK**.
6. In the Sharing tab, perform the following tasks:
   a. Click **Share this folder** and click **Permissions**.
   b. Grant Full control to *ce_os_user* and *ce_os_group*.
   c. Remove all other users and groups in the **Permission entries** table.
   d. Click **OK**.

Configuring a Windows-based file server for a UNIX client using NFS:

To configure Windows Services for NFS, use the procedures in Microsoft documentation.
1. Do all the steps in "Configuring a Windows-based file server for a Windows client using CIFS."
2. Use the procedures in Microsoft documentation to configure Windows Services for NFS to expose `fsa1`.

Tip:

- Windows Services for NFS is an optional Windows component bundled with Windows Server 2003 R2.
- As part of configuring Windows Services for NFS, you must set up a mapping of Windows users and groups to UNIX users and groups. When setting up the mapping for ce_os_user and ce_os_group, you must specify the same UID (UNIX user ID) and GID (UNIX group ID) that these accounts have on the machine where Content Engine Server is installed.

Configuring the remote access protocol on the client machine:

When configuring the remote file access protocol (NFS or CIFS), the client machine is the one where Content Engine Server or Content Search Engine are running. Configuring the remote access protocol (NFS or CIFS) means designating a directory (where content is be stored) so that it appears to be on the a local file system of the client machine.

To configure remote access protocol:

To configure Windows-based Content Engine Server to talk to a Windows file server via CIFS:

1. If both Content Engine Server and the file server are in the same Windows domain, no action is required. If they are in different domains, establish access to the file server machine from the machine where you will install Content Engine Server before creating the file storage areas.

Preparing storage areas to support content search functionality:

To prepare your system for index area creation, each file storage area that will be full-text indexed must be accessible by both the Content Engine and the Autonomy K2 server that will perform the full-text indexing.

Content Engine Operating System User (*ce_os_user*) must have permissions to the following directories:
- File storage area directories (read and write)
- Collections directory (read and write)
- Collections temp directory (read and write)

K2 Operating System User (*k2_os_user*) must have permissions to the following directories:
- File storage area directories (read)
- Collections directory (read and write)
- Collections Temp directory (read and write)

If Content Engine and Autonomy K2 are installed on separate machines, then the file store must be mounted and accessible to both the Content Engine and the Autonomy K2 servers. For Windows, the file store path must be a UNC path name accessible to the Autonomy K2 machine. For UNIX, the file store must be NFS-mounted on the Autonomy K2 machine with the same name as it appears on all other machines.

Related tasks

"Creating Content Engine operating system accounts" on page 10
The IT Administrator must create several operating system accounts required by
FileNet P8 components during installation.

"Creating Content Search Engine accounts" on page 36
If you are installing Content Search Engine, you must create new or designate
existing Autonomy K2 security accounts.

Security administrator installation tasks

The Security administrator must prepare the security environment for IBM FileNet
P8 Platform, including planning the security environment, configuring the
directory server, and creating accounts.

Review all rows assigned to the Security administrator (SA) in the Installation and
Upgrade Worksheet. While you complete the following preparation tasks, provide
values for the rows that are appropriate to your installation.

With the **Data → Filter → AutoFilter** command enabled, as it is by default in the
shipping worksheet file (p8_worksheet.xls), perform the following actions to
quickly see only the properties assigned to a particular Role:

- Click the **AutoFilter** drop-down arrow in the **Role** column header and select
  SA.
- Further filter the result set by clicking the **AutoFilter** drop-down arrow in any
  of the other columns and selecting a value or clear a filter by selecting **All**.

  "Security planning considerations"
  Information in this section is provided to assist in the security planning process
  but is not a complete description of any security feature or level of support.

  "Configuring directory server" on page 28
  The Security administrator must perform certain configurations on the directory
  server that will provide the authentication repository for your IBM FileNet P8
  system.

  "Specifying IBM FileNet P8 accounts" on page 29
  Included in this topic you will find a set of tables that describe all the accounts
  you must specify to set up IBM FileNet P8 components.

Security planning considerations

Information in this section is provided to assist in the security planning process
but is not a complete description of any security feature or level of support.

For complete information about IBM FileNet P8 security, consult the IBM FileNet
P8 help topic **FileNet P8 Administration → Enterprise-wide Administration →
FileNet P8 Security**. The Security Help sections that are especially applicable to
installation are Authentication, Directory Service Providers, and Users and Groups.

Authentication and authorization are separate processes.
Authentication (logon security) is separate from authorization (object and
process security). You must configure your JAAS login on the Content
Engine application server so that any user or group that can successfully
log on to IBM FileNet P8 resources can also be authorized to work within
IBM FileNet P8 interfaces, using the Content Engine directory service
provider connection.

Configuration Manager captures configuration information to create your
application server authentication provider; or you can use an
authentication provider that already exists on the application server.

Immediately following the initial Content Engine deployment, you will use Enterprise Manager to configure the Content Engine authorization by creating a Directory Configuration.

Logins are done through JAAS.
IBM FileNet P8 uses Java Authentication and Authorization Service (JAAS) for authentication, which is a process that occurs between a J2EE client application, a J2EE application server, and one or more JAAS login modules. This process does not involve any IBM FileNet P8 code.

IBM FileNet P8 Platform uses JAAS for authentication only, not for authorization on stored objects. Also, it does not support Java Security Manager.

Determine single sign-on (SSO) requirements.
Content Engine ability to use JAAS-based authentication means that if a single sign-on (SSO) provider writes a JAAS LoginModule for a supported application server, then clients of IBM FileNet P8 applications hosted in that application server can use that SSO solution. Where necessary, this Information Center describes SSO configurations that you must make, for example to the web.xml file, but it does not provide specific instructions for installing or configuring your SSO provider. See the IBM FileNet P8 help topic **FileNet P8 Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Authentication** → **Single sign-on integrations via JAAS**. Also see *Single Sign-On Solutions for IBM FileNet P8* at ibm.com/redbooks for configuration information.

Determine Kerberos applicability.
You can use Kerberos for SSO authentication between IBM FileNet Enterprise Manager and Content Engine, provided you use Windows Active Directory as the directory server. See the IBM FileNet P8 help topic **FileNet P8 Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Authentication** → **Kerberos for Content Engine**.

For information, see the IBM FileNet P8 help topic **FileNet P8 Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Authentication**.

Decide how many authentication realms you require.
At least one authentication realm is required, which you create during an initial installation by running Configuration Manager's Configure LDAP task. For an explanation of how to configure multiple realms, see the IBM FileNet P8 help topic **FileNet P8 Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **How to** → **Configure multiple realms**.

Make sure that you have a directory service provider in place.
Authentication in IBM FileNet P8 Platform is provided by the following supported directory servers:

• Microsoft Windows Active Directory

You can find detailed reference information in the IBM FileNet P8 help topic **FileNet P8 Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Directory Service Providers**.

IBM FileNet P8 supports only homogenous directory server environments. In other words, a single IBM FileNet P8 domain can be configured to use only one of the supported directory servers.

Understand the users and groups required for IBM FileNet P8.

All general administrative users and groups needing access to IBM FileNet P8-based applications must reside in one of the supported directory servers. This *Plan and Prepare Your Environment for IBM FileNet P8* provides instructions for creating the administrative accounts required for installation and initial configuration. For detailed reference information, see the IBM FileNet P8 help topic **FileNet P8 Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Users and groups**.

Process Engine delegates authentication to Content Engine.

Start with the IBM FileNet P8 Platform 4.0.0 release, Process Engine no longer has a direct connection to a directory server for authentication purposes, as it did in earlier releases. Instead, it delegates authentication tasks to Content Engine.

(WebLogic only) Any WebLogic authentication provider must be dedicated to IBM FileNet P8.

For performance reasons, do not share any authentication provider used by WebLogic for deployed IBM FileNet P8 components with applications used for other purposes.

You can configure Content Engine to use email or UPN for login

You can assign the directory server's email attribute or, for Active Directory, the userPrincipalName (UPN) to be the user short name used for login. Instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* provide a link to a procedure that explains how to do this.

Configuring directory server

The Security administrator must perform certain configurations on the directory server that will provide the authentication repository for your IBM FileNet P8 system.

"Configuring Windows Active Directory"
You can configure Microsoft Windows Active Directory to be the directory service for IBM FileNet P8 Platform.

Configuring Windows Active Directory:

You can configure Microsoft Windows Active Directory to be the directory service for IBM FileNet P8 Platform.

For a complete list of IBM FileNet P8-supported Windows Active Directory features, refer to **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Directory service providers** → **Windows Active Directory**.

In a multi-domain Active Directory environment, a logon will fail for any account whose user name and password in a parent/child domain does not match those in a child/parent domain.

If you have an Active Directory failover configuration, you can configure IBM FileNet P8 to follow this failover sequence whenever Content Engine attempts to authorize an already authenticated user. You can do this during Content Engine installation while running the Create a Directory Configuration wizard, or at any time after. See **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Directory service providers** → **Windows Active Directory** → **Active Directory failover**.

Server Side Sorting (SSS) must be enabled. This is because Process Engine and other IBM FileNet P8 components call on Content Engine to perform searches using a sorted paging mechanism. Without SSS, you will experience errors such as in Workplace when retrieving a document, when trying to open Process Designer via Workplace, or when starting a Connection Point on Application Engine. Note that SSS is normally enabled by default but is sometimes disabled due to concerns with performance.

DNS forwarders provide external DNS lookup functionality. If you are working in an "isolated" network, a DNS forwarder is not required. However, if you want to access the Internet or other network resources, then a DNS forwarder pointing to a DNS server that serves the external resources (for example, the Internet) is required.

To enable DNS forwarders:
1. On the machine that is configured as the Windows DNS Server, log on with an account that can configure the DNS components.
2. Start DNS. For example, on Windows 2003, choose **Start** → **All Programs** → **Administrative Tools** → **DNS** .
3. Right-click the *your_computer_name* container and select **Properties**.
4. Select the **Forwarders** tab and verify the check box for **Enable forwarders** is selected. If this feature is grayed out (unavailable), you must reconfigure your DNS server.
5. If you selected the check box, add an appropriate IP address and click **OK**. This IP address can be the IP address of a DNS server that allows traffic to the Internet.

Specifying IBM FileNet P8 accounts

Included in this topic you will find a set of tables that describe all the accounts you must specify to set up IBM FileNet P8 components.

This task assumes that you have completed the tasks required for your directory server in "Configuring directory server" on page 28.

Although these accounts are collected here under the Security administrator tasks because of their relationship to overall system security concerns, you will notice mention of other administrators who are likely to have the actual responsibility to create the accounts and put them to use.

The following procedures direct you to create or designate the accounts needed to install and configure IBM FileNet P8. For a complete list of the user and group roles, accounts, and responsibilities required to install, configure, and maintain an IBM FileNet P8 system, see the IBM FileNet P8 help topic **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Users and groups**.

Accounts are referred to in documentation in the following ways:
- By a display name; for example, *Database User Name*. An account's display name is how the IBM FileNet P8 user interface, such as an installation program or dialog box, refers to the account. Many accounts have both a display name and a variable.
- By a variable designator; for example *ce_db_user*, using lower-cased italics and underscores. The variable is intended to show that you must designate your own account to act in the role described by the variable. Accounts that do not appear in an interface or configuration file will have only a variable designator.

If you see a reference to an account that you do not understand, search the *Plan and Prepare Your Environment for IBM FileNet P8* guide and find the account table that defines it.

"Creating Content Engine directory server accounts"
Create new or designate existing directory server installation accounts for Content Engine.

"Creating Process Engine directory server accounts" on page 33
Create new or designate existing directory server installation accounts for Process Engine, as shown in the following table.

"Creating Application Engine or Workplace XT accounts" on page 36
Create new or designate existing directory server accounts for Application Engine or Workplace XT.

"Creating Content Search Engine accounts" on page 36
If you are installing Content Search Engine, you must create new or designate existing Autonomy K2 security accounts.

Creating Content Engine directory server accounts:

Create new or designate existing directory server installation accounts for Content Engine.

The Security Administrator should perform this task.

To create directory server accounts:

Create the following users and groups:

User or group name	Description
Bootstrap user name: *ce_bootstrap_admin*	A directory service and application server account that is stored in the `CEMPBoot.properties` file that is archived in the Content Engine EAR file. Also known as Content Engine System User.
	Content Engine uses *ce_bootstrap_admin* to establish a connection with the application server, access the application server's JNDI tree, and look up the data sources for accessing the GCD.
	You should not use this account as an all-purpose administrative account. For example, if you had to log on to some other application using the *ce_bootstrap_admin* account and provided the wrong password several times, thereby exceeding the number of allowable login failures, this account could be locked out of the directory server, which would mean that Content Engine would not start.
	Restriction: If you are deploying Content Engine on an application server with federated user repositories and with multiple realms in your FileNet P8 domain, be sure that no two realms contain the same short name for this user; otherwise, this user will not be able to create the GCD.
	See the IBM FileNet P8 help topic **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Users and groups** for information on how *ce_bootstrap_admin* is the account used to create the initial FileNet P8 domain.

User or group name	Description
GCD Administrator: *gcd_admin*	A directory service account that has Full Control access to the domain object for Content Engine.
	The initial *gcd_admin* is created by Configuration Manager using the account entered into its Create Bootstrap Properties panel for the Bootstrap user (*ce_bootstrap_admin*).
	Log on as *gcd_admin* in order to:
	• Create the GCD by launching the Configure New Domain Permissions wizard the first time you start Enterprise Manager to establish the IBM FileNet P8 domain (see in the *IBM FileNet P8 Platform Installation and Upgrade Guide*).
	• Carry out administrative tasks for the FileNet P8 domain.
	For more information, see the IBM FileNet P8 help topic **System Administration → Enterprise-wide Administration → FileNet P8 Security → Users and groups**.
Object Store Administrator: *object_store_admin*, *object_store_admin_group*	A directory service account that can administer an object store by having Full Control access to it. You can also grant Full Control to an object store to group accounts, thereby making all members of the group object store administrators.
	Use Enterprise Manager's Create an Object Store wizard to specify which user or group accounts should be *object_store_admin_group* (see *inst_ce_create_obj_store* in the *IBM FileNet P8 Platform Installation and Upgrade Guide*). Each object store could have a different set of object store administrators, depending on your security design.

User or group name	Description
Directory service user: Active Directory (Referred to as "Directory service bind user name" in Configuration Manager.) *ce_service_user*	An Active Directory user account that Content Engine uses to connect to Active Directory.
	ce_service_user performs the following roles:
	• Acts as the bind user specified by the application server to search through realms to authenticate a user when the user logs in to a Content Engine client such as Workplace.
	• Acts as the user specified in the GCD that searches users and groups to authorize access to a specific FileNet P8 object after a user has been authenticated.
	Using Active Directory tools, grant *ce_service_user* at least the following permissions:
	• Read access rights (specifically the Read All Properties permission) to the forest-wide configuration directory partition and the domain directory partition in each desired domain in the Active Directory forest. Because Authenticated Users by default is a member of the Pre-Windows 2000 Compatible Access group which has these permissions, you will need to assign the permissions to *ce_service_user* only if the default is modified or Authenticated Users access rights are restricted.
	Provide the fully qualified distinguished name of *ce_service_user* as the directory service bind user name while running Configuration Manager and also when you run the Enterprise Manager Directory Configuration Wizard.

Creating Process Engine directory server accounts:

Create new or designate existing directory server installation accounts for Process Engine, as shown in the following table.

This task should be performed by the Security Administrator. Process Engine can be installed by either a domain or local user. If you choose to install and optionally run Process Engine as a domain user, create the users identified here as optional. You must also create local users and groups for Process Engine installation and the domain users must be added to the local groups. See "Creating Process Engine operating system accounts" on page 13 for procedures for creation of local users and groups.

Create the following users and groups:

User or Group	Description
An optional domain user used to install Process Engine and to start and stop Process Engine software and services (Windows): *fnsw* While your security requirements may dictate that installation be done by a domain user, local users and groups are still required on the Process Engine machine.	This optional domain user can be used to install Process Engine software and to start and stop Process Engine software and services. Even if this domain user is created, you must create a corresponding local user on the Process Engine. Process Engine Windows services are configured by default to start with the *fnsw* local user and permissions on the software directory on the Process Engine are set for the local user and groups. Record this value in the Installation and Upgrade Worksheet.Assign the *fnsw* domain user to the following local groups on the Process Engine: • *fnusr* • *fnadmin* • *fnop* • Administrators Set the following rights to this user in the local security policy on the machine where Process Engine will be installed: • Act as part of the operating system • Adjust memory quotas for a process • Log on as a service (if this domain user will also run the Process Engine Windows services) • Replace a process level token If you choose to install Process Engine while logged on as this user, and run Process Engine SQL scripts from the installation program, ensure that this user has privileges in the database to run the scripts.
The password for the optional *fnsw* domain user (Windows): *fnsw_password*	This password is for the *fnsw* domain user. Record this value in the Installation and Upgrade Worksheet. Process Engine installation program on Windows prompts for the *fnsw_password*. On Windows the IMS ControlService and Process Engine Services Manager service are configured to start with the *fnsw* local user and the *fnsw_password* that is entered during Process Engine installation. Set this password to never expire. Process Engine activity ceases if the *fnsw_password* expires.

User or Group	Description
An optional domain user for Process Engine (Windows): *pe_install_user_domain*	This is a domain user account that can be used to install Process Engine software. Set the following rights to this user in the local security policy on the machine where Process Engine will be installed: • Act as part of the operating system • Adjust memory quotas for a process • Log on as a service (if this domain user will also run the Process Engine Windows services) • Replace a process level token You can log on as the *fnsw* user or create this user or the *pe_install_user_local* user (with the specified permissions) expressly for the purpose of running the Process Engine installation program. If you choose to install Process Engine while logged on as this user, and run Process Engine SQL scripts from the installation program, ensure that this user has privileges in the database to run the scripts.
Process Engine service user: *pe_service_user*	Process Engine uses the *pe_service_user* when connecting to the Content Engine server. This user must belong to the Process Engine Administrator group. Record this value in the Installation and Upgrade Worksheet. The name will be required input when configuring the Process Task Manager.
Process Engine administrators group: *pe_admin_group*	Members of this group automatically have administrative privileges for Process Engine. Record this value in the Installation and Upgrade Worksheet. The name will be required input when configuring the Process Task Manager.
Process Engine configuration group: *pe_config_group*	(Optional) A valid group name. Members of this group automatically have configuration privileges for the Process Engine workflow database. If this group is used to configure security on Process Task Manager, members of this group or of the Process Engine Administrator Group can make configuration changes to the workflow database. If the Process Engine Configuration group is not used during this configuration, anyone can make these changes. Record this value in the Installation and Upgrade Worksheet. The name will be input when configuring the Process Task Manager.

Creating Application Engine or Workplace XT accounts:

Create new or designate existing directory server accounts for Application Engine or Workplace XT.

This task should be performed by the Application Engine Administrator.

To create Application Engine or Workplace XT accounts:

Create the following users and groups. All IBM FileNet Workplace accounts, as well as accounts for other client applications and expansion products that use Content Engine or Application Engine, must have passwords.

User or group name	Description
Application Engine or Workplace XT installer account (Windows):*ae_install_user*, *wpxt_install_user*	The account you will use to log on to a Windows machine and launch the Application Engine or Workplace XT installation program. This account must be a Windows Local administrator or a user with equivalent permissions.
Application server account: *ae_deploy_user*, *wpxt_deploy_user*	This account will have permissions to deploy an application. The account can be the same as the Application Engine or Workplace XT installer account.
Application Engine Administrators: *ae_admin_user*, *wpxt_admin_user*	These accounts will serve in the role of Application Engine administrator. The role applies to both Application Engine or Workplace XT. You will specify these accounts as members of the Application Engine administrator role when you set bootstrap preferences. These accounts must have passwords.

In addition to the requirements above, the installer account (*ae_install_user* or *wpxt_install_user*) and the Application server account (*ae_deploy_user* or *wpxt_deploy_user*) and need read/write/execute permission to these directories and files:

Installation paths

> Grant *ae_install_user* read and write permission to the *ae_install_path*.

> Grant *wpxt_install_user* read and write permission to the *wpxt_install_path*.

WebLogic 9.x

> *BEA_Home*/bea/user_projects/domains/*domain_name*/bin/startWebLogic.sh or startWebLogic.cmd

> *BEA_home*/bea/user_projects/domains/*domain_name*/config/config.xml

WebLogic 10.x

> *BEA_Home*/bea/wlserver_10.0/server/bin/startWLS.sh or start WLS.cmd

> *BEA_home*/bea/user_projects/domains/*domain_name*/config/config.xml

Creating Content Search Engine accounts:

If you are installing Content Search Engine, you must create new or designate existing Autonomy K2 security accounts.

The IT Administrator should perform this task.

To create accounts for Content Search Engine:

Create the following users and groups:

User or group name	Description

K2 Security Group: *k2_sec_group*	Autonomy K2 security group used to secure K2 collections. Later, you will specify this group in the **User Group** field of the Verity Domain Configuration tab when you configure CBR in the Enterprise Manager root domain property sheet.
	
	Windows
	   Active Directory group when Autonomy K2 is installed on a machine in a Windows domain.
	
	   Windows local group when Autonomy K2 is installed on a machine in a Windows workgroup.

K2 Security User: *k2_sec_user*	Autonomy K2 security user account, used by Content Engine when logging onto the Autonomy K2 Master Administration Server to perform Content-Based Retrieval (CBR). Later, you will specify this account in the **Verity Username** field in the Verity Domain Configuration when you configure CBR in Enterprise Manager. This user must be a member of the *k2_sec_group* and must be defined as an authorized K2 administrator in the K2 dashboard. The *k2_sec_user* and *k2_os_user* can be the same user. All permissions listed for both users must be assigned.
K2 Operating System User: *k2_os_user*	Windows services will run as this user. This user must be given administrator privileges for Windows on all machines on which the Autonomy K2 software will be installed. If Content Search Engine is to be in a Windows workgroup, *k2_os_user* must be given the *log on as a service* right in the Local Security Policy. If Content Search Engine is to be in a Windows domain, *k2_os_user* must be given the *log on as a service* in the Domain Security Policy. This user also requires read/write permissions to the K2 collections directory, the collections temp directory and to the file system that contains the file storage areas and full-text index collections. The Autonomy K2 software needs to read the file storage areas and write the full text index collections as part of the full text indexing operation.

Database administrator installation tasks

The Database administrator must prepare the databases required for IBM FileNet P8, including gathering information about data sources, creating databases and database accounts, and installing client software.

- Review all rows assigned to the Database administrator (DBA) in the Installation and Upgrade Worksheet. While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation. (Your organization might have different roles, and some of the responsibilities of listed roles will vary from those assigned by default in this documentation.)

    Tip: With the **Data** → **Filter** → **AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following actions to quickly see only the properties assigned to a particular Role:

    - Click the **AutoFilter** drop-down arrow in the **Role** column header and select **DBA**.
    - Further filter the result set by clicking the **AutoFilter** drop-down arrow in any of the other columns and selecting a value or clear a filter by selecting **(All)**.

    As an alternative, you can use the Customize Worksheet filtering macro, embedded in the worksheet file's Instructions tab.

- If you are installing in a non-English environment, review the considerations and procedures in Appendix A, "Preparing non-English environments for installing IBM FileNet P8 Platform," on page 57 before you begin your preparation tasks.

    "Creating Content Engine database accounts"
    Use your database tools to create new or designate existing database accounts for Content Engine, as shown in the following table.

    "Creating Process Engine database accounts" on page 39
    You must create new or designate existing database accounts for Process Engine.

    "Preparing Microsoft SQL Server" on page 40
    Plan the SQL Server installation and configuration, install the software, and configure database components for IBM FileNet P8 components after reviewing the requirements.

    "Configuring Microsoft SQL Server client for Process Engine" on page 48
    To configure your database clients for Process Engine, you must create ODBC data sources and, if required, install Microsoft SQL Server Client software for Process Engine.

    "Verifying the ability to connect to the database" on page 50
    Verify the ability to connect to the Process Engine database. These procedures can be run after initial configuration of the database and immediately prior to installation of Process Engine software. Run these steps on the database server or the client according to whether the database is local to or remote from Process Engine.

Creating Content Engine database accounts

Use your database tools to create new or designate existing database accounts for Content Engine, as shown in the following table.

Create the following users and groups:

User or group name	Description
Database user name (Microsoft SQL Server):*ce_db_user*	The database accounts that Content Engine uses to access SQL Server, depending on whether you use one account for all Content Engine databases, or use one (for example, *ce_db_user*) for the GCD database and different accounts for each object store (for example, *ce_db_user1* , *ce_db_user2* , and so on). *ce_db_user* should be a SQL Server account. It does not have to be an account in the configured directory service. Grant each *ce_db_user* at least the following database access permissions: • public • db_owner Add these accounts to SQL Server's master database and grant the public role to each. When you perform the procedure "Configuring the JDBC distributed transaction components" on page 44, these accounts will also be granted the SqlJDBCXAUser role.

Creating Process Engine database accounts

You must create new or designate existing database accounts for Process Engine.

You create database users on Oracle and SQL Server databases for Process Engine by running SQL scripts. The scripts create default user names and passwords, or they can be edited to change the user names and passwords for your installation. Database permissions are also granted from the SQL scripts.

Create the following users and groups:

User or group name	Description
The database runtime user for Process Engine (SQL Server): *f_sw*	The database runtime user for Process Engine. The user is defined when the CreatePEinstallSP_1.sql script runs. The default value is set to f_sw. The script can be edited to change the user name. Record the value for this user in the Installation and Upgrade Worksheet.
The password for the database runtime user for Process Engine (SQL Server): *f_sw_password*	The database runtime user's password for Process Engine. The password is set when the CreatePEinstallSP_1.sql script runs. The default value is set to filenet. The script can be edited to change the password. Record the value for this user in the Installation and Upgrade Worksheet.

User or group name	Description
The database maintenance user for Process Engine (SQL Server): *f_maint*	The database maintenance user for Process Engine. The user is defined when the `CreatePEinstallSP_1.sql` script runs. The default value is set to f_maint. The script can be edited to change the user name. Record the value for this user in the Installation and Upgrade Worksheet.
The password for the database maintenance user for Process Engine (SQL Server): *f_maint_password*	The database maintenance user's password for Process Engine. The password is set when the `CreatePEinstallSP_1.sql` script runs. The default value is set to change$this_obnoxiou$_password. The script can be edited to change the password. Record the value for this user in the Installation and Upgrade Worksheet.

Preparing Microsoft SQL Server

Plan the SQL Server installation and configuration, install the software, and configure database components for IBM FileNet P8 components after reviewing the requirements.

"Microsoft SQL Server planning considerations"
Determine whether Microsoft SQL Server database components will be dedicated to individual IBM FileNet P8 components or shared, and whether the database will be local or remote, after reviewing the requirements

"Verifying that Microsoft SQL Server is ready for IBM FileNet P8" on page 41
To prepare your databases before installing IBM FileNet P8, you must install the SQL Server software and configure the database components for your installation.

"Running Process Engine pre-installation scripts (SQL Server)" on page 44
Run SQL scripts to create Process Engine database users, passwords, and stored procedures.

Microsoft SQL Server planning considerations:

Determine whether Microsoft SQL Server database components will be dedicated to individual IBM FileNet P8 components or shared, and whether the database will be local or remote, after reviewing the requirements

In a shared configuration, the IBM FileNet P8 components use the same instance. In a dedicated configuration, Content Engine, Process Engine, and Rendition Engine use separate instances.

The following components can optionally share a SQL Server instance:
- Content Engine (object store databases and GCD database)
- Process Engine
- Rendition Engine

Process Engine and the optional Process Analyzer cannot share an instance. For information aboutSQL Server requirements for Process Analyzer Engine, see the *IBM FileNet Process Analyzer Installation and Upgrade Guide*

SQL Server databases can be dedicated or shared by one or more of the following IBM FileNet P8 components:

- Dedicated to Content Engine
- Dedicated to Process Engine
- Shared by two or more of Content Engine, Process Engine, and Rendition Engine
- Shared with other applications

A database should not be shared by multiple Content Engine object stores. Backup and support activities are simplified if each object store has a separate database.

For a SQL Server database dedicated to Rendition Engine, see the *IBM FileNet Rendition Engine Installation and Upgrade Guide*.

A database is local if it is on a server where you will also be installing Content Engine, Process Engine, or Rendition Engine. A database is remote if it is on a separate server from the component that uses that database.

Process Engine requires an ODBC data source for connection to the SQL Server database whether the database is local or remote. See "Configuring Microsoft SQL Server client for Process Engine" on page 48 for details on configuring the data source.

Determine when to run SQL scripts for Process Engine. You must run a number of SQL scripts to create a stored procedures. These stored procedures are then used to create database users and other stored procedures required for Process Engine production. The Process Engine installation program will complete only if all these scripts run successfully. Run the scripts manually, before starting Process Engine installation, or run them from the Process Engine installation program. See "Running Process Engine pre-installation scripts (SQL Server)" on page 44 for information about execution modes and associated security requirements. If you run these scripts from the Process Engine installation program, install the SQL Server client software on the Process Engine

Determine the maximum size of the content elements your users store. The size affects setting up database storage areas or file storage areas. When you create an object store, a database storage area is provided by default where you can store content as database BLOBs. You can also create one or more file storage areas to store content on local or remote file systems. If your users store large individual documents or other content elements, use only file storage areas. Otherwise, users can encounter memory-related errors when retrieving or indexing the large content. Controlled tests with limited concurrency exhibited errors when run with files that were 300 MB or larger. Factors affecting this file-size limitation include driver and application server memory demands, other activity such as concurrent retrieval, indexing of large content, and JVM memory allocations.

Verifying that Microsoft SQL Server is ready for IBM FileNet P8:

To prepare your databases before installing IBM FileNet P8, you must install the SQL Server software and configure the database components for your installation.

"Installing and configuring Microsoft SQL Server" on page 42
Install and configure SQL Server software and create database instances for Content Engine and Process Engine.

"Creating a Microsoft SQL Server database for the GCD" on page 43
Create a Microsoft SQL Server database for the GCD database, which is required for Content Engine installation.

"Creating a Microsoft SQL Server database for Content Engine object stores" on page 43
Create a Microsoft SQL Server database for a Content Engine object store. Each object store you create requires an empty database. Create the database with an initial size of 200 MB minimum.

"Creating a Microsoft SQL Server database for Process Engine" on page 43
Create a Microsoft SQL Server database and filegroup for Process Engine.

"Modifying the tempdb database for Process Engine" on page 43
Verify that the space allocated for the tempdb is at least 80 MB and that the collation settings match the Process Engine database.

"Enabling XA transactions" on page 43
Configure the Windows server to enable XA transactions.

"Configuring the JDBC distributed transaction components" on page 44
You must enable XA transactions on every Microsoft SQL Server that will have a Content Engine database.

Installing and configuring Microsoft SQL Server:

Install and configure SQL Server software and create database instances for Content Engine and Process Engine.

Create and configure one or more database instances using these requirements.
1. Create one or more database instances for use by IBM FileNet P8 software or verify that such instances already exist.
2. If you create a new instance, indicate an appropriate name based on whether Content Engine (object store), GCD, Process Engine, or Rendition Engine will use the instance.
3. Verify that the authentication mode is Mixed Mode.
4. Select the database collation settings: Choose one of the following from the collation options:
   - Dictionary order, case-insensitive, for use with 1252 Character Set (or any case-insensitive SQL Server collation). Case-insensitive collation is the Microsoft default and the setting most used in IBM FileNet P8 environments (because it offers search results without regard to character case).
   - Dictionary order, case sensitive, for use with 1252 Character Set (or any case-sensitive SQL Server collation). Select case-sensitive SQL Server collation only if your site requires (and will continue to require) searches that must differentiate uppercase from lower-case characters. If you plan to use the Content Engine with CFS/IS, you must configure case sensitive. The FileNet Image Servicesdatabase is configured as case sensitive and the Content Engine database must match.

   Important: Select your SQL Server collation setting carefully. Switching collation settings after installation can be difficult and time-consuming, especially if you want to switch from case-sensitive to case-insensitive collation after significant user activity. If you have a case-sensitive database and you want to perform a case-insensitive search (programmatically or otherwise), you might encounter serious performance degradation on SQL Server because the database cannot use column, or property, indexes in these cases.
5. See *IBM FileNet P8 Hardware and Software Requirements* for required operating system and database patch sets and service packs. Verify that the required service pack is installed before proceeding.

6. Optional: If you want to disable the Named Pipes protocol for the database instance to be used by Process Engine, wait until after the Process Engine installation and configuration is complete. Disabling this protocol too early might cause Process Engine initialization to fail.

7. Record the following values in the Installation and Upgrade Worksheet:
   - The database server name, which is needed during Content Engine installation and configuration for the GCD and object stores. It is also needed to create the ODBC data source for Process Engine.
   - The database instance name, which is needed to create the ODBC data source for Process Engine.

Creating a Microsoft SQL Server database for the GCD:

Create a Microsoft SQL Server database for the GCD database, which is required for Content Engine installation.

1. Create the database with a minimum size of 100 MB.
2. Record the GCD database name and the TCP/IP port number in the *Installation and upgrade worksheet*.

Creating a Microsoft SQL Server database for Content Engine object stores:

Create a Microsoft SQL Server database for a Content Engine object store. Each object store you create requires an empty database. Create the database with an initial size of 200 MB minimum.

Record the object store database name and the TCP/IP port number in the *Installation and upgrade worksheet*.

Creating a Microsoft SQL Server database for Process Engine:

Create a Microsoft SQL Server database and filegroup for Process Engine.

Create a database and filegroup for the Process Engine. Do not use the primary filegroup.

1. Create a database with a minimum size of 200 MB.
2. Assign a name to the database. The default name assigned by the Process Engine installation program is VWdb.
3. Assign a new file name for the database, for example, vw_data.
4. Specify a filegroup, for example, vwdata_fg. Do not use the Primary filegroup.
5. If you will be using the region recovery feature, create an additional filegroup for every region that is configured for recovery.
6. Record the values for the Process Engine database, database password, and filegroup in the *Installation and upgrade worksheet.*

Modifying the tempdb database for Process Engine:

Verify that the space allocated for the tempdb is at least 80 MB and that the collation settings match the Process Engine database.

Enabling XA transactions:

Configure the Windows server to enable XA transactions.

Perform these steps on every Microsoft SQL Server that will contain a Content Engine database.

1. From Control Panel, open Administrative Tools, and then open Component Services.
2. Expand Component Services, right-click **My Computer**, and then select Properties.
3. (Windows 2008 only) Expand **Distributed Transaction Coordinator** and right-click **Local DTC**.
4. Click the **MSDTC** tab, and then click **Security Configuration**.
5. Select the **Enable XA Transactions** check box, and then click **OK** to restart the Microsoft DTC service.
6. Click **OK** again to close the Properties dialog box, and then close Component Services.
7. Stop and then restart the Microsoft SQL Server.

Configuring the JDBC distributed transaction components:

You must enable XA transactions on every Microsoft SQL Server that will have a Content Engine database.

Take the following steps to enable XA transactions for Content Engine databases.

1. Download the Microsoft SQL Server JDBC Driver that is referenced in the *IBM FileNet P8 Hardware and Software Requirements* document for Content Engine SQL Server databases.
2. Copy the `sqljdbc_xa.dll` from the JDBC installation directory to the `binn` folder of the instance, although a pre-2.0 version of the driver also functions correctly from the `tools\binn` folder. If you are on a 32-bit processor, use the *sqljdbc_xa.dll* file in the x86 folder. If you are on a 64-bit processor, use the *sqljdbc_xa.dll* file in the x64 folder.
3. Log on as the sa administrator or as a user with equivalent permissions and execute the database script *xa_install.sql* on the master database on every SQL Server instance that will participate in distributed transactions.

   Important: Use SQL Server database credentials, not Windows credentials, to log on. Windows Integrated Logon to SQL Server is not supported with IBM FileNet P8.
   This script installs *sqljdbc_xa.dll* as an extended stored procedure and creates the SqlJDBCXAUser role in the Master database.
4. Add each database account (*ce_db_user*) that Content Engine uses to access SQL Server to the SqlJDBCXAUser role. This action grants permissions to those accounts to participate in distributed transactions with the JDBC driver.
5. Click **Component Services > My Computer > Properties > MSDTC > Security Configuration** and ensure that XA transactions are enabled.

Running Process Engine pre-installation scripts (SQL Server):

Run SQL scripts to create Process Engine database users, passwords, and stored procedures.

The scripts to create Process Engine database users, passwords, and stored procedures for Microsoft SQL Server databases can run in several ways.

- Run the scripts manually on the database server, before you run the Process Engine installation program. You can change the default runtime and maintenance users and passwords in the scripts before you run the scripts.
- Run the scripts automatically from the Process Engine installation program, allowing the installation program to prompt for the sa password for SQL Server. You can change the default runtime and maintenance users and passwords in the scripts before you run the Process Engine installation program.
- Run the scripts automatically from the Process Engine installation program by using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local database.

  "Scenarios for manually running SQL scripts for a Process Engine Microsoft SQL Server database"
  Database user names and passwords are defined in SQL scripts and the same user names and passwords must be entered during Process Engine installation. Several scenarios are used to describe setting either default or custom values for the user names and passwords.

  "Process Engine SQL scripts for Microsoft SQL Server" on page 46
  All the SQL scripts are located in the root of the Process Engine installation directory. The following scripts are for new Process Engine installations with SQL Server databases.

  "Editing SQL scripts and changing default passwords (SQL Server)" on page 46
  Edit SQL scripts to change default passwords for the SQL Server database runtime and maintenance users for Process Engine.

  "Running SQL scripts manually for a new installation (SQL Server)" on page 47
  Run the preinstallation scripts and optionally modify the runtime and maintenance user names and passwords.

Scenarios for manually running SQL scripts for a Process Engine Microsoft SQL Server database:

Database user names and passwords are defined in SQL scripts and the same user names and passwords must be entered during Process Engine installation. Several scenarios are used to describe setting either default or custom values for the user names and passwords.

Scenario 1: Run SQL scripts manually before installing by using all default users and passwords

To run the required SQL scripts manually before running the Process Engine installation program. All default users and passwords are defined in the scripts.

1. Before you install Process Engine, run the scripts manually on the database server. Do not change user names or passwords in the scripts.
2. When you install Process Engine, use default users for the database accounts. Do not set aliases.
3. Leave the f_sw and f_maint password fields blank to use the default values that were assigned when the scripts ran.
4. Reset the f_sw and f_maint passwords by running the Xdbconnect utility. Running Xdbconnect changes both the encrypted version of the password and the password in the database. Xdbconnect works only if the passwords in the encrypted file and the database match after installation.

Scenario 2: Run SQL scripts manually before installing with nondefault users and passwords

To run the required SQL scripts manually before running the Process Engine installation program, but nondefault users and passwords are set:

1. Edit the scripts to change the user names and passwords.
2. Run the scripts manually on the database server.
3. Install Process Engine and specify that aliases will be configured. The user names that are set when the scripts ran must be defined as the aliases for the f_sw and f_maint users. None of the fields for alias names can be left blank in the Process Engine installation window. However, you can use default values as needed.
4. Set the f_sw and f_maint passwords during Process Engine installation to the match the passwords that were set when the scripts were run manually.
5. Do not reset passwords by running Xdbconnect. Because nondefault passwords were used, there is no need to change them immediately after installation.

Process Engine SQL scripts for Microsoft SQL Server:

All the SQL scripts are located in the root of the Process Engine installation directory. The following scripts are for new Process Engine installations with SQL Server databases.

Script name	Action
CreatePEinstallSP_1.sql	Creates the PE_createDbUsers stored procedure in the Process Engine database.
CreatePEinstallSP_2.sql	Creates a stored procedure that is called during Process Engine installations.
CreatePEinstallSP_3.sql	Creates the fn_error stored procedure.

The Process Engine scripts run with the following options:

run_time_user
> This option value is one of the following:
> - f_sw - if you use the default database users when you run the Process Engine installation program.
> - alias for f_sw - if you define aliases for the operating system and database users when you run the Process Engine installation program.

maintenance_user
> This option value is one of the following:
> - f_maint - if you use the default database users when you run the Process Engine installation program.
> - alias for f_maint - if you define aliases for the operating system and database users when you run the Process Engine installation program.

PE_database_name (SQL Server)
> The database name that is entered during Process Engine installation.

DSN (SQL Server only)
> The ODBC data source name that is entered during Process Engine installation.

Editing SQL scripts and changing default passwords (SQL Server):

Edit SQL scripts to change default passwords for the SQL Server database runtime and maintenance users for Process Engine.

The Process Engine database runtime and maintenance users and their passwords are defined by running a series of scripts. The scripts can optionally be edited to change the default values to meet your requirements. Both passwords are defined in the `CreatePEinstallSP_1.sql` script and can be changed using the steps in this topic.

1. Open the `CreatePEinstallSP_1.sql` script with a text editor.
2. Modify the following line to change the default password for the user called run_time_user where 'filenet' is the default password. Passwords must be enclosed in single quotation marks.

   ```
   set @passwd1 = 'filenet'
   ```

3. Modify the following line to change the default password for the maintenance user where 'change$this_obnoxiou$_passwrd' is the default password.

   ```
   set @passwd2 = 'change$this_obnoxiou$_passwrd'
   ```

Running SQL scripts manually for a new installation (SQL Server):

Run the preinstallation scripts and optionally modify the runtime and maintenance user names and passwords.

Run the three SQL scripts to create users, passwords, and stored procedures. Optionally, modify the default user names and passwords created by the scripts. Record the values for the user names and passwords in the Installation and Upgrade Worksheet to refer to when installing Process Engine software.

1. Copy the `CreatePEinstallSP_1.sql`, `CreatePEinstallSP_2.sql`, and `CreatePEinstallSP_3.sql` script files from the Process Engine software package to the local disk on the SQL Server database server.
2. Create a text file named `input.txt` to identify the runtime and maintenance users. The file contains the following text:

   ```
   PE_createDbUsers 'run_time_user', 'maintenance_user', 'PE_database_name'
   ```

   The following is an example of the contents of the `input.txt` file.

   ```
   PE_createDbUsers 'f_sw', 'f_maint', 'VWdb'
   ```

   In this example:
   - The run_time_user is the default f_sw
   - The maintenance_user is the default f_maint
   - The database is VWdb

3. Edit the `CreatePEinstallSP_1.sql` script file to change the default passwords. Locate and modify the following two lines:

   ```
   set @passwd1 = 'filenet'
   set @passwd2 = 'change$this_obnoxiou$_passwrd'
   ```

   where **@passwd1** is the runtime user's password and **@passwd2** is the maintenance user's password.

4. Run the `CreatePEinstallSP_3.sql` script first, then run the `CreatePEinstallSP_2.sql` script, and then run the `CreatePEinstallSP_1.sql` script. The first example shows execution of the `CreatePEinstallSP_3.sql` script on the database server.

   ```
   osql -E -d instance_name/PE_database_name
   -i CreatePEinstallSP_3.sql -n -o output3.log
   ```

Instead of using the parameter **-E**, you can use the following parameter:

```
-U sa -P sa password
```

The next example shows execution of `CreatePEinstallSP_2.sql` and `CreatePEinstallSP_1.sql` from a Process Engine server that is remote from the database server.

```
osql -D DSN -i CreatePEinstallSP_2.sql -n -o output2.log
osql -D DSN -i CreatePEinstallSP_1.sql -n -o output1.log
```

The **-D** *DSN* parameter is the ODBC data source name. You do not need this parameter if the scripts are running on the database server.

The **-d** parameter is the database instance name and database name for Process Engine. If the default instance is being used, it does not need to be included in this parameter.

The **-E** or alternative **-U** and **-P** parameters are not required if the scripts are running remotely from the database server.

If the scripts run successfully, nothing is logged in the `output3.log`, `output2.log`, and `output1.log` files.

5. Run the script to create the runtime and maintenance users, using the `input.txt` file created earlier. The passwords set in the `CreatePEinstallSP_1.sql` script file will be assigned to the users. This example shows execution of the script from a Process Engine that is remote from the database server.

```
osql -D DSN -i input.txt
```

Configuring Microsoft SQL Server client for Process Engine

To configure your database clients for Process Engine, you must create ODBC data sources and, if required, install Microsoft SQL Server Client software for Process Engine.

"Creating the Process Engine ODBC data source and testing the connection" Process Engine connections to a Microsoft SQL Server database at run time are handled through an ODBC data source. The 32-bit ODBC data source is required for both local and remote databases and must be created on the Process Engine server.

"Installing Microsoft SQL Server client software for Process Engine" on page 49 Install Microsoft SQL Server client software if the Process Engine database is remote and if preinstallation scripts will be executed from the Process Engine installation program.

Creating the Process Engine ODBC data source and testing the connection:

Process Engine connections to a Microsoft SQL Server database at run time are handled through an ODBC data source. The 32-bit ODBC data source is required for both local and remote databases and must be created on the Process Engine server.

Create a 64-bit ODBC data source if the preinstallation SQL scripts run from the Process Engine installation program on a 64-bit Windows server. This is in addition to the 32-bit ODBC data source required for Process Engine run time software.

The following steps apply on a 32-bit version of the operating system.

On a 64-bit operating system, locate and manually run the 32-bit version, which is typically located in `C:\WINDOWS\SysWOW64\odbcad32.exe`.

To create the Process Engine ODBC data source and test the connection:

1. Click **Start** → **Program** → **Administrator Tools** → **Data Source (ODBC)**.

2. Click **Add** on the System DSN page.

3. Select **SQL Server** as the driver to use for the new data source and click **Finish**.

4. Enter a name and description for the data source. The name is required input for the Process Engine installation program when configuring for a SQL Server database.

5. Choose the SQL Server to connect to from the list of servers and click **Next**. If you see only a server name in the list, the connection is to the default instance. If there are named instances in the database, the server/instance name is in the list.

6. Click **Next**

7. Choose **SQL Server authentication**.

8. Select the option to get default settings for additional configuration options by connecting to the SQL Server.

9. Enter the login ID and Password to connect to the database. The database login ID information does not need to be for an administrator, and it is only used to connect to the database to get the default values for the remaining settings required to configure the data source.

10. Click **Next**

11. Change the default database to be the Process Engine database that was created earlier in "Microsoft SQL Server planning considerations" on page 40.

12. Enable **Use ANSI null, paddings, and warning** and enable **Use ANSI quoted identifiers**.

13. Enable **Perform translation for character data** and click **Finish**.

14. Verify the settings for the data source configuration and click **Test Data Source**. If the test is successful click **OK**. Otherwise resolve any problem before continuing.

15. Double-click **SQL Server** on the Connection Pooling page.

16. Select **Don't pool connection to this driver** and click **OK**.

17. Click **OK** on the ODBC Data Source Administrator window to finish configuration of the data source.

18. On the summary window, click **Test Data Source**.

Installing Microsoft SQL Server client software for Process Engine:

Install Microsoft SQL Server client software if the Process Engine database is remote and if preinstallation scripts will be executed from the Process Engine installation program.

You must run a number of SQL scripts for the SQL Server database to create database users and stored procedures.

- If the scripts are executed manually on the remote database server before you install Process Engine, you do not need to install SQL Server client.

- If the scripts are executed from the Process Engine installation program, SQL Server client software must be installed on the Process Engine server.

- If SQL Server client software is installed on the Process Engine Server, that software can be removed after successful installation and configuration of Process Engine software.

To install Microsoft SQL Server client software:

1. Log on to the Process Engine server with an account that has local administrator privileges.

2. Install **Workstation components, Books Online and development tools** and from the **Advanced options**, select **Client Components**.

3. When the installation is complete, start the **Client Network Utility** and clear the **Automatic ANSI to OEM conversion** check box on the DB-Library Options page.

Verifying the ability to connect to the database

Verify the ability to connect to the Process Engine database. These procedures can be run after initial configuration of the database and immediately prior to installation of Process Engine software. Run these steps on the database server or the client according to whether the database is local to or remote from Process Engine.

"Verifying the Process Engine database connection (Microsoft SQL Server)" You need to verify that the Microsoft SQL Server database instance used by Process Engine is accessible. You must know both the Process Engine database and filegroup names. Make whatever corrections are necessary before proceeding.

Verifying the Process Engine database connection (Microsoft SQL Server):

You need to verify that the Microsoft SQL Server database instance used by Process Engine is accessible. You must know both the Process Engine database and filegroup names. Make whatever corrections are necessary before proceeding.

1. Log on to the Process Engine as a member of the local Administrators group or a user with equivalent permissions. The user you log on as must also be a database administrator. If the database is remote, the SQL connection must also be a trusted connection.

2. At a command prompt, enter:

   ```
   osql -E -D DSN
   ```

   where *DSN* is the ODBC data source name

   This command puts Process Engine into osql interactive mode.

3. At the osql prompt, enter:

   ```
   1> use VWdb
   2> go
   ```

   where *VWdb* is the Process Engine database name

   This command verifies that the Process Engine database has been created. If you get another prompt with no error, you are attached to that database.

4. Verify that the correct Process Engine filegroup was created. At the osql prompt, enter:

   ```
   1> select substring(groupname,1,20) from sysfilegroups where groupname=
      'defined filegroup'
   2> go
   ```

   where *defined filegroup* is the default filegroup

   A listing of the Process Engine filegroups will display, for example:

   ```
   vwdata_fg
   ```

5. Exit osql by typing **exit**.

Application Server administrator installation tasks

The Application Server Administrator must prepare the application servers for IBM FileNet P8 Platform, including planning deployment, creating administrative accounts, and configuring JDBC drivers for both Content Engine and Application Engine.

- Review all rows assigned to the Application Server Administrator (ASA) in the Installation and Upgrade Worksheet. While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation.

  Tip: With the **Data** → **Filter** → **AutoFilter** command enabled, as it is by default in the shipping worksheet file (`p8_worksheet.xls`), perform the following actions to quickly see only the properties assigned to a particular Role:
  - Click the **AutoFilter** drop-down arrow in the **Role** column header and select ASA.
  - Further filter the result set by clicking the **AutoFilter** drop-down arrow in any of the other columns and selecting a value or clear a filter by selecting (All).

- If you are installing in a non-English environment, review the considerations and procedures in Appendix A, "Preparing non-English environments for installing IBM FileNet P8 Platform," on page 57 before you begin your preparation tasks.

- Review the following planning considerations:
  - "Application server planning considerations" on page 4

  "Creating Content Engine application server accounts"
  Create new or designate existing application server accounts for Content Engine, as shown in the following table.

  "Configuring WebLogic Server for Content Engine" on page 52
  You need to configure WebLogic Server after installing it on the machine where you are going to install and deploy Content Engine.

  "Configuring WebLogic Server for Application Engine or Workplace XT" on page 53
  You must install Oracle WebLogic Server on the machine where you are going to install and deploy Application Engine or Workplace XT.

  "Configuring Process Engine clients for ORB" on page 54
  Process Engine clients require either the IBM or the Sun Object Request Broker (ORB).

  "Configuring the documentation server" on page 55
  If you want to have a local documentation installation instead of linking to the information center at www.ibm.com, plan to install IBM FileNet P8 documentation on an application server. Deploying the help as a Web application ensures that you can access online help from within IBM FileNet P8 applications and use the full-text search feature.

Creating Content Engine application server accounts

Create new or designate existing application server accounts for Content Engine, as shown in the following table.

User/Group	Description
Application server administrator user name *ce_appserver_admin*	Create an application server administrative account that can log on to Content Engine's application server administration console.
Administrative Console user *ce_appserver_console_admin*	

Configuring WebLogic Server for Content Engine

You need to configure WebLogic Server after installing it on the machine where you are going to install and deploy Content Engine.

Before installing and deploying Content Engine on a WebLogic Server machine, you must create a WebLogic Server domain and install JDBC drivers (The drivers must be installed on the WebLogic Server machine whether your database is collocated or not).

To configure WebLogic Server:

1. Use the WebLogic Configuration Wizard to create a WebLogic Server domain for Content Engine. In the examples below, use the domain name FNCEDomain. Keep the following in mind as you configure the domain:

    a. Set the server start mode to Production mode.

    b. Select the appropriate Java Development Kit (JDK) for your environment, as specified in the *IBM FileNet P8 Hardware and Software Requirements*.

2. Use the WebLogic Administration Console to set the following:

    a. Optional: Create a WebLogic Server authentication provider. You can use the Content Engine Configuration Manager tool to create a WebLogic Server authentication provider later on, or you can create the provider now by using the WebLogic Console.

    Important: In some situations (for example, if you have a single-sign-on provider, such as Netegrity SiteMinder), Configuration Manager cannot configure a WebLogic authentication provider.

    b. For performance reasons set the parameters that control searches within the authentication provider, as shown in the following table:

Parameter	Value	Description
Group Membership Searching	unlimited	Group searches are unlimited in depth
Max Group Membership Search Level	0	Only direct group members are found

    If you encounter performance problems, change the Group Membership Searching parameter value to limited.

    c. If you do not use Configuration Manager to configure your LDAP settings, set the Control Flag value. Set the Control Flag to REQUIRED to allow logons to FNCEDomain by LDAP-authenticated users in the Default Authenticator who are not in FNCEDomain's active security realm. The Configuration Manager tool will set this flag when you run the Configure LDAP task.

    d. If you are using multiple authentication providers in an Active Directory environment of multi-forest domains, reorder (as needed) the list of providers so that the most-frequently-used provider is first in the list, and the least-frequently-used is last. Reordering is necessary to prevent logon failures when IBM FileNet P8 Workplace is being accessed by many users simultaneously.

    e. Specify the following heap sizes for the JVM:

    Initial Java heap size (-Xms):
        512 MB

> **Maximum Java heap size (-Xmx):**
>> 1024 MB

3. Refer to the *IBM FileNet P8 Hardware and Software Requirements* guide for information on the JDBC driver file for the database type that you need for the GCD or for an object store you will be creating later.

4. Depending on your database, use one of the following procedures to install the JDBC drivers.

   Microsoft SQL Server

   a. Download and unzip SQL Server Driver JDBC Driver from Microsoft Support to a directory `jdbc_path` on your application server machine, such as:

      Windows
      > `C:\jars`

   b. Perform the following step, depending on your operating system type.

      Windows
      > Edit the file `startWebLogic.cmd` (by default, in the directory `C:\bea\user_projects\domains\bin\`*FNCEDomain*) for the WebLogic domain you created. Insert the following two lines immediately after the first occurrence of the line `CLASSPATH=...`
      >
      > `set JDBC_PATH=`*jdbc_path*`\sqljdbc_1.0\enu\sqljdbc.jar`
      > `set CLASSPATH=%JDBC_PATH%;%CLASSPATH%`

5. If your application server uses the IBM JVM, edit the JAVA_OPTIONS variable to improve performance.

   Windows
   > Immediately before the `set SAVE_JAVA_OPTIONS=%JAVA_OPTIONS%` line in the `startWebLogic.cmd` file, insert this line. Do not type any line breaks.
   >
   > `set JAVA_OPTIONS=%JAVA_OPTIONS% -Dprogram.name=%PROGNAME%`
   > `-Dfilenet.pe.peorb.pool.min=2 -Dfilenet.pe.peorb.pool.max=5`

6. Stop and then start WebLogic Server.

7. Give the Configuration Manager user, *config_mgr_user*, the following permissions:

   a. Read, write, and execute permission on the domain directory `../users_projects/domains/`*your_domain*.

   b. Read and execute permission on the `../common/bin` directory.

Configuring WebLogic Server for Application Engine or Workplace XT

You must install Oracle WebLogic Server on the machine where you are going to install and deploy Application Engine or Workplace XT.

Application Engine or Workplace XT can be collocated with Content Engine as long as the server is appropriately sized. However, each instance of Application Engine or Workplace XT and each instance of the Content Engine must run in its own JVM. For assistance in sizing your system, contact your service representative.

To configure WebLogic Server:

1. Verify that the application server is set to use JSESSIONID as the default cookie name. To avoid forcing end users to log in individually to applets such as Process Designer, Search Designer, and Process Simulator, configure the application server to use JSESSIONID as cookie name, and not use

application-unique cookie names. Using JSESSIONID is typically the default setting for the supported application servers. Both Application Engine and Workplace XT use cookie names for passing session information between Application Engine or Workplace XT and the client browser.

2. Create a WebLogic Server domain before installing and deploying Application Engine or Workplace XT. Refer to your WebLogic Server documentation for detailed instructions.

   Important: You will perform further configuration on Workplace XT after you perform the installation.

3. Enable trust between the Content Engine WebLogic domain and the Application Engine or Workplace XT WebLogic domain. If you have already configured a trust on the Content Engine, perform these steps only on the Application Engine or Workplace XT server. You must use the same password as the one used on the Content Engine.

   a. Prepare the WebLogic Administration Console for editing.

   b. In the advanced security settings for your domain, enter a password for the domain. You must enter the same password for both the Content Engine domain and the Application Engine or Workplace XT domain.

   c. Save and activate your changes.

   d. Restart the server if needed.

   e. Repeat this procedure in each domain for which you want to enable trust.

   Important: If you are enabling this feature in a managed server environment, you must stop the Administration server and all the Managed Servers in both domains and then restart them. If this step is not performed, servers that were not rebooted will not trust the servers that were rebooted. Refer to your BEA documentation for more information.

Configuring Process Engine clients for ORB

Process Engine clients require either the IBM or the Sun Object Request Broker (ORB).

This applies to the following configurations:

- J2EE application server clients such as Workplace or Workplace XT
- Content Engine when using the workflow subscription processor to launch workflows
- Non-J2EE and custom applications

The default ORB varies by application server, so in most instances no changes are required. However, in certain configurations you must override the defaults as follows:

Application server	Changes
Oracle WebLogic Server with the Sun JVM	No changes are required.

Application server	Changes

WebLogic Server with the JRockit JVM	Replace the default ORB with the Sun ORB by adding the following code to your application server startup file. Enter without carriage returns:
	
	`set JAVA_OPTIONS=%JAVA_OPTIONS%`
	`-Dorg.omg.CORBA.ORBClass`
	`=com.sun.corba.se.impl.orb.ORBImpl`
	`set JAVA_OPTIONS=%JAVA_OPTIONS%`
	`-Dorg.omg.CORBA.ORBSingletonClass`
	`=com.sun.corba.se.impl.orb.ORBSingleton`

Configuring the documentation server

If you want to have a local documentation installation instead of linking to the information center at www.ibm.com, plan to install IBM FileNet P8 documentation on an application server. Deploying the help as a Web application ensures that you can access online help from within IBM FileNet P8 applications and use the full-text search feature.

Ensure that you have a supported application server. The application server can be one of the servers that you already prepared for Content Engine or Application Engine.

After you have chosen the application server for the documentation, you can specify the value for the documentation URL. This is useful when recording the installation properties for the other IBM FileNet P8 components.

To determine the documentation URL:

1. Choose the application server for the IBM FileNet P8 documentation.
2. Make a note of the server URL. The format of the URL depends on the documentation package you choose:

Table 1. Documentation URLs

FileNet P8 Documentation package	URL
Installed Eclipse-based information center	http://*yourdocserver*:*port#*/p8docs/index.jsp
Installed non-Eclipse-based documentation package	http://*yourdocserver*:*port#*/ecm_help/_start_here.htm

3. Search the Installation and Upgrade Worksheet for documentation server properties in the **Property** or **Parameter** column.
4. Enter the URL value into the worksheet for each instance of that property for the components you are going to install.

Appendix A. Preparing non-English environments for installing IBM FileNet P8 Platform

To run IBM FileNet P8 components in a non-English environment, certain conditions must be met. Review the following considerations and tasks, organized by administrator role, if you plan to run IBM FileNet P8 in a non-English environment.

"Application Server administrator"
To support Unicode UTF-8 characters, all FileNet P8 domain application servers must be properly configured and must have all fix packs installed.

"Security administrator" on page 58
The IBM FileNet P8 security administrator installation role includes configuring and maintaining directory servers.

"Database administrator" on page 58
The IBM FileNet P8 database administrator installation role includes configuring database installations and tablespaces, and creating database accounts.

"IT administrator" on page 59
Depending on the operating system, the IT administrator installs either a localized version of the operating system, or the operating system language pack.

"IBM FileNet P8 administrator" on page 60
The FileNet P8 administrator starts Process Engine services and configures Process Task Manager for Application Engine and Workplace XT.

"Limitations on installing in a non-English environment" on page 60
There are certain limitations on installing IBM FileNet P8 in non-English environments.

Application Server administrator

To support Unicode UTF-8 characters, all FileNet P8 domain application servers must be properly configured and must have all fix packs installed.

"Configuring character encoding (WebLogic Server)"
Set the appropriate codesets to UTF-8.

"WebLogic Server" on page 58
(WebLogic Server 9.x) To fix the BEA CR298435 character corruption problem in FileNet P8, apply WebLogic Server 9.2 MP1 on WebLogic Server 9.x.

Configuring character encoding (WebLogic Server)

Set the appropriate codesets to UTF-8.

To configure the correct character encoding:

Set the following Interop-Orb-Protocol (IIOP) attributes:
- Default Char Codeset to "UTF-8"
- Default Wide Char Codeset to "UTF-8"

WebLogic Server

(WebLogic Server 9.x) To fix the BEA CR298435 character corruption problem in FileNet P8, apply WebLogic Server 9.2 MP1 on WebLogic Server 9.x.

Security administrator

The IBM FileNet P8 security administrator installation role includes configuring and maintaining directory servers.

"Extended characters and user names"
Note the following considerations for localized IBM FileNet P8 accounts.

Extended characters and user names

Note the following considerations for localized IBM FileNet P8 accounts.

- With Microsoft Active Directory, Process Engine supports:
  - Extended characters in user names and passwords for all Latin1, Latin2, and Arabic languages
  - Extended characters in passwords for double-byte languages (Users names are not supported due to Microsoft limitations.)
- Process Engine does not support extended characters in LDAP attributes for authentication purposes. These attributes include, but are not limited to, such items as cn (common name), ou (organizational unit), or dc (domain component). ASCII characters are required for these attributes.
- Process Engine does not support non-ASCII user names. Therefore, do not use non-ASCII user names when starting the Process Task Manager.
- WebDAV and the SSO environment also do not support Non-ASCII user names.
- The Content Engine locale must match directory server locale to manage non-ASCII user names correctly.
- Support for a Windows non-ACSII user name is limited to the Process Engine database code page. For example, if the Process Engine database code page is Japanese, Process Engine can support only Japanese and English user names. If the Process Engine database code page is Latin1, Process Engine can support all Latin1 user names, such as English, French, German, Spanish, and so forth.

Database administrator

The IBM FileNet P8 database administrator installation role includes configuring database installations and tablespaces, and creating database accounts.

"Installing Microsoft SQL Server"
During installation, the Microsoft SQL Server installer program detects the Windows regional setting and sets the Microsoft SQL Server language setting accordingly. Use the regional setting selected by the installation program throughout the entire Microsoft SQL Server installation.

Installing Microsoft SQL Server

During installation, the Microsoft SQL Server installer program detects the Windows regional setting and sets the Microsoft SQL Server language setting accordingly. Use the regional setting selected by the installation program throughout the entire Microsoft SQL Server installation.

Microsoft does not recommend changing the selected regional setting unless you have to match the regional setting to the collation of another instance of Microsoft SQL Server or to the Windows regional setting of another computer. Localized

versions of Microsoft SQL Server are available in French, German, Spanish, Italian, Japanese, Korean, and Simplified and Traditional Chinese.

Restriction: Process Engine does not support Unicode UTF-8 encoded characters when using Microsoft SQL Server. Process Engine stores characters according to the Process Engine Windows regional setting.

Using a binary sort for Process Engine is faster when configuring Microsoft SQL Server. However, IBM does support other collation settings. (A specific list of supported sort orders is not available. The collation setting depends on which sort order you need to use in your production environment.)

Queried information returns in the selected sort order. The sort order can be either binary or based on a combination of sort options, such as language, case-insensitive (ci), case-sensitive (cs), accent-insensitive (ai), and accent-sensitive (as).

The collation designator is a collation name that is based on a Windows regional setting. When you specify a collation designator, you should specify a sort order as well.

For example:

French - Latin 1
SQL Collations: Dictionary order, case-sensitive, for use with 1252 Character Set (or any case-sensitive MS-SQL collation).

IT administrator

Depending on the operating system, the IT administrator installs either a localized version of the operating system, or the operating system language pack.

"Operating system considerations"
In addition to any operating system platforms, the IT administrator must consider the FileNet P8 Platform components that will be installed in a non-English environment.

"Microsoft Windows" on page 60
Use the localized Windows version when available. If the localized version is not available, use the English version with the appropriate regional setting.

Operating system considerations

In addition to any operating system platforms, the IT administrator must consider the FileNet P8 Platform components that will be installed in a non-English environment.

Application Engine or Workplace XT

Application Engine and Workplace XT can be installed:
- In any locale on any of the supported UNIX platforms
- On any localized version of Windows or in any region on the English version of Windows

Remember: The Application Engine and Workplace XT setting must match the Process Engine setting when the Process Task Manager is started. Otherwise, workflows can experience unexpected problems such as errors related to the way characters display.

Content Engine

Content Engine can be installed:
- In any locale on any of the supported UNIX platforms
- On any localized version of Windows or in any region on the English version of Windows

Process Engine

Process Engine must be installed on:
- A UTF-8 locale and a UTF-8 database character set when used with Oracle or DB2
- On any localized version of Windows for the supported language or in any region on the English version of Windows

Process Engine has limited language support on Windows. Process Engine supports only those languages that are supported by the Windows code page and a corresponding database character set. For example, if the active Windows code page and the corresponding database character set support Latin 1 languages, Process Engine will support such languages as English, French, Spanish, and Italian (and other Latin 1 languages) because these languages share a common character code page. However, Japanese cannot be supported in this scenario because the Japanese language uses a different character code page. Conversely, if the active Windows code page and the corresponding database character set support Japanese (and English) will be the only languages supported by Process Engine.
- All Windows character code pages support English in addition to their primary languages.
- The Java Runtime Environment version must be 1.5.0 or higher for the language capability to function correctly.

Microsoft Windows

Use the localized Windows version when available. If the localized version is not available, use the English version with the appropriate regional setting.

Use the Regional Options Control Panel to change the regional setting. For more information, see the Windows help system.

IBM FileNet P8 administrator

The FileNet P8 administrator starts Process Engine services and configures Process Task Manager for Application Engine and Workplace XT.

Limitations on installing in a non-English environment

There are certain limitations on installing IBM FileNet P8 in non-English environments.

Oracle WebLogic Server

WebLogic Server 8.1 is not supported on a Turkish operating system.

GB18030 support for Windows Process Engine

Process Engine installed on Windows does not support GB18030 Simplified
Chinese characters.

Publishing

When the watermark is checked in the Publishing Style Template Manager on the
Content Engine server, double-byte documents cannot be published. This is a
known independent software vendor issue.

Process Engine Modeler

To import a user defined XSD file containing non-English characters in Process
Designer, run the following command to convert characters to the Unicode
encoding format to match a UNIX Process Engine in a UTF-8 locale, then import
the XSD file into Process Designer.

```
Java -cp pe.jar filenet.vw.toolkit.utils.FileConverter /in filename /out outfilename
```

Appendix B. IBM FileNet P8 ports

The tables below list the port numbers used by IBM FileNet P8 components.

Content Engine ports

Port name	Default number	Notes
LDAP	389	This port is on the directory server, specified on the Content Engine server for authentication.
SSL	636	This port is on the directory server, specified on the Content Engine for authentication through SSL.
LDAP Global Catalog	3268	
WebLogic	7001	This port is on the WebLogic Server for Content Engine for communication with Content Engine by clients.
Microsoft SQL Server	1433	This port is on the Microsoft SQL Server database server, for communication with the database by Content Engine.

Process Engine ports

Port name	Default number	Notes
SMTP (E-mail Notification)	25	This port is on the SMTP server.
Microsoft SQL Server	1433	This port is on the database server, for connection between Process Engine and Microsoft SQL Server.
TMS	32768 (TCP)	Task Manager service. TM_daemon listens for requests from initfnsw running on the same or a different system analogous to COR_Listen listening for RPCs.
COR	32769 (TCP)	Courier service. COR_Listen listens on this port for incoming RPC requests.
NCH	32770 (UDP)	nch is the NCH deamon. NCH_daemon listens on this port. Pre-4.1.2 listened for broadcasts, etc. Post 4.1.2, only listens for old print servers to verify NCH is up.

Port name	Default number	Notes
fn_snmpd	161 (UDP)	IBM FileNet Image Services Simple Network Management Protocol Daemon. It listens for SNMP requests from the native OS snmp daemon. The native snmp daemon listens on this port and communicates with fn_snmpd through other local port. fn_snmpd does not listen on this port.
fn_trapd	35225 (UDP)	fn_trapd is the IBM FileNet Image Services trap daemon, which listens for notifications of the end of Image Services background processes running on the server and sys_logs information.
Rules Listener	32774 (TCP/IP)	This port is on the Process Engine server.
Process Engine Communication Port (IOR port)	32776	This port is on the Process Engine server.
BPM Web Services Reliable messaging client port	32767 (TCP) 32767 (UDP)	This port is on the Process Engine server.
Registry Port	32771	This port is on the Process Engine server and is used by the Component Manager.
Event Port	32773	This port is on the Process Engine server.
Process Engine Broker Port	32777	This port is on the Process Engine server, and is for the vworbbroker process handling RPCs from the PE API clients.

Application Engine ports

Port name	Default number	Notes
WebLogic	7001	This port is on the WebLogic Server for clients to connect to Application Engine.
WebLogic SSL	7002	This port is on the WebLogic Server for clients to connect to Application Engine through SSL.

Port name	Default number	Notes
Process Engine (RMI)	32771	This port is on the Windows Process Engine server for Process Task Manager to communicate with the Windows Process Engine Services Manager. This port is on the Application Engine server for Process Task Manager to communicate with the Component Managers and the Windows Process Application Engine (or Workplace XT) Services Manager.
Component Manager (Event Port)	32773	This port is on the Application Engine server, and is used when the Component Manager (running on the Application Engine Server) is configured to be triggered by events, instead of polling.
Web Services Reliable Messaging Client Port	xxxx - assigned by the client on the Application Engine Server, for a particular Component Manager instance.	This port (or a number of ports, one per Component Manager instance) is on the Application Engine server to respond to the WS-Reliable Messaging requests.

Autonomy K2 search engine ports

Port name	Default number	Notes
K2 Master Administration Server	9950 (default)	This port is located on the machine that hosts the K2 Master Administration Server machine. It is the software port for K2 Master Administration Server.
K2 Dashboard	9990 (default)	This port is located on the machine that hosts the K2 Master Administration Server. It is used specifically for the URL access for the K2 Dashboard.
K2 Index Server	9960 - 9979 (recommended range)	This port is located on the machine where you created the K2 Index Server service. It is for communication between the K2 Index Server and P8 Content Engine.

Port name	Default number	Notes
K2 Broker Server	9900 - 9909 (recommended range)	This port is located on the machine where you created the K2 Broker Server service. It is for communication between the K2 Broker Server and P8 Content Engine.
K2 Server (Search Server)	9920 - 9949 (recommended range)	This port is located on the machine where you created the K2 Server (Search Server) service. It is for communication between the K2 Server (Search Server) and Content Engine.
K2 Ticket Server	9910 - 9919 (recommended range)	This port is located on the machine where you created the K2 Ticker Server service. It is for communication between the K2 Ticket Server and Content Engine.

Rendition Engine and Content Engine ports for Liquent

Port name	Default number	Notes
Liquent input port	2867 (COM Repository only)	This port allows for distributed processing of jobs on the Rendition Engine servers when there are more than one Rendition Engine server. A Rendition Engine server or Rendition Engine client (Content Engine Publishing server) on which a job is submitted makes the load balancing decision on where to run the job. If the job is sent to another Rendition Engine server for execution then it uses the other Rendition Engine server's input port to do so.

Port name	Default number	Notes
Liquent notify port	2868	This port is also related to the distributed processing of jobs on the Rendition Engine servers. When a job is sent to another Rendition Engine server then the originating Rendition Engine server or Rendition Engine client (Content Engine Publishing server) is notified on this port by the other Rendition Engine Rendition Engine server when the other Rendition Engine Rendition Engine server has finished processing the job. One case of this is when the Rendition Engine Rendition Engine server is notifying the Rendition Engine client (Content Engine Publishing server) that a conversion job has completed.
Liquent event port	2869	This port is used by each Rendition Engine server to send or receive events. These events are used by the Rendition Engine servers to keep each other informed of current activities. The Rendition Engine or Liquent Domain Manager uses these events for the job status display.
Liquent admin port	2870	This port is used by each Rendition Engine server for internal administrative functions. It's primary use is for each Rendition Engineserver to publish its current activity statistics (for example, how busy the business services are) for use by other Rendition Engine servers to make load balancing decisions.
Liquent file transfer port	2871	This port is used for transferring source, temporary, and result files between Rendition Engine servers as well as between Rendition Engine servers and Rendition Engine clients (Content Engine Publishing servers).

Port name	Default number	Notes
Liquent job queue port	2872	This port is for job queuing by the render business service on all Rendition Engine servers.
Microsoft SQL Server	1433	When the selected Rendition Engine database is Microsoft SQL Server, then this port is used by the Rendition Engine servers and the Rendition Engine clients (Content Engine Publishing server) to access the Rendition Engine database.

Tip: If the Liquent port number assigned to the Rendition Engine/Liquent software conflicts with the port number required by another application or service that runs on the Rendition Engine server or the Content Engine Publishing server, then the default values can be changed in the Rendition Engine/Liquent Domain Manager. The above port numbers are the default values set by the Rendition Engine installer for the COM Repository in the Rendition Engine/Liquent domain, which represents the Rendition Engine server itself. When a Content Engine Publishing server is configured to point to the Rendition Engine server, one must create a "Java Repository" in the Rendition Engine/Liquent domain that represents the Content Engine Publishing server as a Java client to the Rendition Engine server, and this repository will also have the same default port number values.

The database port number is not specified directly by the Rendition Engine server's Liquent software, but it is specified in the Rendition Engine client/Content Engine Publishing server's Rendition Engine Connection (for the JDBC connection from the Content Engine Publishing code to the Rendition Engine database server).

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
3-2-12, Roppongi, Minato-ku, Tokyo 106-8711 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml (www.ibm.com/legal/copytrade.shtml).

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Other company, product, and service names may be trademarks or service marks of others.

Index

IBM. ®

Program Number: 5724-R76
                 5724-R81

GC19-2824-01