

PREPARING TO INSTALL ADOBE® LIVECYCLE® ES4 (SERVER CLUSTER)

Legal notices

For legal notices, see http://help.adobe.com/en_US/legalnotices/index.html.

Contents

Chapter 1: About This Document

1.1 Conventions used in this document	1
1.2 Additional information	3

Chapter 2: Introduction to Application Server Clusters

2.1 About clustering application servers	4
2.2 Terminology	5
2.3 Clustering LiveCycle	6
2.4 Supported topologies	8
2.5 Unsupported topologies	9

Chapter 3: System Requirements

3.1 Prepare your server environment	11
3.2 System requirements	11
3.3 LDAP configuration	27
3.4 Global document storage directory	27
3.5 Synchronizing clock times	29
3.6 Installation considerations	29
3.7 LiveCycle IPv6 support	30
3.8 Central Migration Bridge service	33

Chapter 4: LiveCycle Credentials and Certificates

4.1 Obtaining the Reader Extensions Rights credential	35
4.2 Obtaining digital certificates for use with Digital Signatures	36

Chapter 5: Creating the LiveCycle Database

5.1 Database configuration requirements	37
5.2 Minimum database user permissions	38
5.3 Creating an Oracle database	38
5.4 Creating a SQL Server database	39
5.5 Creating a DB2 database	43
5.6 Creating a MySQL database	46

Chapter 1: About This Document

Adobe® LiveCycle® Enterprise Suite 4 (ES4) is an enterprise server platform that helps you automate and streamline business processes. LiveCycle comprises the following components:

- J2EE-based Adobe® LiveCycle® Foundation 11 provides server capabilities and runtime environment
- Tools to design, develop, and test LiveCycle applications
- Modules and Services are deployed on LiveCycle server and provide functional services

For more information about the LiveCycle architecture and capabilities, see [LiveCycle Overview](#).

This document is part of a larger documentation set available at [LiveCycle Documentation page](#). It is advised that you start with the preparing guide and then move on to installation and configuration guide depending on whether you are performing a fresh installation (single server or cluster setup) or upgrading your existing LiveCycle deployment. For Turnkey deployment, which is only for evaluation purposes, see [Installing and Deploying LiveCycle using JBoss Turnkey](#).

1.1 Conventions used in this document

The following naming conventions are used for common file paths.

Name	Description	Default value
<i>[LiveCycle root]</i>	The installation directory that is used for all LiveCycle modules. The installation directory contains subdirectories for Adobe® LiveCycle® Configuration Manager, the LiveCycle SDK, and each LiveCycle module installed (along with the product documentation). This directory also includes directories relating to third-party technologies.	Windows: C:\Adobe\Adobe LiveCycle ES4\ AIX, Linux, and Solaris: /opt/adobe/adobe_lifecycle_es4/
<i>[appserver root]</i>	The home directory of the application server that runs the services that are part of LiveCycle.	JBoss on Windows: C:\jboss JBoss on Linux and Solaris: /opt/jboss JBoss Enterprise Application Platform on Windows: C:\jboss-eap-<version>\jboss-as JBoss Enterprise Application Platform on Linux and Solaris: /opt/jboss-eap-<version>/jboss-as WebSphere on Windows: C:\Program Files\IBM\WebSphere\AppServer WebSphere on Linux and Solaris: /opt/IBM/WebSphere/AppServer WebSphere on AIX: /usr/IBM/WebSphere/AppServer, or, /opt/IBM/WebSphere/AppServer WebLogic on Windows: C:\Oracle\Middleware\wlserver_<version> WebLogic on Linux and Solaris: /opt/Oracle/Middleware/wlserver_<version>
<i>[WL_Home]</i>	The install directory for WebLogic as specified for the WL_HOME environment variable.	WebLogic on Windows: C:\Oracle\Middleware\ WebLogic on Linux and Solaris: /opt/Oracle/Middleware
<i>[appserverdomain]</i>	The domain that you configured on WebLogic. The default domain is called <i>base_domain</i> .	WebLogic on Windows: C:\Oracle\Middleware\user_projects\domains\base_domain WebLogic on Linux and Solaris: /opt/Oracle/Middleware/user_projects/domains/base_domain
<i>[WebSphere ND root]</i>	The install directory for WebSphere Application Server, Network Deployment	WebSphere on Windows: C:\Program Files\IBM\WebSphere\AppServer WebSphere on Linux and Solaris: /opt/IBM/WebSphere/AppServer WebSphere on AIX: /usr/IBM/WebSphere/AppServer

About This Document

Name	Description	Default value
<i>[server_name]</i>	The name of the server configured on your WebLogic or WebSphere server	WebLogic: server1 WebSphere: server1
<i>[profile_name]</i>	The profile name for the JBoss application server.	Adobe pre-configured JBoss: lc_<db-name> Manual or Downloaded Jboss: all
<i>[dbserver root]</i>	The location where the database server is installed	Depends on the database type and your specification during installation

Most of the information about directory locations in this document is cross-platform (all filenames and paths are case-sensitive on AIX, Linux, and Solaris). Any platform-specific information is indicated as required.

1.2 Additional information


The resources in this table can help you learn about LiveCycle.

For information about	See
LiveCycle and the modules in general	LiveCycle Overview
Other services and products that integrate with LiveCycle	LiveCycle Developer Center
All documentation that is available for LiveCycle	LiveCycle documentation

Chapter 2: Introduction to Application Server Clusters

2.1 About clustering application servers

A *cluster* is a group of application server instances that run simultaneously, which act like a single system, enabling high availability and load balancing. Within a cluster, multiple server instances can run on the same computer (known as a *vertical cluster*) or can be located on different computers (known as a *horizontal cluster*), or they can form a combination of both horizontal and vertical clusters. With clustering, client work can be distributed across several nodes instead of handled by a single application server.



In a clustered configuration, application server instances are server members of the cluster, all of which must have identical application components deployed on them. However, other than the configured applications, cluster members do not have to share any other configuration parameters. For example, you can cluster multiple server instances on one computer, with a single instance on another computer, provided they are all running the same version of application server.

By clustering, you can achieve one or more of the following benefits. How you implement clustering determines which benefits are achieved:

- Failover
- Load balancing
- Scalability

2.1.1 Failover

Failover allows one or more application server instances to act as backup to a failed application server instance and resume processing the task, therefore enabling another application server to carry on processing. However, if an application server fails during a transaction, the backup application server does not recover the state of the failed instance. If a server fails when a user enters data into a form, for example, the data may have to be reentered.

2.1.2 Load balancing

Load balancing is a technique used to distribute work across a number of systems so that no single device is overwhelmed. If one server starts to get congested or overloaded, requests are forwarded to another server that has more capacity.

2.1.2.1 Application server load balancing

Application server load balancing is useful for managing the load between application servers. Application servers can be configured to use a weighted round-robin routing policy that ensures a balanced routing distribution based on the set of server weights that are assigned to the members of a cluster. Configuring all servers in the cluster to have the same weight produces a load distribution in which all servers receive approximately the same number of requests. Weighting some servers more heavily sends more requests to these servers than to those that are weighted less heavily.

Preferred routing configurations can also be configured to ensure, for example, that only cluster members on that node are selected (using the round-robin weight method) and cluster members on remote nodes are selected only if a local server is not available. Application server load balancing is best used when balancing is needed between tiers.

2.1.2.2 Web server load balancing

Web server load balancing is useful for queuing and throttling requests. For the Apache HTTP Server, the most commonly used method for load balancing is Round-Robin DNS.

Round-Robin DNS is a relatively simple method of load balancing, where a domain name system (DNS) server provides a name to address resolution and is always involved when a host name is included in a URL. A Round-Robin DNS server can resolve one single host name into multiple IP addresses such that requests for a single URL (containing a host name) actually reference different web servers. The client requests a name resolution for the host name but, in fact, receives different IP addresses, therefore spreading the load among the web servers. In a simple configuration, the Round-Robin DNS server cycles through the list of available servers.

2.1.3 Scalability and availability

Scalability in a cluster means that an administrator can increase the capacity of the application dynamically to meet the demand without interrupting or negatively impacting service. Clusters allow administrators to remove nodes from a cluster in order to upgrade components, such as memory, or to add nodes to the cluster without bringing down the cluster itself.

2.2 Terminology

Each application server vendor uses specific terminology, which is defined here to avoid confusion.

2.2.1 JBoss

server: Represents an instance of a Java™ virtual machine (JVM™).

cluster: Represents a logical grouping of multiple application servers for administration, application deployment, load balancing, and failover purposes.

2.2.2 WebLogic

server: Represents an instance of a Java™ virtual machine (JVM™).

machine: Represents a physical system running one or more instances of WebLogic Server.

cluster: Represents a logical grouping of multiple application servers within a domain for administration, application deployment, load balancing, and failover purposes.

2.2.3 WebSphere

server: Represents an instance of a Java™ virtual machine (JVM™).

node: Represents a physical system running one or more instances of WebSphere Application Server.

cell: Represents a logical grouping of multiple nodes for administrative purposes.

federation: The process of joining a stand-alone WebSphere node to a WebSphere cell.

cluster: Represents a logical grouping of multiple application servers within a cell for administration, application deployment, load balancing, and failover purposes.

2.3 Clustering LiveCycle

If you install a LiveCycle component on an application server cluster, here are some things you must know:

- Each application server in the cluster on which you deploy LiveCycle must have a homogeneous topology (that is, all nodes in the cluster must be configured identically). You can ensure that all modules are configured identically by configuring run-time properties in the single-installation staging area.
- The configuration is deployed by using the single entity approach; all nodes in a cluster are deployed as if deploying to a single node.

Setting up a clustered environment for LiveCycle involves the tasks that follow.

2.3.1 Clustering LiveCycle on JBoss

You must perform the following tasks to deploy LiveCycle on a JBoss Application Server cluster:

- 1 Prepare the hardware as per your requirements.
- 2 Install the operating system and update it with all necessary patches and service packs.
- 3 Install and configure the database server.
- 4 Install and configure the application server cluster.
- 5 Install LiveCycle modules and configure LiveCycle for deployment.
- 6 Deploy LiveCycle to JBoss Application Server.

2.3.2 Clustering LiveCycle on WebLogic

You must perform the following tasks to deploy LiveCycle on a WebLogic Server cluster:

- 1 Prepare the hardware as per your requirements.
- 2 Install the operating system and update it with all necessary patches and service packs.
- 3 Install and configure the database server.
- 4 Install and configure the application server cluster.
- 5 Install LiveCycle modules and configure LiveCycle for deployment.
- 6 Deploy LiveCycle to WebLogic Server.

2.3.3 Clustering LiveCycle on WebSphere

You must perform the following tasks to deploy LiveCycle on a WebSphere Application Server cluster:

- 1 Prepare the hardware as per your requirement.
- 2 Install WebSphere Application Server Network Deployment software on the server that will manage the cluster
Note: Select default environment(Cell), if you want to install deployment manager and an application server profile on the same machine.
- 3 Install WebSphere Application Server Network Deployment software on any other server that will be used as a cluster node.
Note: In steps 2 and 3, you can also select None and create Deployment Manager or Application Server profiles later using command line utility.
Note: For cluster, use Network Deployment software only. Base software version is only for standalone application server. As per IBM EULA, user are not allowed to use Base software with Network Deployment software.
- 4 Install the HTTP Server (IBM HTTP Server) and start Webservers.
- 5 Install WebSphere Update Installer on each server/node and deploy the required Fixpack.
Before you deploy LiveCycle, you must update to the supported version of WebSphere as specified in [Adobe LiveCycle ES4 Supported Platforms](#). See the [WebSphere Support Website](#) for Fix Pack installation instructions and to access the WebSphere updates.
- 6 Install Websphere Installation Manager on each server/node and deploy the JPA 2.0 (if not already installed on your WebSphere).
You must install WebSphere Feature Pack for JPA 2.0. See, [WebSphere Support Website](#) and search Installing and uninstalling the feature pack to download and install Feature Pack.
- 7 Set up the cluster:
 - Start the Deployment Manager.
 - Start all servers on all nodes.
 - Federate nodes to the Deployment Manager.
 - Create the cluster using existing application server profiles or create member using application server template.
 - Start the cluster.
- 8 Configure cluster resources.
- 9 (For manual deployment) Deploy applications to both the cluster and HTTP Server.

10 Generate the WebSphere HTTP plug-in and propagate it.

2.4 Supported topologies

The following sections discuss various topologies, both clustered and non-clustered, that you can employ. For additional information about configuring your application server in a cluster, go to one of the following websites, as applicable to your application server:

- (WebLogic Server) [Supported Configurations for WebLogic Platform](#)
- (WebSphere Application Server) [WebSphere Application Server \(Distributed platforms and Windows\)](#)
- (JBoss Application Server) [High Availability Enterprise Services via JBoss Clusters](#)

2.4.1 Combined web, application, and database servers

This topology consists of a web server, an application server, and a database server on the same node. This topology is the simplest one and must be used for development only.

2.4.2 Combined web and application servers with separate database server

This topology can be considered for production in case the load on the user interface (including the web tier) is minimal, with a small number of users.

Combining the web and application servers means that all Enterprise JavaBeans™ (EJB) look-ups are local, and therefore reduces the overhead of doing a remote look-up. Also, this topology reduces the network overhead of a round trip between the web tier and the application tier.

However, with both servers on the same node, if the web tier is compromised, both tiers are compromised. If the web tier experiences a heavy load, the application server processing is affected and vice versa. User response time is usually affected in situations when users need to wait a significant amount of time to get a page back due to all server resources (that is, CPU and/or memory) being consumed by the application server.

2.4.3 Single web server with combined application and database server

The simplest topology that should be considered for a production environment is a web server and combined application server with a database server. Use this topology only if you are sure that your database load will be minimal. In this scenario, the web server is providing a redirection to the application server. The advantages of this topology are low cost, low complexity, and no need for load balancing. The disadvantages of this topology are little redundancy, low scalability, inability to perform updates and upgrades, and possible low performance due to too many CPU processes.

2.4.4 Separate web, application, and database servers

This topology is the most common in production systems because it allows allocation of separate resources to each of the tiers. In this case, the web server acts as a proxy to the web tier on the application server that hosts the web components. This level of indirection provides additional security by securing the application server even if the web server is compromised.

2.4.5 Adding additional web servers

You can add additional web servers for scalability and failover. When using multiple web servers, the WebLogic/WebSphere HTTP plug-in configuration file must be applied to each web server. Failure to do so after introducing a new application will likely cause a “404 File Not Found” error to occur when a user tries to access the web application.

2.4.6 Adding additional application servers

This topology is used in most large-scale production systems where the application servers are clustered to provide high availability and, based on the topology, failover and load balancing.

Clustering application servers has these benefits:

- You can use cheaper hardware configurations and still achieve higher performance
- You can upgrade software on servers without down time
- Provides higher availability (that is, if one server fails, the other nodes in the cluster pick up the processing)
- Provides the ability to leverage load-balancing algorithms on the web server (by using load balancers) as well as on the EJB tier for processing requests

LiveCycle components are typically CPU-bound. As a result, performance gains are better achieved by adding more application servers than by adding more memory or disk space to an existing server.

2.4.7 Multiple JVMs

Vertical scaling of multiple JVMs offers the following advantages:

Increased processing power efficiency: An instance of an application server runs in a single JVM process. However, the inherent concurrency limitations of a JVM process prevent it from fully using the memory and processing power of multi-CPU systems. Creating additional JVM processes provides multiple thread pools, each corresponding to the JVM process that is associated with each application server process. This correspondence avoids concurrency limitations and lets the application server use the full processing power of the computer.

Load balancing: Vertical scaling topologies can use the WebLogic Server or WebSphere Application Server workload management facility.

Process failover: A vertical scaling topology also provides failover support among application server cluster members. If one application server instance goes offline, the other instances on the computer continue to process client requests.

2.5 Unsupported topologies

The following topologies are not supported for LiveCycle.

Splitting the web container from the EJB container

Splitting LiveCycle Server into presentation/business logic tiers and running them on distributed computers is not supported.

Geographically distributed configuration

Many applications locate their systems geographically to help distribute the load and provide an added level of redundancy. LiveCycle does not support this configuration because LiveCycle components cannot be pulled apart to run on different hosts; LiveCycle is deployed as a monolithic application.

Most of the information about directory locations in this document is cross-platform (all file names and paths are case-sensitive on Linux and UNIX). Platform-specific information is indicated as required.

Chapter 3: System Requirements

3.1 Prepare your server environment

Do the following tasks to prepare your server environment:

- 1 Verify that the installation media is not damaged. It is recommended that you copy the installer files to the hard disk of the computer where you intend to install LiveCycle, instead of installing directly from the installation media.
Also verify the MD5 checksum of the installation media matches the checksum posted at [Adobe Website](#).
- 2 Read “[3.2 System requirements](#)” on page 11, and ensure that your software infrastructure is compliant.
- 3 Read System requirements, and ensure that your hardware infrastructure is in place, as per your requirements.
- 4 Install the operating system and update with all necessary patches and service packs.
- 5 Install and configure the database server.
- 6 Install and configure the application server.
- 7 Ensure that your target computer has at least the required amount of RAM as described in “[3.2.1 Minimum hardware requirements](#)” on page 11.

3.2 System requirements

Click [here](#) to view the system requirements for a single-server deployment.

3.2.1 Minimum hardware requirements

This table provides the minimum hardware requirements that LiveCycle supports.

Operating system	Minimum hardware requirement
Microsoft Windows Server® 2008 R1/R2 (64-bit architecture)	Intel® Pentium® 4, 2.8 GHz processor or equivalent VMWare ESX 3.0 or later RAM: 4 GB (64-bit OS with 64-bit JVM) Free disk space: 5.4 GB of temporary space plus 3.4 GB for LiveCycle
Sun Solaris 10 (64-bit architecture)	UltraSPARC® IIIi, 1.5 GHz processor Solaris Containers (Zones) partitioning RAM: 4 GB (64-bit OS with 64-bit JVM) Free disk space: 5.4 GB of temporary space plus 3.4 GB for LiveCycle

Operating system	Minimum hardware requirement
IBM AIX 6.1, 7.1	P5 pSeries 520 (Model 52A) 9131-52A, 1.8 GHz processor LPAR partitioning RAM: 4 GB (64-bit OS with 64-bit JVM) Free disk space: 5.4 GB of temporary space plus 3.4 GB for LiveCycle
SUSE Linux Enterprise Server 11.0 (64-bit edition only)	Dual core, 1 GHz processor VMWare ESX 3.0 or later RAM: 4 GB (64-bit OS with 64-bit JVM) Free disk space: 5.4 GB of temporary space plus 3.4 GB for LiveCycle
Red Hat Enterprise Linux AP or ES 5.5 or 6 (64-bit edition only)	Dual core, 1 GHz processor VMWare ESX 3.0 or later RAM: 4 GB (64-bit OS with 64-bit JVM) Free disk space: 5.4 GB of temporary space plus 3.4 GB for LiveCycle

3.2.1.1 Intel x86 compatibility

On supported Windows and UNIX -like environments, LiveCycle supports Intel and AMD64 compatible chipsets running either 32-bit or 64-bit supported operating systems.

Note: 32-bit platforms are supported only for development and evaluation purposes.

3.2.1.2 Recommended hardware requirements

For a small production environment:

Intel environments: Pentium 4, 2.8 GHz or greater. Using a dual core processor will further enhance performance.

Sun SPARC environments: UltraSPARC V or later.

IBM AIX environments: Power4 or later

Memory requirements: 4 GB of RAM.

Note: It is recommended to add additional 2 GB of RAM if you are running the publish instance of Correspondence Management Solution on the same machine.

3.2.1.3 Recommended hardware requirements for client-side computers

Adobe® LiveCycle® Workbench 11

- Disk space for installation:

1.5 GB for Workbench only 1.7 GB on a single drive for a full installation of Workbench, Adobe® LiveCycle® Designer 11, and the samples assembly 400 MB for temporary install directories - 200 MB in the user temp directory and 200 MB in the Windows temporary directory

Note: If all of these locations reside on a single drive, there must be 1.5 GB of space available during installation. The files copied to the temporary directories are deleted when installation is complete.

- Memory for running Workbench: 2 GB of RAM
- Hardware requirement: Intel® Pentium® 4 or AMD equivalent, 1 GHz processor
- Minimum 1024 X 768 pixels or greater monitor resolution with 16-bit color or higher

- TCP/IPv4 or TCP/IPv6 network connection to the LiveCycle server

Note: You must have Administrative privileges to install Workbench on Windows. If you are installing using a non-administrator account, the installer will prompt you for the credentials for an appropriate account.

Adobe® LiveCycle® Designer 11

- A minimum of Adobe® Acrobat® 9 Professional, Acrobat 9 Standard, or Adobe Reader® 9. To benefit from the new features in Designer 11, you must have Acrobat XI Professional, Acrobat XI Standard, or Reader XI.
- Adobe® Flash® Player 11.1 or later.
- (Optional) Adobe® Flash® Builder® 4.5 or later.
- (Optional) Adobe® Flex® SDK 4.1 (required for customizing form guide components that are shipped with Designer).

Note: Use the Flex SDK 4.1 included with the Designer installer only. Do not use any other version of Flex SDK obtained from the Adobe web site.

End-user hardware requirements

- Adobe® LiveCycle® Workspace 11: 1 GB of RAM (includes requirements for Adobe Flash and Adobe Reader)
- Adobe Flash Player 11.1 and later: 512 MB of RAM (1 GB recommended)
- Adobe Reader 9 and later: 128 MB of RAM (256 MB recommended)

Note: For more information, see [Adobe LiveCycle ES4 Supported Platforms](#).

3.2.2 Supported Platforms

For detailed information about the supported combinations of operating systems, application server, databases, JDKs, and so on, see [Adobe LiveCycle ES4 Supported Platforms](#).

3.2.3 Additional requirements for Linux and Solaris

Note: On Linux and Solaris machines, LiveCycle installer uses the JDK installed on the machine. Therefore, you must ensure to install the supported JDK version. On other operating systems, the installer uses the JVM bundled with the installer.

3.2.3.1 Installing and configuring UTF-8

When installing LiveCycle on AIX, Linux, or Solaris operating systems, you must install and configure the US English version of UTF-8 locale if it is not already installed. You will need the install media (CDs or DVDs) for the operating system to perform this task.

Note: On Linux platforms, this locale is installed by default and is called `en_US.utf8`. It can be verified by using the `locale -a` command.

Installing UTF-8 on AIX

- 1 Verify the US English UTF-8 locale is not installed by typing `locale -a` in a command prompt. Verify that the command output does not contain the entry `EN_US.UTF-8`.
- 2 Access the AIX SMIT utility (in text mode) by typing `smitty mle_add_lang` at the root in the command prompt.
- 3 On the screen that appears, select **UTF-8 US English (United States) [EN_US]** from both the **CULTURAL CONVENTION** and **LANGUAGE TRANSLATION** drop-down lists.

Note: Keep the `INPUT DEVICE/DIRECTORY` as the default `/dev/cd0` setting.

System Requirements

- 4 Press **Enter** to proceed. A message such as the following will appear:

```
installp: Device /dev/cd0 not ready for operation.
Please insert media and press Enter to continue.
```

- 5 Insert the appropriate AIX install disk in the disk drive.
- 6 When the command is complete, exit the SMIT utility and type `locale -a` to verify that `EN_US.UTF-8` is set as the locale.

Installing UTF-8 on Solaris

- 1 Verify the US English UTF-8 locale is not installed by typing `locale -a` in a command prompt. Verify that the command output does not contain the entry `EN_US.UTF-8`.

- 2 Insert the Solaris install CD #1 in the disk drive and mount it to an appropriate location, for example:

```
/cdrom/sol_10_807_sparc/s0
```

- 3 Type the following command as root: `localeadm -a nam -d /cdrom/sol_10_807_sparc/s0`

Note: This command installs all locales in the North America (`nam`) region even if you specify only the `en_US.UTF-8` locale.

- 4 When the command is complete, type `locale -a` to verify that `EN_US.UTF-8` is set as the locale.

Note: See this link for [FAQs on Solaris Locales](#).

3.2.3.2 Solaris

Note: Ensure that X Window libraries are installed on your operating system. This is required for Adobe® LiveCycle® PDF Generator 11 and Adobe® LiveCycle® Forms Standard 11. See the documentation for your operating system for more information.

Important: Do not use the Solaris `tar` command to extract files or errors (such as missing files) will occur. Download the GNU tar tool and use it to extract all files on a Solaris environment.

3.2.3.3 Linux

On Linux operating systems, ensure the following:

- **All Linux distributions:** Ensure that X Window libraries are installed on your operating system. This is required for PDF Generator and Forms. See documentation for your operating system for more information.
- **SUSE Linux:** You must install the `glibc-locale-32bit` library that ships with SUSE Linux Enterprise Server; otherwise, LiveCycle will not generate PDF files. This library file is not installed by default, you must use YaST to install it. (See the [SUSE Linux Enterprise Server documentation](#) for details.)

If you plan to install LiveCycle on SUSE Linux 11, you must also install the `libstdc++-libc6.2-2.so.3` libraries. SUSE Linux 11 does not include these libraries by default. For more information, see this [Novell Web](#) page. These libraries are required for running Adobe Central Pro Output Server

3.2.3.4 Configuring the file limit values on Solaris and Linux

To avoid StuckThread issues on a Solaris or Linux environment, add or increase the `rlim` values in the `/etc/system` file.

- 1 **(Linux)** Locate and open the `/etc/security/limits.conf` file.
(Solaris) Locate and open the `/etc/system` file.
- 2 **(Linux)** Add the following lines to the `/etc/security/limits.conf` file:

System Requirements

```
<app_group> soft nfile 8192
<app_group> hard nfile 8192
```

Replace `<app_group>` with the user group who will run the application server. You may also replace `<app_group>` with an asterisk (*) to match all users and user groups.

(**Solaris**) Locate and modify the `rlim` values in the `/etc/system` file as follows:

`set rlim_fd_cur`: The initial (soft) maximum number of file descriptors per process. Set this value to 8192 or more.

`set rlim_fd_max`: The hard maximum number of file descriptors per process. Set this value to 8192 or more. (This modification is required only if the default value is lower than 8192). You must have super user privileges to change this value.

Note: *The `rlim_fd_max` value must be equal to or greater than the `rlim_fd_cur` value.*

- 3 Save and close the file.
- 4 Restart your computer.

Verify the updated settings

- 1 Launch a new shell.
- 2 Type `ulimit -n` and press **Enter**.
- 3 Verify the value returned matches the `rlim` values you have set.

3.2.4 Privileges required to install on Windows

When installing on Windows, you must use an account that has administrator privileges. If you run the installer using a non-administrator account, enter the credentials of an account that has administrator privileges. Turn off the UAC to run the installation and configuration process.

3.2.5 Configuring Windows installation

On 64-bit Windows Server 2008 R1, Windows Server 2008 R2 or Vista operating systems, modify the Admin Approval Mode security option as follows:

- 1 (On Windows 2008 R1 and Vista) Go to **Start > Control Panel > Administrative Tools > Local Security Policy > Local Policies > Security Options**.
(On Windows 2008 R2 and Windows 7) Go to **Start > Administrative Tools > Local Security Policy > Local Policies > Security Options**.
- 2 Locate **User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode** and set it to **Elevate without prompting**.
- 3 Restart your computer.

3.2.5.1 Disable the Windows UAC on Vista and Windows Server 2008 R1

- 1 To access the System Configuration Utility, go to **Start > Run** and then enter **MSCONFIG**.
- 2 Click the **Tools** tab and scroll down and select **Disable UAC**.
- 3 Click **Launch** to run the command in a new window.
- 4 When finished, close the command window and close the System Configuration window.
- 5 Restart your computer.

To enable the UAC again, repeat the steps above and select **Enable UAC** before clicking Launch.

3.2.5.2 Disable the Windows UAC on Windows Server 2008 R2 and Windows 7

- 1 To access the System Configuration Utility, go to **Start > Run** and then enter **MSCONFIG**.
- 2 Click the **Tools** tab and scroll down and select **Change UAC Settings**.
- 3 Click **Launch** to run the command in a new window.
- 4 Adjust the slider to the **Never notify** level.
- 5 When finished, close the command window and close the System Configuration window.
- 6 Restart your computer.

To enable the UAC again, repeat the steps above and adjust the slider to a desired level before restarting your computer.

3.2.6 Additional requirements for PDF Generator

Note: You cannot use the Shared Printer Protocol for the `SendToPrinter` API on Windows 2008 machines that have PDF Generator deployed on them. Use alternate protocols like CIFS or Direct IP.

3.2.6.1 User account for Windows

You must use a user account with administrator privileges for the following tasks:

- Installing Microsoft Office
- Installing PDF Generator
- Installing Acrobat for PDF Generator
- Running the application server process

Note: When you add a users for PDF Generator, grant the user running the application server with the Replace a process level token privilege.

3.2.6.2 User account for Linux and Solaris

You must use a user account with administrator privileges for the following tasks:

- Installing PDF Generator
- Running the application server process
- Running the `sudo` command

Note: When you add a users for PDF Generator, grant the user running the application server with the Replace a process level token privilege.

3.2.6.3 Using 64-bit application servers with PDF Generator

If you are using a 64-bit application server on a system with PDF Generator, ensure that a 32-bit Java 6 JDK is installed in addition to the 64-bit one the application server uses. Set the environment variable `JAVA_HOME_32`. This variable is required to point to a 32-bit JDK on systems where a 64-bit application server is in use. The specific path varies based on the installation directory you specified and the operating system you are installing on.

Note: You need to install the 32-bit Sun JDK and configure `JAVA_HOME_32` to point to the directory where it resides. Review Sun Java 6 Release Notes > Supported System Configurations and download the 32-bit version for your operating system. However, if you are using AIX, do not set `JAVA_HOME_32`.

System Requirements

Important: Ensure that `JAVA_HOME_32` is set only as an environment variable and is not included in the `PATH`. If `JAVA_HOME_32` is included in the `PATH`, Java core dumps may appear during EAR deployment or when you restart the server.

3.2.6.3.1 Set the Windows JAVA_HOME_32 variable

- 1 Select **Start > Control Panel > System**.
- 2 Click the **Advanced System Settings** tab.
- 3 Click **Environment Variables** and, under System Variables, click **New**.
- 4 Enter the environment variable `JAVA_HOME_32`. This value is the directory that contains the JDK. For example, type the following:

```
C:\Program Files (x86)\Java\jdk1.6.0_31
```

3.2.6.3.2 Set the JAVA_HOME_32 variable on Linux or Solaris

Set the `JAVA_HOME_32` variable for the supported JDK for Bourne and Bash shells as shown in this example:

```
JAVA_HOME_32=/opt/jdk1.6.0_31
export JAVA_HOME_32
```

3.2.6.4 Native file conversion software installation

Before you install PDF Generator, install the software that supports the file formats for which PDF conversion support is required and manually activate the licenses for the software using the user account that is used for running the application server process.

Note: On cluster configurations, you must activate one license on each LiveCycle Server of your cluster for each native application that PDF Generator supports.

Refer to the individual licensing agreement for each native application that your LiveCycle deployment will support, and ensure that your LiveCycle deployment meets the licensing requirements specified. Typically, each LiveCycle user who will use native application support must also have an activated license on their own computer for the native application.

PDF Generator can be extended to convert these additional file types to PDF files by using the following applications:

Note: PDF Generator supports only English, French, German, and Japanese versions of the supported applications.

- Microsoft Office 2007, 2010 (DOC, DOCX, XLS, XLSX, PPT, PPTX, RTF, TXT)
- Microsoft Office Visio 2007, 2010 (VSD)
- Microsoft Publisher 2007, 2010 (PUB)
- Microsoft Project 2007, 2010 (MPP)
- AutoCAD files (DWG, DWF) converted through Acrobat
- Corel WordPerfect X4 (WPD)(English only)
- Adobe FrameMaker® 8.0 (FM)
- Adobe PageMaker® 7.0 (PMD, PM6, P65, PM)
- OpenOffice 3.3 (ODT, ODP, ODS, ODG, ODF, SXW, SXI, SXC, SXD, SXM)
- Paper Specification document (XPS) converted through Acrobat

Note: LiveCycle supports only 32-bit editions of all the above mentioned software.

Note: OpenOffice 3.3 or later must be installed on the server to convert the documents created in version 3.3.

System Requirements

Note: On Linux platform, OpenOffice must be installed under /root user. If OpenOffice is installed for specific users, PDFG might not be able to convert OpenOffice documents.

Note: LiveCycle PDF Generator does not support DWF and DWG files created with AutoCAD 2013.

Note: Do not use these software applications on the server. Using these can lead to failed PDF Generator conversions.

You do not need to install a native software application to convert the following native file formats:

- Print files (PS, PRN, EPS)
- Web files (HTML)
- Image files (JPEG, GIF, BMP, TIFF, PNG)
- Flash videos (SWF, FLV)XML

3.2.6.5 Installing Acrobat for PDF Generator

Install Acrobat XI Pro before running the LiveCycle installer. Ensure that you launch Acrobat at least once after installing it to avoid PDF Generator configuration issues. Dismiss all modal dialog boxes that appear on launching Acrobat.

Note: Ensure that Acrobat is installed using the same user account that you will use to install LiveCycle.

However, if LiveCycle is installed and Acrobat XI Pro is not installed, install Acrobat XI Pro and then run the Acrobat_for_PDFG_Configuration.bat script, located in the folder [LiveCycle root]\pdfg_config. Otherwise, PDF conversions will fail.

The Configuration Manager sets the Acrobat_PATH (case-sensitive) environment variable automatically. You can also choose to set it manually, see “3.2.6.8 Setting environment variables” on page 19. Restart your application server after setting the environment variable.

Note: Depending on the JDK that you use, download the required JAR files from either Sun or IBM website.

3.2.6.6 Configure Acrobat to use SHX fonts (Windows only)

Note: Perform these steps to configure Acrobat if you want PDF Generator to use SHX fonts to convert AutoCAD DWG files without installing AutoCAD. Also, these steps need to be performed for all user accounts configured in Administration Console.

- 1 Open Acrobat.
- 2 Select **Edit > Configurations**.
- 3 Select **Convert to PDF > Autodesk AutoCAD**.
- 4 Click **Edit Settings**.
- 5 Click **Configuration Preferences**.
- 6 Click **Browse** next to the SHX Font File Search Path and specify the path to the SHX font file.
- 7 Click **OK** on each opened dialog.

3.2.6.7 QuickTime 7

PDF Generator requires that QuickTime 7 (Player or Pro) be installed if you want to convert video embedded in files, such as PowerPoint presentations to PDF multimedia files. This application is available from the Apple Downloads site.

3.2.6.8 Setting environment variables

You must set the environment variables in Windows if you plan to create PDF documents from applications such as FrameMaker, Photoshop, PageMaker, and WordPerfect.

The names of these environment variables are listed here:

- FrameMaker_PATH
- Notepad_PATH
- OpenOffice_PATH
- PageMaker_PATH
- WordPerfect_PATH
- Acrobat_PATH

These environment variables are optional and need to be set only if you plan to use the corresponding application to convert PDF files through PDF Generator. The value of the environment variable should contain the absolute path of the executable that is used to start the corresponding application.

For example, the variable `FrameMaker_PATH` may contain the value `C:\Program Files\Adobe\FrameMaker7.2\FrameMaker.exe`. However, `OpenOffice_PATH` is different from others. This variable must be set to the OpenOffice installation folder (instead of the path to the executable). A typical value of `OpenOffice_PATH` on Windows would be `C:\Program Files (x86)\OpenOffice.org 3`.

Paths for Microsoft Office applications such as Word, PowerPoint, Excel, Visio, Publisher, and Project or for AutoCAD are not required. The Generate PDF service starts these applications automatically if they are installed on the server.

To set the Windows environment variables

- 1 Select **Start > Control Panel > System**.
- 2 Click the **Advanced System Settings** tab.
- 3 Click **Environment Variables** and, under System Variables, click **New**.
- 4 Enter the environment variable name you need to set (for example, `FrameMaker_PATH`). This value is the directory that contains the executable file. For example, type the following:

```
C:\Program Files\Adobe\FrameMaker7.2\FrameMaker.exe
```

3.2.6.9 Configuring PDF Generator on a Remote Machine

In case of a cluster, LiveCycle is installed only on one machine. Perform the following steps to configure PDF Generator on other machines in the cluster:

- 1 On the remote machine, if an earlier version of Acrobat is installed, uninstall it by using Add or Remove Programs in the Windows Control Panel.
- 2 Install Acrobat XI Pro by running the installer.
- 3 From the machine where LiveCycle is installed, copy `pdfg_config` and `plugins` folders to the remote machine under any directory.
- 4 On the remote machine, open `/pdfg_config/Acrobat_for_PDFG_Configuration.bat` file for editing.
- 5 Locate and comment the `goto locationerror` line.

Before

```
goto locationerror
```

After

```
rem goto locationerror
```

6 Save and close the Acrobat_for_PDFG_Configuration.bat file.

7 Open the command prompt and run the following command:

```
Acrobat_for_PDFG_Configuration.bat <Path of the pdfg_Configuration folder>
```

3.2.6.10 Service Control Manager command line tool

Before you complete an automatic installation of PDF Generator on Windows, ensure that the Service Control Manager command line tool, `sc.exe`, is installed in the Windows environment. Some Windows servers do not have this software preinstalled. By default, the `sc.exe` file is installed in the `C:\Windows\system32` directory. Most OS installations have this tool installed. If you do not have the tool installed, it is available in the Windows Resource Kit for your specific version of Windows. To confirm that the tool is installed on your server, type `sc .exe` from a command prompt. The tools usage is returned.

Note: The LiveCycle must be running as Windows service for the PDF Generator to work properly.

3.2.6.11 Headless mode configuration

If you are running PDF Generator in a headless mode environment (that is, on a server without a monitor, keyboard, or mouse), the x11 libraries must be installed. Some flavors of Linux do not install these libraries by default; therefore, you must obtain the libraries and install them manually.

Note: Activating x11 forwarding on a shell session causes the SOAP UI to create UI elements during SOAP requests, leading to request failures. To avoid request failures, you must add the `-Djava.awt.headless=true` JVM argument.

3.2.6.12 Enabling multi-threaded file conversions and multi-user support for PDF Generator

By default, PDF Generator can convert only one OpenOffice, Microsoft Word, or PowerPoint document at a time. If you enable multi-threaded conversions, PDF Generator can convert more than one of the documents concurrently by launching multiple instances of OpenOffice or PDFMaker (which is used to perform the Word and PowerPoint conversions).

Note: Multi-threaded file conversions (through Microsoft Office) are only supported for Microsoft Word 2007 or 2010 and PowerPoint 2007 or 2010.

Note: Microsoft Excel, Publisher, Project, and Visio files are not converted simultaneously. During conversion, `EXCEL.exe`, `PUBLISHER.exe`, `PROJECT.exe`, and `VISIO.exe` are watched in the task manager. The Visio file is processed only after processing the Excel file is complete.

Each instance of OpenOffice or PDFMaker is launched using a separate user account. Each user account that you add must be a valid user with administrative privileges on the LiveCycle Server computer. For more information, see “[3.2.5 Configuring Windows installation](#)” on page 15

In a clustered environment, the same set of users must be valid for all nodes of the cluster.

After your LiveCycle Server is configured, add LiveCycle user accounts in Administration Console. See the User accounts for multi-threaded file conversions section in the LiveCycle installation guide for your application server. To enable multiuser support for native files and OpenOffice files on a Windows environment, add a minimum of three users with the following permissions.

When you add users for PDF Generator native conversions, grant the user running the application server with the Replace a process level token privilege. For more information, see “[3.2.6.14 Granting the Replace a process level token privilege \(Windows only\)](#)” on page 22

Dismiss initial dialogs and disable automatic updates for native applications

Converting native files from PDF Generator requires dismissing any initial registration, activation, and Improvement program dialogs with the option to not show them again. Automatic updates for these applications also needs to be disabled as these update dialogs can cause failures on a running server.

The dialogs and automatic update need to be disabled for the user running the server and all user accounts configured under PDFG Accounts for multi-user support. The dialogs need to be dismissed for the following applications if installed on the server:

- Microsoft Word (Windows)
- Microsoft Excel (Windows)
- Microsoft PowerPoint (Windows)
- Microsoft Project (Windows)
- Microsoft Publisher (Windows)
- Microsoft Visio (Windows)
- OpenOffice (Windows/ Solaris / Linux)
- Adobe Acrobat (Windows)
- Adobe FrameMaker (Windows)
- Adobe PageMaker (Windows)
- Autodesk AutoCAD (Windows)
- Corel WordPerfect (Windows)

Note: Ensure that you launch Adobe Acrobat Distiller at least once for all the PDFG user accounts configured on the server.

Disable error reporting on Windows Server 2008 (Optional but recommended)

While converting a document to PDF using PDF Generator on Windows Server 2008, Windows may report that the executable has encountered a problem and needs to close. However, it does not impact the PDF conversion as it continues in the background.

To avoid receiving the error, you can disable the Windows error reporting. For more information on disabling error reporting, see <http://technet.microsoft.com/en-us/library/cc754364.aspx>.

Additional configuration required for OpenOffice on Linux or Solaris

- 1 Add entries for additional users (other than the administrator who runs the LiveCycle Server) in the `/etc/sudoers` file. For example, if you are running LiveCycle as a user named `lcadm` and a server named `myhost`, and you want to impersonate `user1` and `user2`, add the following entries to `/etc/sudoers`:

```
lcadm myhost=(user1) NOPASSWD: ALL
```

```
lcadm myhost=(user2) NOPASSWD: ALL
```

This configuration enables `lcadm` to run any command on host `'myhost'` as `'user1'` or `'user2'` without prompting for password.

- 2 Allow all the LiveCycle users to make connections to the LiveCycle Server. For example, to allow a local user named `user1` the permission of making the connection to the LiveCycle Server, use the following command:

```
xhost +local:user1@
```

Ensure that the session with which the application server started should not get closed.

For more details, refer to xhost command documentation.

- 3 Restart the server.

3.2.6.13 Multi-user support for PDF Generator

To enable multi-user support for native files and OpenOffice files on a Windows environment, a minimum of three users with the following permissions must be added. On a Linux or Solaris platform, create at least one user.

Platform	User permissions
Windows 2008 Server	Users with administrative privileges, Read/write permission on LiveCycle temporary directory, PDF Generator temporary directory and application server installation directory.
Linux and Solaris	Users with <code>sudo</code> privileges Read/write permission on LiveCycle temporary directory, PDF Generator temporary directory, and application server installation directory.

Note: For clusters, the users you create must have the above permissions on all nodes.

When you add users for PDF Generator native conversions, you must grant the user running the application server with the *Replace a process level token* privilege. See “[3.2.6.14 Granting the Replace a process level token privilege \(Windows only\)](#)” on page 22.

3.2.6.14 Granting the Replace a process level token privilege (Windows only)

User account that are used to start the application server should be part of the local administrators group and requires the *Replace a process level token* privilege. To provide *Replace a process level token* privilege:

- 1 Click Start > Run, and then type `gpedit.msc`.
- 2 On the Group Policy dialog box, select **Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignment**, and double click **Replace a process level token**.
- 3 Click **Add User or Group**, add the Windows user account that is used to open the command prompt from which the application server is started.
- 4 Restart Windows, and then start the application server.

3.2.6.15 Symbolic link on Linux platform

To substitute required fonts in a HTML-to-PDF conversion on the Linux platform, PDF Generator creates a symbolic link that point to the `/usr/share/X11/fonts` directory.

Sometimes the user running the application server might not possess permissions that are necessary to create a symbolic link. On such systems; create a symbolic link `/usr/lib/X11/fonts` that point to the `/usr/share/X11/fonts` directory.

3.2.6.16 Symbolic link on Solaris 11 platform

In Solaris 11 some of the fonts required for HTML to PDF conversion are moved from `/usr/openwin/lib/X11/fonts` location to `/usr/share/fonts` location. To allow PDF Generator to access these fonts, create a symbolic link at `/usr/openwin/lib/X11/fonts`, referring to `/usr/share/fonts` location. To perform HTML-to-PDF conversion on Solaris 11 platform, do the following:

- 1 Open the Terminal Window

2 Run the following command:

```
ln -s /usr/share/fonts /usr/openwin/lib/X11/fonts/usr_share_fonts
```

3.2.6.17 Additional requirement for Windows Server 2008 and Windows Server 2008 R2

Install the hotfix for your Windows Server 2008 version as described in the following articles:

<http://support.microsoft.com/kb/972616> for Windows Server 2008

<http://support.microsoft.com/kb/2815716> for Windows Server 2008 R2

3.2.6.18 Additional requirements for Red Hat Enterprise Linux 6 (RHEL6)

PDF Generator requires additional RPM packages and fonts to perform conversions on RHEL6. Perform the following steps to configure the PDF Generator on RHEL6:

1 Install the following RPM packages from RHEL6 installation media:

- glibc-2.12-1.25.el6.i686.rpm
- nss-softokn-freebl-3.12.9-3.el6.i686.rpm
- libX11-1.3-2.el6.i686.rpm
- libxcb-1.5-1.el6.i686.rpm
- libXau-1.0.5-1.el6.i686.rpm
- zlib-1.2.3-25.el6.i686.rpm
- libXext-1.1-3.el6.i686.rpm
- fontconfig-2.8.0-3.el6.i686.rpm
- expat-2.0.1-9.1.el6.i686.rpm
- freetype-2.3.11-6.el6_0.2.i686.rpm
- libSM-1.1.0-7.1.el6.i686.rpm
- libICE-1.0.6-1.el6.i686.rpm
- libuuid-2.17.2-12.el6.i686.rpm
- libXrandr-1.3.0-4.el6.i686.rpm
- libXrender-0.9.5-1.el6.i686.rpm
- libXinerama-1.1-1.el6.i686.rpm

2 In your browser, open website <http://cgit.freedesktop.org/xorg/font/ibm-type1/>

3 Download the compressed file `font-ibm-type1-1.0.3.tar.gz` or `font-ibm-type1-1.0.3.zip`. The compressed file contains required fonts.

4 Extract the downloaded zip file to the `/usr/share/fonts` directory.

3.2.6.19 Configuring user accounts for multi-threaded file conversions

By default, PDF Generator can convert only one OpenOffice, Microsoft Word, or PowerPoint document at a time. If you enable multi-threaded conversions, PDF Generator can convert more than one of the documents concurrently by launching multiple instances of OpenOffice or PDFMaker (which is used to perform the Word and PowerPoint conversions).

Note: Only Microsoft Word 2007 and Microsoft PowerPoint 2007 are supported with multi-threaded file conversions. Microsoft Excel 2003 or 2007 versions are not supported.

If you need to enable multi-threaded file conversion, you must first perform the tasks outlined in the “Enabling multi-threaded file conversions” section of the Preparing to Install or Upgrade guide available on the [LiveCycle documentation](#).

For Linux and Solaris users, you must create users and configure the system to remove the password prompts. The following section outlines the method to create a user and perform additional configurations.

3.2.6.19.1 Add user account

- 1 In Administration Console, click **Services > LiveCycle PDF Generator 11 > User Accounts**.
- 2 Click **Add** and enter the user name and password of a user who has administrative privileges on the LiveCycle Server. If you are configuring users for OpenOffice, dismiss the initial OpenOffice activation dialogs.

Note: If you are configuring users for OpenOffice, the number of instances of OpenOffice cannot be greater than number of user accounts specified in this step.

- 3 Restart the LiveCycle Server.

Note: Ensure that the added user account is defined for all the nodes of the cluster.

3.2.6.19.2 Additional configuration required for OpenOffice on Linux or Solaris

- 1 Add user accounts as described above.
- 2 Add entries for additional users (other than the administrator who runs the LiveCycle Server in the `/etc/sudoers` file. For example, if you are running LiveCycle as a user named `lcadm` on a server named `myhost`, and you want to impersonate `user1` and `user2`, add the following entries to `/etc/sudoers`:

```
lcadm myhost=(user1) NOPASSWD: ALL
lcadm myhost=(user2) NOPASSWD: ALL
```

This configuration enables `lcadm` to run any command on host ‘`myhost`’ as ‘`user1`’ or ‘`user2`’ without prompting for password.

- 3 Allow all the users that you added via Add a user account to make connections to the LiveCycle Server. For example, to allow a local user named `user1` the permission of making the connection to the LiveCycle Server, use the following command:

```
xhost +local:user1@
```

For more details, refer to `xhost` command documentation.

- 4 Enable `requiretty` in the `/etc/sudoers` file.
- 5 Restart the server.

Note: Ensure the session with which the application server is started, is kept open. On closing the session some conversions may fail intermittently.

3.2.7 Additional requirements for Connector for Documentum

If LiveCycle is connecting to Documentum, you must install Document Foundation Classes on machine hosting LiveCycle.

3.2.8 Additional requirements for Connector for IBM Content Manager

- DB2 Universal Database Client
- IBM Information Integrator for Content (II4C)

See “Post-Deployment Activities” chapter in the Installing and Deploying LiveCycle document for your application server.

3.2.8.1 Configure the connection for a single IBM Content Manager datastore

- 1 Start the DB2 Configuration Assistant.
- 2 Click **Selected > Add Database Using Wizard**.
- 3 Select **Manually Configure a Connection to a Database** and click **Next**.
- 4 Select **TCP/IP** and click **Next**.
- 5 Specify the following TCP/IP communication options and then click **Next**:
 - In the **Host Name** box, type the host name of the server hosting DB2 Content Manager.
 - Leave the Service Name box empty.
 - In the **Port Number** box, type the port number. The default DB2 Content Manager port number is 50000.
- 6 In the **Database Name** box, type the IBM Content Manager datastore name and, in the **Database Alias** box, type the alias name for the datastore and then click **Next**.
- 7 Click **Next** to accept the default data source settings.
- 8 In the **Operating System** list, select the operating system you are using and then click **Next**.
- 9 Specify the following system options and then click **Next**:
 - In the **System Name** box, type the server name hosting DB2. If you click Discover, DB2 Content Manager searches for the system name you specified and, if the system is not found, all of the DB2 instances are listed.
 - In the **Host Name** box, type the name of the host, or click View Details to show the domain and IP address of the system you named in the previous step.
 - In the **Operating System** list, select the operating system (Windows, Linux, or AIX) on which you deployed DB2 Content Manager.
- 10 (Optional) To specify Security options, select **Use Authentication Value in Server’s DBM Configuration** and click **Finish**.
- 11 In the Test Connection dialog box, test the connection as required.

3.2.8.2 Configure connections for multiple IBM Content Manager datastores

- 1 Configure the initial connection by following the steps in “[3.2.8.1 Configure the connection for a single IBM Content Manager datastore](#)” on page 25.
- 2 Add additional database connections by modifying the cmbicmsrvs.ini file (the file that stores the datastore information) as follows:
 - From a command prompt window, change the directory to *[II4C home]/bin* (for example, C:\Program Files\db2cmv8\ on Windows **or** /opt/IBM/db2cmv8 on AIX, Linux, or Solaris).
 - Run the cmbenv81.bat (Windows) or cmbenv81.sh (AIX, Linux, or Solaris) file to set the environment and the classpath for the Java Utilities of II4C.
 - Change the directory to *[II4C working directory]/cmgmt/connectors* where *[II4C working directory]* is one of the following paths:
 - (Windows) C:/Program Files/db2cmv8
 - (Linux, AIX) /home/ibmcmadm
 - (Solaris) /export/home/ibmcmadm

System Requirements

- Run the command

```
java com.ibm.mm.sdk.util.cmbsrvsictm -a add -s <library server database name> -sm <database schema name>
```

where <library server database name> is the same as Database Alias configured in step 6 above.

Note: The following procedure allows users without DB2 rights to share the connection credentials through the cmbicmenv.ini file.

3.2.8.3 Configure a multiuser connection to the IBM Content Manager datastore

- 1 From a command prompt window, change the directory to *[II4C home]/bin* (for example, C:\Program Files\db2cmv8\ on Windows **or** /opt/IBM/db2cmv8 on AIX, Linux, or Solaris).
- 2 Run the cmbenv81.bat (Windows) or cmbenv81.sh (AIX, Linux, or Solaris) file to set the environment and the classpath for the Java Utilities of II4C.
- 3 Change the directory to *[II4C working directory]/cmgmt/connectors*, where *[II4C working directory]* is one of the following paths:

(Windows) C:/Program Files/db2cmv8

(Linux, AIX) /home/ibmcmadm

(Solaris) /export/home/ibmcmadm

- 4 Run the command

```
java com.ibm.mm.sdk.util.cmbenvicm -a add -s <library server database name> -u <database user ID> -p <database password>
```

where <library server database name> is the same as Database alias configured in step 6 above.

3.2.9 Additional requirements for Connector for IBM FileNet

These requirements are optional and required only if you are installing Adobe® LiveCycle® 11 Connector for IBM® FileNet.

IBM FileNet 4.0

If LiveCycle is connecting to IBM FileNet 4.0 Content Engine, you must install the Content Engine Java Client. Use the IBM FileNet 4.0 content engine client installer located by default in C:\Program Files\FileNet\CEClient. During installation, select only the Java client component on the component selection screen.

For IBM FileNet 4.0 Process Engine settings, copy the pe.jar file from the Process Engine directory to the computer that will host LiveCycle. Create the directory C:\FileNetPE\files and copy the pe.jar file there. The Process Engine client install directory is now C:\FileNetPE.

IBM FileNet 4.5

If LiveCycle is connecting to IBM FileNet 4.5 Content Engine, you must install the Content Engine Java Client. Use the IBM FileNet 4.5 content engine client installer located by default in C:\Program Files\FileNet\CEClient. During installation, select at least one of the components from Application Engine or Process Engine on the component selection screen.

For IBM FileNet 4.5 Process Engine, you must install the IBM FileNet 4.5 Process Engine Client located by default in C:\Program Files\FileNet\BPMClient. During installation, select the **Other** option on the component selection screen.

IBM FileNet 5.0

If LiveCycle is connecting to IBM FileNet 5.0 Content Engine, you must install the Content Engine Java Client. Use the IBM FileNet 5.0 content engine client installer located by default in C:\Program Files\FileNet\CEClient. During installation, select at least one of the components from Application Engine or Process Engine on the component selection screen.

For IBM FileNet 5.0 Process Engine, you must install the IBM FileNet 5.0 Process Engine Client located by default in C:\Program Files\FileNet\BPMClient. During installation, select the Other option on the component selection screen.

3.2.10 Additional requirement for Forms, Output and ConvertPDF on Windows based cluster

Microsoft Visual C++ 2008 SP1 Redistributable Package (x86) is installed along with the LiveCycle installation. For Cluster, LiveCycle installation is not performed on all the nodes so the redistributable package is not installed all the nodes.

On Windows based clusters, install Microsoft Visual C++ 2008 SP1 Redistributable Package (x86) or latest available redistributable package on all the nodes of cluster.

3.3 LDAP configuration

This configuration is optional and required only if you are using an LDAP directory to authenticate users.

If you do not have an existing LDAP server and database, install and configure your LDAP server and database according to the vendor's documentation. Make note of the LDAP administrator name and password to use during the LiveCycle configuration process. Configure LiveCycle to connect with the LDAP database after you install and deploy your services that are part of LiveCycle. This configuration is done by using the User Manager service.

3.4 Global document storage directory

The global document storage (GDS) directory is used to store long-lived files that are used within a process as well as critical LiveCycle product components. The lifetime of long-lived files is intended to span multiple restarts of a LiveCycle system, and can span days and even years. These files may include PDF files, policies, or form templates.

You must create the GDS directory before you initialize the LiveCycle database. See [“3.4.1 Location of the global document storage directory”](#) on page 28 for information.

On cluster configurations, create a shared file system for the GDS directory on any computer that is accessible by each node in the LiveCycle cluster, and ensure that all nodes of the cluster have read and write permissions for the directory.

On cluster configurations, you must specify the shared directory that you create for GDS. You cannot use the default location in a cluster environment.

Note: Create a shared file system for the GDS directory on any computer that is accessible by each node in the LiveCycle cluster, and ensure that all nodes of the cluster have read and write permissions for the directory.

3.4.1 Location of the global document storage directory

You configure the location of your GDS directory with Configuration Manager after you install LiveCycle. The GDS directory you specify should be highly available and should have low access time to enhance performance. If the GDS directory is on a shared network drive, it is recommended that you specify the location in UNC style as `\\computer_name\GDS`.

If you must change the GDS directory location after completing the installation (see [LiveCycle Administration Help](#)), you should plan an appropriate location for the GDS directory.

Important: *Module deployment will fail on Windows if the GDS directory is at the drive root (for example, D:\). For GDS, you must make sure that the directory is not located at the root of the drive but is located in a subdirectory. For example, the directory should be D:\GDS and not simply D:\.*

3.4.2 Sizing factors for the global document storage directory

The size of the global document storage directory depends on expected LiveCycle usage factors for the deployment. You should allocate a minimum of 10 GB of disk space for the GDS directory.

The following factors also affect the sizing:

- The typical volume of documents that LiveCycle processes. Processing high volumes of documents requires a larger GDS directory.
- The typical size of documents that LiveCycle processes. Processing large documents requires a larger shared GDS directory.
- The complexity of documents that LiveCycle processes. Processing complex documents, such as documents that are processed by multiple services that are part of LiveCycle, require a larger GDS directory.

On cluster configurations, in addition, documents that are processed by different members of the cluster require a larger GDS directory.

3.4.3 Securing the global document storage directory

Access to the GDS directory must be secure. The long-lived documents in this directory may contain sensitive user information, such as information that requires special credentials when accessed by using the LiveCycle SDK or user interfaces.

Use a security method that is appropriate to your operating system. It is recommended that only the operating system account that is used to run the application server has read and write access to this directory.

Note: *Deleting files or directories from the GDS directory can render the LiveCycle server inoperative.*

3.4.4 Backing up the global document storage directory

The global document storage directory should be backed up to allow administrators to restore LiveCycle in case of failure.

If the global document storage directory becomes unavailable or is lost due to failure, LiveCycle will not run until the GDS directory and database are restored by a consistent back up or LiveCycle is reinitialized with a new installation.

3.5 Synchronizing clock times

You must ensure that all computers in a horizontal cluster synchronize their clock times regularly. Your LiveCycle installation may encounter problems if the node times differ by more than a few seconds.

Apply the standard time synchronization practices employed by your network to all computers of the LiveCycle cluster.

3.6 Installation considerations

3.6.1 Disabling Virus scans (Windows only)

To improve the speed of the installation, you may choose to disable any on-access virus scanning software for the duration of the installation.

3.6.2 Installing from network drives

It is recommended that you install LiveCycle only from the installation media or a local disk drive. Attempting to install the software over a network results in considerable delays in starting and installing. It is also likely that installing from a network drive will add to the length of the directory path, which will cause the LiveCycle installer to prevent the installation from proceeding.

3.6.3 Using LiveCycle with a Luna HSM cluster

When using a SafeNet Luna ethernet-attached Hardware Security Module (HSM) cluster, you must ensure HAOnly mode is enabled on the device.

3.6.3.1 Enable HAOnly mode on the Luna device

1 Use the vtl tool shipped with the Luna client to determine if HAOnly mode is enabled. Type:

```
vtl haAdmin -HAOnly -show
```

2 1) If HAOnly mode is not enabled, type:

```
vtl haAdmin -HAOnly -enable
```

3.6.4 Manual use of Acrobat restricted

If you installed the PDF Generator for native document conversion, use of the bundled Acrobat installation is restricted to the Generate PDF service and is not licensed for any other use.

3.6.5 Temporary directory

LiveCycle requires a temporary directory to store documents that are larger than the maximum inline size set for document objects. For detailed information on how you can improve the performance of your LiveCycle Server by specifying an appropriate value for this setting, see this [blog on the Document Max Inline Size performance knob](#).

(AIX, Linux, and Solaris only) If a non-root user is running the application server, the user must have read, write, and execute permissions on the specified temporary directory.

Important: (For cluster deployments only) Future upgrades might fail if you specify a shared network directory as the temporary directory.

3.6.5.1 Temporary directory on cluster configurations

Horizontal cluster: The location of the temporary directory can be either unique or shared for a horizontal cluster. Create a unique temporary directory on each node of the cluster to eliminate network traffic to a shared drive or directory.

Vertical cluster: Create a shared temporary directory. Typically, the shared folder is on the same local system as the vertical cluster.

In addition to the shared temporary folder, each server creates its unique subdirectories to store temporary data. The subdirectory name is based on the corresponding server name.

(JBoss Only) To ensure each JBoss instance has a unique name, open the run.bat file, and manually modify the `-Dadobeidp.serverName=<servername>` JVM argument. The default value is `server1`.

3.7 LiveCycle IPv6 support

LiveCycle includes IPv6 support. The default configurations defined in the installation documentation for LiveCycle set IPv4 as the default IP protocol because this protocol has the most compatibility with third-party infrastructure.

Do not enable IPv6 unless your deployment must use it. The number of supported platform configurations is reduced when enabling IPv6 support with LiveCycle. You should verify that all third-party software, hardware, and networks that you plan to use have IPv6 support before you attempt to enable IPv6.

Note: If you are enabling CIFS in an IPv6 environment, you must explicitly enable IPv6 configuration after you configure your LiveCycle installation using Configuration Manager. See “Enabling CIFS in IPv6 mode” in the guide for your application server.

3.7.1 Supported IPv6 configurations

Not all infrastructure components support IPv6. For example, Oracle database does not support IPv6. You can use these databases by configuring the connection between the application server and the databases with IPv4, and the remaining communications over IPv6.

Check with your component vendor if IPv6 is supported.

3.7.2 IPv6 implementation guidelines

When you use IPv6 implementation either partially or fully, keep the following points in mind:

- After installing LiveCycle, do not use the option to start the Configuration Manager directly from the LiveCycle installer. Instead, navigate to the `[LiveCycle root]\configurationManager\bin\IPv6` directory, and run the IPv6-specific script (`ConfigurationManager_IPv6.bat` or `ConfigurationManager_IPv6.sh`) to launch the Configuration Manager.
- If you have chosen to validate the application server configuration using the Configuration Manager, the validation will fail after you enable IPv6 for the application server. You can ignore this error message during the process. After you restart the application server in the IPv6 mode, the application server can connect to the database.

System Requirements

- (*WebLogic only*) You need to enable IPv6 for the managed servers only. The Admin Server can continue to run on IPv4, and can be accessed with its IPv4 address. However, the managed server that you have started in IPv6 environment can only be accessed through its IPv6 address or a hostname resolved through DNS.
- (*WebLogic only*) Even if you are running Configuration Manager on the same computer that hosts the application server, you must provide the listen address of the managed server for bootstrapping and deploying LiveCycle modules. This listen address must be the DNS name that resolves to IPv6 address of the computer.
- To have a pure IPv6 communication with the database server, modify both EDC_DS and IDP_DS connection settings to use the hostname of the database which resolves to a numeric IPv6 address.
- If you are installing LiveCycle on to a server cluster, you must map the numeric IPv6 addresses of each cluster node to the computer's host name in DNS or in the `hosts` file on each cluster node. The `hosts` file is located at:
 - Solaris: `/etc/inet/ipnodes`
 - Windows: `C:\Windows\system32\drivers\etc\hosts`
 - Linux: `/etc/hosts`
- Many software components such as database drivers do not completely support numeric IPv6 addresses. So, it is recommended that you use a DNS-resolved hostname instead of numeric IPv6 addresses.
- Ensure that name used for mapping IPv6 is added to the CSRF filter section. If the name is not added, see Preventing CSRF attacks in [Administration Help](#).

Note: Name used for mapping IPv6 must not contain square brackets (`[]`).

- In an IPv6 environment, if you are using Microsoft SQL Server, you should specify the database server IP address in the following format. Note that in this string, `;serverName` is a keyword, and so must not be replaced with the actual server name.

```
jdbc:sqlserver://;serverName=<IPv6 address>;portNumber=<port>;databaseName=<db_name>
```

Here, instead of the numeric IPv6 address, you can specify the hostname of the SQL Server database.

3.7.3 Configuring IPv6 for JBoss

- 1 You can download and install JBoss from <http://www.jboss.org/jbossas/downloads/> or obtain the `jboss.zip` file from the third-party directory on the installation media and extract the bundled JBoss.
- 2 Modify `adobe-ds.xml` and the database-specific data source configuration file to connect to the LiveCycle database.
- 3 Modify the `login-config.xml` file to connect to the LiveCycle database.
- 4 For clusters, modify the `jgroups-channel-factory-stacks.xml` files to enable IPv6. Refer to *Configuring LiveCycle Application Server Clusters Using JBoss*.
- 5 Modify the following files to enable IPv6:
 - (**JBoss on Windows**) `[appserver root]\bin\run.conf.bat`
 (**JBoss on other platforms**) `[appserver root]\bin\run.conf`
 - Change `-Djava.net.preferIPv4Stack=true` to `-Djava.net.preferIPv6Stack=true`.
 - Add the `-Djava.net.preferIPv6Addresses=true` argument.
 - (**JBoss 5.2 only on Windows**) `[appserver root]\bin\run.bat`
 (**JBoss 5.2 only on other platforms**) `[appserver root]\bin\run.sh`
 - Delete `-Djava.net.preferIPv4Stack=true`.

System Requirements

- `[appserver root]\bin\service.bat`
 - Replace `- b 0.0.0.0` with `-b <hostname mapped to IPv6 address>`.
- 6 Launch Configuration Manager by invoking the `[LiveCycle root]\configurationManager\bin\IPv6\ConfigurationManager_IPv6.bat` or `ConfigurationManager_IPv6.sh` script.
- 7 In the Configuration Manager, select the steps to configure EAR files, bootstrap and deploy LiveCycle modules.
- 8 After the Configuration Manager process is completed, copy these EAR files to the `[appserver root]\server\<profile_name>\deploy` directory.
- 9 Start JBoss from a command line.
- 10 Provide Configuration Manager hostname of the computer that is mapped to its IPv6 address and then bootstrap the application server to deploy the LiveCycle modules.

3.7.4 Configuring IPv6 for WebLogic

- 1 Install LiveCycle using the installer.
- 2 Do not launch the Configuration Manager when the installer finishes. Launch Configuration Manager by invoking the `[LiveCycle root]\configurationManager\bin\IPv6\ConfigurationManager_IPv6.bat` or `ConfigurationManager_IPv6.sh` script.
- 3 Select the options to configure LiveCycle EARs, WebLogic application server, and verify the application server settings using Configuration Manager.

Note: You will get an error message that the Configuration Manager validation of data source has failed. This is because the application server is not yet started in the IPv6 mode, and the data source is now configured in the IPv6 mode. You can ignore this warning at this stage.

- 1 From the WebLogic Server Administration Console, change the application **Server Start** arguments of the managed server to enable IPv6.
 - Change `-Djava.net.preferIPv4Stack=true` to `-Djava.net.preferIPv6Stack=true`.
 - Add the `-Djava.net.preferIPv6Addresses=true` argument.
- 2 Change the listen address of the managed server to enable it using IPv6 address.
 - In the WebLogic Server Administration Console, select **Environment** > **Servers** > `[managed server name]`**Configuration** tab.
 - In the Listen Address field, enter the hostname of the computer. Ensure that this hostname resolves to the IPv6 address of this computer.
- 3 Save the changes, and then restart the managed server.
- 4 Launch Configuration Manager by invoking the `[LiveCycle root]\configurationManager\bin\IPv6\ConfigurationManager_IPv6.bat` or `ConfigurationManager_IPv6.sh` script.
- 5 In the Configuration Manager, select the steps to deploy EAR files, bootstrap, deploy LiveCycle modules.
- 6 Enter the same hostname as provided in the listen address field of the managed server.

Note: Even if you are running Configuration Manager on the same computer, you must provide the listen address of the managed server for bootstrapping and deploying LiveCycle modules.

3.7.5 Configuring IPv6 for WebSphere

- 1 Install LiveCycle using the installer script. After the installation is complete, do not start the Configuration Manager when the installer prompts.

System Requirements

- 2 Navigate to the `[LiveCycle root]\configurationManager\bin\IPv6` directory, and run the IPv6-specific script (`ConfigurationManager_IPv6.bat` or `ConfigurationManager_IPv6.sh`) to launch LCM.
- 3 Use the Configuration Manager options to configure EAR and the application server.
- 4 Follow the steps in the Configuration Manager to configure the application server. While configuring database, provide hostname of database that is mapped to IPv6 address.
- 5 Validate the application server configuration by Configuration Manager. Ignore the warning if data source validation fails. You can validate data sources from WebSphere Administrative Console.
- 6 On the WebSphere Administrative Console, select **Servers > Server Types > WebSphere application servers > [server name] > Java and Process Management > Process definition > Java Virtual Machine**. In **Generic JVM arguments**, add the `-Djava.net.preferIPv6Addresses=true` argument.
- 7 Select **Servers > Server Types > WebSphere application servers > [server name] > Java and Process Management > Process definition > Java Virtual Machine**, and click **Custom Properties**. Change `java.net.preferIPv4Stack=true` to `java.net.preferIPv6Stack=true`.
- 8 Manually deploy the EAR files to WebSphere Application Server by using WebSphere Administrative Console. Configured EAR files are available at `[LiveCycle root]/configurationManager/export` folder.
- 9 Restart WebSphere Application Server.
- 10 Navigate to the `[LiveCycle root]\configurationManager\bin\IPv6` directory, and run `ConfigurationManager_IPv6.bat` or `ConfigurationManager_IPv6.sh` to launch Configuration Manager.
- 11 In the Configuration Manager, select the options to bootstrap and deploy LiveCycle modules. Provide the application server hostname that is mapped to IPv6 address.

Note: After you start the application server in an IPv6 environment, (with the flag `-Djava.net.preferIPv6Stack=true`), you can access it only through its IPv6 address or hostname mapped to IPv6 address.

3.8 Central Migration Bridge service

The Central Migration Bridge service helps you to migrate existing applications from your Adobe Central Pro Output Server or Adobe Web Output Pak products to work with the Output service. The Central Migration Bridge service allows you to use your current IFD/MDF templates, data transformation scripts, and DAT files in a LiveCycle environment as part of an overall migration strategy.

Note: The Central Migration Bridge is useful only if you have existing Central Pro applications to migrate.

Authorized use of Central Migration Bridge

To use the Central Migration Bridge service, you must have a valid license for Central Pro Output Server 5.7 or an executed Central Pro Output Server 5.7 migration agreement. To install Central Pro Output Server 5.7, use your existing media and existing product authorization code (PAC). Your PAC is for a specific operating system platform; if this differs from the operating system platform on which you are installing LiveCycle, you will need to acquire a PAC for that operating system. Contact your Adobe representative for more information regarding migration, or acquiring Central Pro Output Server 5.7 media or PAC.

Installation considerations

The Central Migration Bridge service interacts directly with the Central Pro (version 5.7) executable files. You must have Central Pro installed on the same server as LiveCycle, but it is not a prerequisite to installing LiveCycle (that is, it can be installed before or after LiveCycle). Consult the Central Pro documentation set for installation instructions.

System Requirements

Important: Do not start Central Pro or change its properties to run automatically.

On Windows, the Central Pro service *Adobe Central Output Server* is installed as a manual service. Do not run the service or change its properties to run automatically.

On AIX, Linux, or Solaris, do not start the Central Pro daemon *jfdaemon*. If you edited the start-up script of your computer to launch *jfdaemon* when restarting the computer, change the script to prevent the daemon from being launched automatically. (See the Central Pro installation documentation.) Do not start Central from the command line by launching the *jfserver* process.

Note: The LiveCycle User that invokes the Central Migration Bridge service must have access rights to the Central Pro install directory as well as execute permissions on the Central Pro executable files.

LiveCycle turnkey installations

If you are installing and configuring a LiveCycle turnkey environment using Express mode, the Central Migration Bridge service will be installed and configured by default - you will not be prompted for any input.

Note: Ensure that the Adobe Central Pro product is installed in the default directory.

LiveCycle custom installations

If you are installing and configuring LiveCycle using the Custom mode (Partial turnkey or Manual), Configuration Manager will prompt you to include Central Migration Bridge in the deployment.

By default, the service will use the default Central Pro installation path. If Central Pro is installed in a different location, update the [Central Install Dir] configuration for the Central Migration Bridge Service by going to Administration Console.

Note: In a clustered environment, you must install Central Pro on each node of the cluster.

When you have finished installing LiveCycle, if Central Pro is not installed in the default location, do the following steps to point LiveCycle to the correct directory:

- 1 Log in to Administration Console.
- 2 Click **Services > Applications and Services > Service Management**.
- 3 Click the **Central Migration Bridge:1.0** service.
- 4 Type the correct path to the Central Pro install directory.
- 5 Click **Save**.

Note: This configuration can also be made in Workbench. See *Creating and managing processes*.

Chapter 4: LiveCycle Credentials and Certificates

This chapter describes how to do the following tasks:

- Obtain the Adobe® LiveCycle® Reader® Extensions 11 Rights credential
- Obtain digital certificates for use with Adobe® LiveCycle® Digital Signatures 11.

4.1 Obtaining the Reader Extensions Rights credential

The Reader Extensions Rights credential is a digital certificate that is specific to Reader Extensions that enables Adobe Reader usage rights to be activated in PDF documents. If the credential is not installed, Reader Extensions users cannot apply usage rights to documents. You cannot use a standard digital certificate for this function; you must use the dedicated Rights credential.

The Rights credential extends the usage rights of each PDF file that Reader Extensions processes. It is a critical part of the software licensing and should be stored carefully in a secure environment.

The following types of Rights credentials are available:

Customer Evaluation: A credential with a short validity period that is provided to customers who want to evaluate Reader Extensions. Usage rights applied to documents using this credential expire when the credential expires. This type of credential is valid only for two to three months.

Production: A credential with a long validity period that is provided to customers who purchased the full product. Production credentials are unique to each customer but can be installed on multiple systems.

The Rights credential is delivered as a digital certificate that contains the public key, the private key, and the password used to access the credential.

If your organization orders an evaluation version of Reader Extensions, you receive an evaluation Rights credential from the sales representative you ordered the product from or from the website where you downloaded the evaluation product.

If your organization purchases a production version of Reader Extensions, the production Rights credential is delivered by Electronic Software Download (ESD). A production Rights credential is unique to your organization and can enable the specific usage rights that you require.

If you obtained Reader Extensions through a partner or software provider who integrated Reader Extensions into their software, the Rights credential is provided to you by that partner who, in turn, receives this credential from Adobe.

Note: *The Rights credential cannot be used for typical document signing or assertion of identity. For these applications, you can use a self-sign certificate or acquire an identity certificate from a Certificate Authority (CA).*

4.2 Obtaining digital certificates for use with Digital Signatures

Digital certificates are required for use with Digital Signatures. Although you can configure and manage digital certificates after you install and configure LiveCycle, obtaining them before you install ensures that you are ready to use LiveCycle when it is deployed.

Digital certificates are obtained from a Certificate Authority (CA) and sent to you by email or over the web as a certificate file. This certificate file contains the public keys (also called *certificates*) and references to private keys (also called *credentials*) that are used for encrypting and signing documents. Certificates do not contain actual private keys; instead, they contain a reference to the identity of the user who keeps the private keys securely stored in an encrypted file or HSM.

You can use Internet Explorer (Windows) or OpenSSL (AIX, Linux, and Solaris) to export PFX, P12, and CER files for certificates that are stored in any compatible certificate store that is available on your computer. PFX files can be exported only as the certificate store or the credential itself permits. CER files that hold the public key that corresponds to a credential can also be exported from PFX files by using either Internet Explorer or OpenSSL.

Note: *You can configure and manage certificates, credentials, and Certification Revocation Lists (CRLs) for use with LiveCycle by using Trust Store Management, which is accessible through the web-based Administration Console. (See [Administration Help](#).)*

The CRL distribution point describes where you can download the CRL that corresponds to a particular CER or PFX file.

The following file types are supported:

Certificates: DER-encoded X509v3 and base64-encoded certificate (.cer) files. Certificates that verify the trust.xml file can be either DER-encoded or base64-encoded.

Credentials: RSA and DSA credentials up to 4096 bits in standard PKCS12 format (.pfx and .p12 files).

CRLs: Base64-encoded and DER-encoded CRLfiles.

Maintaining the security of private keys (credentials) is critical to ensuring the stability of sensitive information. A physical storage device, often called a *Hardware Security Module* (HSM), typically provides the maximum level of security for private keys. If you do not use a physical device, it is important to store highly sensitive private keys and certificates in encrypted files in a safe place.

Digital Signatures supports the industry-standard PKCS #11 interface to communicate with HSMs. An HSM vendor can provide the resources and tools that you need to install and configure an HSM storage system.

Chapter 5: Creating the LiveCycle Database

5.1 Database configuration requirements

This section describes how to set up the database for use with LiveCycle. This section describes special tables, indexes, and other properties that are required in the LiveCycle database that are not configured by Configuration Manager.

The database will contain these elements:

- LiveCycle services
- LiveCycle run-time configurations
- LiveCycle process data
- Customer process definitions and templates
- Application server managed data

Before you create the database, you must ensure that you read the pre-installation requirements and have the required software installed.

If this is the first installation of LiveCycle, you must create an empty database. All the tables required to support LiveCycle are created when you initialize the LiveCycle database using Configuration Manager. (See the *Installing and Deploying LiveCycle* document for your application server.)

5.2 Minimum database user permissions

Database	Initialization permissions	Runtime permissions
Oracle	CREATE SESSION CREATE CLUSTER CREATE TABLE CREATE VIEW CREATE SEQUENCE UNLIMITED TABLE SPACE	CREATE SESSION UNLIMITED TABLE SPACE (only needed if you do not configure user quotas) CREATE TABLE
MySQL	SELECT INSERT UPDATE DELETE CREATE DROP REFERENCES INDEX ALTER CREATE_TEMP_TABLE LOCK_TABLES	SELECT INSERT UPDATE DELETE
SQL Server - DB level	Create Table Create View Connect	Connect
SQL Server - Schema level	Alter Insert References Select Update Delete	Insert Select Update Delete
DB2	See "5.5.1 DB2 user account" on page 43 for a complete description.	See "5.5.1 DB2 user account" on page 43 for a complete description.

5.3 Creating an Oracle database

If you prefer not to use the default database that was created when you installed Oracle 11g, create a new database by using the Database Configuration Assistant tool. LiveCycle supports Oracle 11g on WebLogic, WebSphere, and JBoss application servers.

Note: You can use the Transaction Processing or General Purpose templates while configuring an Oracle database instance for LiveCycle. If you wish to use the Custom Database template for configuring a database instance, the minimum set of database components you must include are **Oracle JVM** and **Enterprise Manager Repository**.

Do the following when you create your Oracle database:

- Create user quotas to allow the database to grow to accommodate persistent data from LiveCycle applications.
- Enable support for UTF-8 encoding.
- Set Database Character Set to Unicode (AL32UTF8), and the National Character Set is AL16UTF16 (Unicode UTF-16 universal character set).
- Set NLS_LENGTH_SEMANTICS to BYTE (if required). The database initialization fails if you set any other value.
- You must install Oracle using Transaction Processing and set the connection mode for the server to Dedicated Processing.

5.3.1 User account and rights

Create a new user account on the database and assign it the following system privileges:

- CREATE SEQUENCE
- CREATE VIEW
- UNLIMITED TABLESPACE
- CREATE TABLE
- CREATE CLUSTER
- CREATE SESSION

Important: (WebLogic only) When using WebLogic and an Oracle database, ensure that the database user name does not have a number as the first character, nor a hyphen (-) within the user name. Otherwise, the database does not bootstrap properly.

Note: For deployments on AIX, Linux, or Solaris, the user name must not exceed 8 characters; on Windows, it must not exceed 12 characters.

You need the following information when you configure the data source on the application server:

- SID (Service ID)
- User name and password of the Oracle user account
- Host name or IP address of the database server
- Oracle LISTENER port number (default is 1521)

For information about using Oracle 11g, see the appropriate user documentation.

5.4 Creating a SQL Server database

You can create a SQL Server database that LiveCycle will use to store run-time and configuration data. For information about creating a SQL Server database, refer to the SQL Server documentation. LiveCycle supports SQL Server 2005 SP2 and 2008 R2 (English and Japanese) on JBoss, WebLogic, and WebSphere.

Create a SQL Server database, and create a user account and assign it DB_OWNER privileges for use when configuring the data source on the application server. For information about creating the database and user, see the SQL Server documentation.

You need the following information when you configure the data source on the application server:

- Database name
- User name and password of the SQL Server user account
- Host name or IP address of database server
- SQL Server port number

5.4.1 Set up SQL Server for LiveCycle

Before you create the LiveCycle database, optimize SQL Server by changing these settings.

5.4.1.1 Increase memory

The default SQL Server settings do not aggressively allocate memory. This situation significantly affects performance on most deployments of a SQL Server database.

Note: This section is recommended but optional.

- 1 Using Microsoft SQL Server Management Studio, connect to the database server where you will host the LiveCycle database.
- 2 Right-click the database server connection and select **Properties**.
- 3 Select the **Memory** page and enter a size in the **Minimum Server Memory (in MB)** box that is equal to the size of the free memory on the server.
- 4 Restart the SQL Server database.

5.4.1.2 Set the processor priority

On dedicated database servers, which are recommended for production installations of LiveCycle, the SQL Server process is configured so that it does not consume too much of the system CPU resources.

Note: This section is recommended but optional.

- 1 Using Microsoft SQL Server Management Studio, connect to the database server where you will host the LiveCycle database.
- 2 Right-click the database server connection and select **Properties**.
- 3 Select the **Processors** page and select **Boost SQL Server Priority**.
- 4 Restart the SQL Server database.

5.4.1.3 Increase the recovery interval

This setting specifies the amount of time the deployment waits for recovery after a crash. The SQL Server default setting is one minute. Increasing this setting to a larger value improves performance because it causes the server to write changes from the database log to the database files less frequently. This setting does not compromise the transactional behavior; however, it does affect the size of the log file that is replayed on startup.

Note: This section is recommended but optional.

- 1 Using Microsoft SQL Server Management Studio, connect to the database server where you will host the LiveCycle database.
- 2 Right-click the database connection and select **Properties**.
- 3 Select the **Database Settings** page and type 5 in the **Recovery Interval (Minutes)** box.

4 Restart the SQL Server database.

5.4.1.4 Integrated security

Note: This is an optional configuration.

If you are using SQL Server integrated security, you can set your SQL Server database to Mixed Mode or Windows Authentication Mode. However, if you are using Windows Authentication Mode, you must configure integrated security on Windows to establish a trusted connection with SQL Server.

- For JBoss, see Configure Integrated Security on Windows in [Preparing to Install LiveCycle on Single Server](#).
- For WebLogic, see Configuring SQL Server database connectivity in [Installing and Deploying LiveCycle for WebLogic](#).
- For WebSphere, see Configuring SQL Server database connectivity in [Installing and Deploying LiveCycle for WebSphere](#).

5.4.2 Sizing your SQL Server database

The default database sizes that SQL Server provides are too small for LiveCycle. Even if the database is set to auto-grow, unintended effects can occur, such as reduced performance when the database grows or the growth begins to fragment the disk. It is best to preallocate the database size at creation to reflect your deployment requirements:

Medium size deployments: Environments where the LDAP directory has approximately 100,000 users and 10,000 groups. Set Database Data Initial Size to 1 GB, and set autogrowth to 250 MB.

Large size deployments: Environments where the LDAP directory has approximately 350,000 users and more than 10,000 groups. Set Database Data Initial Size to 2 GB, and set autogrowth to 1 GB.

Note: Database growth is always restricted to a certain size. Administrators should monitor the resource usage of the LiveCycle database to ensure that it does not lose its restricted space or the space available on the disks where the database resides.

5.4.3 Creating the LiveCycle database user, schema, and login

You are now ready to create the LiveCycle database user, schema, and login.

Important: Ensure that you use the `SQL_Latin1_General_CP1_CI_AS` collation (or the `Japanese_CI_AS` collation if your database will run in a Japanese environment) when you create the database instance for LiveCycle. Any other collation may cause your database initialization to fail. The collation for your LiveCycle database instance can be different from the collation used when creating the SQL Server database.

1 Using Microsoft SQL Server Management Studio, click **Server**, and then right-click **Database** and select **New Database**.

2 Enter the database name of your choice.

Note: The database name is very important, and the name chosen must be consistently used in the following procedures where a reference to `database_name` exists.

3 (SQL Server 2005 Only) Ensure that **Use Full-Text Indexing** is deselected.

4 In the **Database Data Initial Size MB** box, enter the appropriate value:

- For small development or small production systems, specify 200 MB.
- For larger systems, see “5.4.2 Sizing your SQL Server database” on page 41.

5 In the **Database DataAutogrowth** box, enter 50%.

- 6 In the **Database LogInitial Size** box, enter the appropriate value:
 - For small development or small production systems, specify 20 MB.
 - For larger systems, see “5.4.2 Sizing your SQL Server database” on page 41.
- 7 In the **Database LogAutogrowth** box, enter 50%.
- 8 Click **OK** to create the database.

5.4.3.1 Create the LiveCycle user in SQL Server

In the following procedure, *[database_name]* represents the name you specified when you created your database, and *[database_username]* represents the name you must specify for the new user.

- 1 Using Microsoft SQL Server Management Studio, connect to the database server where you created the LiveCycle database.
- 2 Click **Server** > **Security**, and then right-click **Logins** and select **New Login**.
- 3 Enter the login name *[database_username]*, and then select **SQL Server Authentication** and type a new password.
- 4 Ensure that **Enforce Password Expiration, User must change password on next login** is also deselected.
- 5 Leave the default database as **Master**, and click **OK**.
- 6 Click **Server** > **Databases** > *[database_name]* > **Security**, and then right-click **Schemas** and select **New Schema**.
- 7 In the **Schema Name** box, type *[database_username]*, and click **OK**.
- 8 Click **Server** > **Databases** > *[database_name]* > **Security**, and then right-click **Users** and select **New User**.
- 9 In the New User dialog box, type the login name and user name *[database_username]*.
- 10 Set the default schema to *[database_username]* and click **OK**.
Note: The schema name should be the same as the [database_username].
- 11 (SQL Server 2008 R2) In the **Schemas Owned By This User** area, select the schema *[database_username]*.
- 12 Click **Server** > **Databases** > *[database_name]* > **Security**, right-click the *[database_username]* schema, and select **Properties** > **Permissions**.
 - a Click **Search** in Users or Roles and type *[database_username]* and click **OK**.
 - b In the **Explicit** tab, grant the following permissions:
 - Alter
 - Insert
 - Reference
 - Select
 - Update
 - Delete
 - c Click **OK**.
- 13 Right-click **Server** > **Databases** > *[database_name]*, right-click the *[database_username]* schema, and select **Properties** > **Permissions**.
 - a In the **Explicit** tab, grant Create Table, Create View, and Connect permissions.
 - b Click **OK**.

5.4.3.2 Associate the LiveCycle user with the database

After you create the LiveCycle user, associate it with the LiveCycle database.

- 1 Click **Security > Logins**, and then right-click `[database_username]` and select **Properties**.
- 2 In Login Properties, on the General page, set the user's default database to `[database_name]`.
- 3 Select the **User Mapping** page and, in the Users Mapped To This Login section, verify that `[database_name]` is selected, **User** is set to `[database_username]`, and **Default Schema** is set to `[database_username]`.
- 4 Ensure that `[database_name]` is selected in the Users Mapped To This Login table, and ensure that **public** is selected in the **Database Role Membership For**`[database_name]` table and then click **OK**.

5.4.3.3 Set the isolation level for the LiveCycle database

LiveCycle requires a specific isolation level to manage deadlocking when long-running transactions occur at the same time as numerous shorter reads.

Important: You must set the isolation level for MS SQL Server to avoid deadlocking issues.

- 1 Click **Databases**, and then right-click `[database_name]` and select **New Query**.
Note: `[database_name]` represents the name you specified when you created your database.
- 2 In the Query panel, type the following text:

```
ALTER DATABASE [database_name]
SET READ_COMMITTED_SNAPSHOT ON
GO
```
- 3 Click **Execute**. A response is displayed in the messages panel.

5.5 Creating a DB2 database

Create a DB2 database by running one of the scripts provided in this section. The script is tuned for a system that will use 1 GB of memory for the database. If your system has less memory dedicated for the database, see the appropriate DB2 documentation for details about configuring your system settings. LiveCycle supports DB2 with WebSphere and WebLogic.

You need the following information when you configure the data source on the application server:

- Database name
- User name and password of the DB2 user account
- Host name or IP address of the database server
- DB2 port number

5.5.1 DB2 user account

LiveCycle requires a dedicated system database account. If the account is the schema owner no other privileges are required. The schema owner can alter, create and drop any object in the schema. If the account is not the schema owner then the ALTERIN, CREATEIN AND DROPIN schema privileges are required. The user will also have to belong to the DB2USERS group.

The DB2 user account needs to have schema owner rights to create tables both at product installation time and system runtime. The application server may add further requirements to the database user and schema due to its internal use of the database.

For information about creating a user account, see the DB2 documentation.

Note: For deployments on AIX, Linux, or Solaris, the user name must not exceed 8 characters; on Windows, it must not exceed 12 characters.

5.5.2 Create the DB2 database

Note: The following text contains formatting characters for line breaks. When you copy this text to a location outside this document, remove the formatting characters.

```
-- Create a database, using a custom USERSPACE1 definition
-- Configures the database for typical LC use. For large number of connections, add:
--     num_remote_apps
-- to the AUTOCONFIGURE section
-- To use: Globally replace (case sensitive) DB_NAME with the name of the database to be
created.
-- Modify the FILE locations based on system configuration
-- Note: The size of the file has to be specified in pages (4K pages in this case) when used
as part of the create database command
create database DB_NAME using codeset utf-8 territory default
USER TABLESPACE MANAGED BY DATABASE USING (FILE'C:\Db2\DB_NAME\DATA_4K_1.db2' 8192) AUTORESIZE
YES
AUTOCONFIGURE using workload_type mixed isolation CS admin_priority performance apply DB and
DBM;
-- Create required buffer pools. Let the DB manage the sizes based on load.
-- Creates 8K pool for LC
-- Let DB2 manage the sizes based on usage patterns
connect to DB_NAME;
CREATE BUFFERPOOL DB_NAME_BP8K IMMEDIATE SIZE AUTOMATIC PAGESIZE 8192;
ALTER BUFFERPOOL IBMDEFAULTBP IMMEDIATE SIZE AUTOMATIC;
connect reset;
-- Create required tablespaces for tables requiring 8K and 32K pools.
connect to DB_NAME;
-- 8K tablespace used by a number of tables.
-- Note that the majority of data is stored in the 4K tablespace (created as part of the databse
create)
CREATE TEMPORARY TABLESPACE DB_NAME_TEMP_8K IN DATABASE PARTITION GROUP IBMTEMPGROUP PAGESIZE
8192 MANAGED BY SYSTEM
USING ('C:\Db2\DB_NAME\TEMP_8K') PREFETCHSIZE AUTOMATIC BUFFERPOOL DB_NAME_BP8K;
-- Start off large enough to allow a reasonable test to run without needing to extend the space
-- Should be modified to fit expected usage patterns
CREATE LARGE TABLESPACE DB_NAME_DATA_8K IN DATABASE PARTITION GROUP IBMDEFAULTGROUP PAGESIZE
8192 MANAGED BY
DATABASE USING (FILE'C:\Db2\DB_NAME\DATA_8K.db2'1G)
PREFETCHSIZE AUTOMATIC BUFFERPOOL DB_NAME_BP8K AUTORESIZE YES INCREASESIZE 1 G;
commit work;
connect reset;
deactivate database DB_NAME;
activate database DB_NAME;
```

1 Make the following changes to the script:

- Replace the instances of *dbname* and *DBNAME* with the name you want for the LiveCycle database.
- If you are creating multiple database instances on the same host, create an uniquely named buffer pool (for example, *BP8K_1*) for each database instance, using a buffer pool size such that all will not exceed 10% of total physical memory. For example, on a machine with 1GB of physical memory, the total buffer pool size should not exceed 100MB (“SIZE 100000”)
- Replace *DB2_root* with the path to the root directory where DB2 is installed.
- Ensure that no commands include line breaks and each command is terminated by a semicolon (;).
- Change 9000 in the following line based on your database size:

```
(FILE'DB2_root\DBNAME_DATA'9000)
```

This number specifies the minimum number of pages required to initialize the database. You can also change this number by using the DB2 administration tools after you initialize the database.

2 Save the text file in a location that DB2 Command Line Processor can access.

3 Open a DB2 command prompt and type the following command to run the script:

```
db2 -tf <path_to_script_file>/<script_file_name>
```

5.5.3 Configure DB2 for concurrent usage

1 Open the DB2 Control Center:

- (Windows) Select **Start > Programs > IBM DB2 > General Administration Tools > Control Center**.
- (AIX, Linux, and Solaris) From a command prompt, enter the command `db2jcc`.

2 In the DB2 Control Center object tree, click **All Databases**.

3 Right-click the database created for LiveCycle products and click **Configuration Advisor**.

4 Follow the steps in the Configuration Advisor wizard and set the following properties:

DB2 Property	Required Value
Workload type	Mixed
Average number of SQL transactions per unit of work	Less than 10
Database Administration Priority	Faster transaction performance
Populated Database	Set according to the current state of the database instance. The instance is not populated if it has not yet been initialized by the Configuration Manager.
Isolation Level	Read Committed

5 Click **Finish**.

5.5.4 Add a user to the new database

1 Log in to the IBM Control Center.

2 Click *[database_name]* > **User and Group Objects > DB Users**.

3 Select **Add User** and select the user.

- 4 Under Authorities, select **Connect to database**, **Create tables**, and **Database administration authority**, and then select **Apply**.
- 5 Click **OK**.

5.6 Creating a MySQL database

Note: The information contained in this document is meant for users who are installing MySQL manually and not for a turnkey installation. See [Installing and Deploying LiveCycle for JBoss using Turnkey](#).

Use the MySQL tools to create a MySQL 5 database for use with LiveCycle and a MySQL user account that the application server can use to connect to the database. You also must modify the MySQL database server configuration. For information about creating the database and user account, see the MySQL documentation. LiveCycle supports MySQL 5 with JBoss.

Note: MySQL does not support the use of special characters or spaces in the user name or password. Ensure that your user name and password adhere to this restriction. Also, to avoid Scheduler errors, do not use a dash (-) in the MySQL database name.

Note: On AIX, Linux, or Solaris, you must set the `lower_case_table_names` system variable to 1 to ensure that table names are case-insensitive. On Windows systems, this parameter does not have any effect. For more information about setting table names for case-sensitivity, see <http://dev.mysql.com/doc/refman/5.1/en/identifier-case-sensitivity.html>.

Note: To set up a database schema and new users on MySQL by using a graphical user interface (GUI), you must install MySQL Administrator tool.

When installing the MySQL database, you must specify UTF-8 character encoding in the Variables section of the MYSQL UI.

You need the following information when you configure the data source on JBoss:

- User name and password of the MySQL user account
- Host name or IP address of database server
- MySQL port number

5.6.1 Limitation on document usage in processes for MySQL databases

The MySQL database limits the amount of data that can be stored in a table row. Processes that involve multiple instances of documents can require more storage space than the MySQL row size limitation. Because LiveCycle stores data for a process instance in a single table row, processes that involve large amounts of data can exceed the MySQL limitation on row size, causing errors to occur.

If you run a process that involves many instances of documents and errors occur in the database, you can redesign the process so that the operations that handle documents execute in several subprocesses. For example, you can redesign your process so that its operations are executed in two processes. The first process includes the first half of the operations of the original process, and the second process includes the last half of the operations. The first process must include the invoke operation of the second process so that the second process runs as a subprocess.

5.6.2 MySQL user account

The MySQL user account that you create must have these privileges to access the tables in the LiveCycle database:

- SELECT

- INSERT
- UPDATE
- DELETE
- CREATE
- DROP
- REFERENCES
- INDEX
- ALTER
- CREATE_TMP_TABLE
- LOCK_TABLES

5.6.3 MySQL initial server configuration requirements

The following configuration setting is required in the `mysqld` section of the `my.ini` file (Windows) or `/etc/my.cnf` file (AIX, Linux, or Solaris):

- `max_allowed_packet=25M`

Note: Failure to configure this setting will result in “packet too large” errors that will cause module deployment to fail.

These configuration settings are recommended in the `mysqld` section of the `my.ini` file (Windows) or `/etc/my.cnf` file (AIX, Linux, or Solaris):

- `key_buffer_size=64M`
- `sort_buffer_size=1M`
- `default-storage-engine= INNODB`
- `innodb_buffer_pool_size=100M`
- `innodb_additional_mem_pool_size=5M`
- `innodb_log_file_size=170M`
- `innodb_log_buffer_size=8M`
- `innodb_flush_log_at_trx_commit=2`
- `query_cache_type=2`
- `default-character-set=utf8`
- `read_buffer_size=1M`
- `read_rnd_buffer_size=1M`