



Escenic Content Engine  
**Installation Guide**

5.2.7.2







Copyright © 2008-2011 Vizrt. All rights reserved.

No part of this software, documentation or publication may be reproduced, transcribed, stored in a retrieval system, translated into any language, computer language, or transmitted in any form or by any means, electronically, mechanically, magnetically, optically, chemically, photocopied, manually, or otherwise, without prior written permission from Vizrt.

Vizrt specifically retains title to all Vizrt software. This software is supplied under a license agreement and may only be installed, used or copied in accordance to that agreement.

### **Disclaimer**

Vizrt provides this publication “as is” without warranty of any kind, either expressed or implied.

This publication may contain technical inaccuracies or typographical errors. While every precaution has been taken in the preparation of this document to ensure that it contains accurate and up-to-date information, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained in this document.

Vizrt’s policy is one of continual development, so the content of this document is periodically subject to be modified without notice. These changes will be incorporated in new editions of the publication. Vizrt may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

Vizrt may have patents or pending patent applications covering subject matters in this document. The furnishing of this document does not give you any license to these patents.

### **Technical Support**

For technical support and the latest news of upgrades, documentation, and related products, visit the Vizrt web site at [www.vizrt.com](http://www.vizrt.com).

### **Last Updated**

21.12.2011





# Table of Contents

<b>1 Introduction</b> .....	7
<b>1.1 Installation Components</b> .....	7
<b>1.1.1 Escenic Components</b> .....	7
<b>1.1.2 Third-Party Components</b> .....	8
<b>1.2 Installation Types</b> .....	9
<b>1.3 Conventions Used in This Manual</b> .....	10
<b>1.3.1 Identifying Host Machines</b> .....	10
<b>1.3.2 Shell Commands</b> .....	11
<b>1.3.3 Standard File Hierarchy</b> .....	11
<b>2 Product Requirements</b> .....	13
<b>2.1 Required Supporting Software</b> .....	13
<b>2.2 Other Requirements</b> .....	13
<b>2.2.1 Shared Filesystem</b> .....	13
<b>2.2.2 SSH/FTP Server</b> .....	14
<b>2.2.3 Apache Ant</b> .....	14
<b>3 Installation Procedure</b> .....	15
<b>3.1 Install Java SE Development Kit (JDK)</b> .....	15
<b>3.2 Install Ant</b> .....	16
<b>3.3 Install Various Utilities</b> .....	16
<b>3.4 Create Escenic Users</b> .....	16
<b>3.5 Create Shared File System</b> .....	17
<b>3.6 Download Content Engine</b> .....	18
<b>3.7 Install LDAP Server</b> .....	19
<b>3.8 Install Database</b> .....	22
<b>3.9 Install Application Server</b> .....	23
<b>3.10 Create Configuration Folder</b> .....	25
<b>3.11 Unpack Content Engine Components</b> .....	26
<b>3.12 Install a Daemon Script</b> .....	27
<b>3.13 Logging</b> .....	28
<b>3.14 Copy Solr Configuration</b> .....	28
<b>3.15 Initialize the Assembly Tool</b> .....	29
<b>3.16 Initialize the Bootstrap Layer</b> .....	29
<b>3.17 Create the Common Configuration Layer</b> .....	30
<b>3.18 Create Host Configuration Layers</b> .....	33



<a href="#">3.19 Install the ece Script</a>	34
<a href="#">3.20 Assemble and Deploy</a>	36
<a href="#">3.21 Verify the Installation</a>	37
<a href="#">3.22 Create a Publication</a>	38
<a href="#">3.23 Test Web Studio</a>	42
<a href="#">3.24 Test Content Studio</a>	43
<a href="#">4 Recommended Modifications</a>	45
<a href="#">4.1 Configure SSL Support</a>	45
<a href="#">4.1.1 Using Self-Signed Certificates</a>	46
<a href="#">4.2 Search Engine Deployment and Configuration</a>	46
<a href="#">4.3 Distributed Memory Cache</a>	48
<a href="#">4.3.1 Install memcached</a>	48
<a href="#">4.3.2 Install the memcached Java Libraries</a>	49
<a href="#">4.3.3 Configure memcached</a>	49
<a href="#">4.3.4 Configure PresentationArticleCache Instances</a>	50
<a href="#">5 Installing on Windows</a>	51
<a href="#">5.1 Windows Software Components</a>	51
<a href="#">5.2 Windows Installation Hints</a>	52
<a href="#">5.2.1 General Advice</a>	52
<a href="#">5.2.2 LDAP on Windows</a>	52
<a href="#">5.2.3 MySQL on Windows</a>	53
<a href="#">5.2.4 Tomcat on Windows</a>	53
<a href="#">5.2.5 JDBC Connector on Windows</a>	54
<a href="#">5.3 Assembly and Deployment on Windows</a>	54



---

# 1 Introduction

This **Installation Guide** is intended to be read by the person responsible for installing the Escenic Content Engine on one or more servers. It should be read together with the **Escenic Content Engine Server Administration Guide**, which contains descriptions of the periodic administration tasks a system administrator needs to carry out once the Content Engine is installed and in operation, since there is considerable overlap between installation and maintenance of the Content Engine.

Both this manual and the **Escenic Content Engine Server Administration Guide** make the following assumptions about the Escenic installation and you, the reader:

- The Content Engine is to be installed on one or more host computers running an operating system listed in [section 1.1.2](#).
- The supporting software components (database, web server, application server and so on) listed in [section 1.1.2](#) are either already installed on these computers, or available for installation.
- You are a suitably qualified system administrator with a working knowledge of both the operating system on which the Content Engine is to be installed and of the components in the supporting software stack.

---

## 1.1 Installation Components

A working Escenic Content Engine installation depends on many components. Some of these components are delivered by Vizrt as part of the product, others are third-party products on which the Content Engine depends.

### 1.1.1 Escenic Components

The components supplied by Vizrt are:

**Escenic Content Engine**

The Content Engine itself, a content management system.

**escenic-admin**

A server administration web application.

**Escenic Web Studio**

A publication administration web application.

**Escenic Content Studio**

A Java desktop application for creating, editing and managing publication content. Although this is a desktop application and runs on end-user's PCs, it is served to users from the Content Engine via Java Webstart, and is therefore part of the Content Engine installation.

**Publications**

The installation includes a few small demo publications for test, demonstration and instructional purposes.

**Assembly tool**

The assembly tool is an ant build file that automates the application assembly and deployment process.

## 1.1.2 Third-Party Components

The Content Engine depends on the following third-party components:

**Java Development Kit (JDK)**

This provides the Java runtime and development environment required by J2EE application servers.

**Database**

Used to store publication content, structure and so on.

**LDAP server**

An LDAP (Lightweight Directory Access Protocol) server is used to store user information (user names, passwords, permissions and so on).

**J2EE Java application server**

The Content Engine is a JEE web application, and therefore runs inside a JEE application server. `escenic-admin`, Web Studio and all Escenic publications are also JEE applications and therefore also run within a JEE application server.

**Web server (optional)**

J2EE applications have built-in web servers that can be used for test and development purposes. For large-scale production installations, however, a separate web server may be advisable.

**JDBC (Java DataBase Connectivity) driver**

This is required to connect the Content Engine to the database.

**Apache Ant**

This is a Java-based build tool that is used by the Content Engine assembly tool.

**Apache Solr**

This is a Java-based search engine. It is used to provide search functionality both in web publications and in Content Studio. Unlike the other third-party components listed here, Solr is actually installed with the Content Engine.

**Caching Server (optional)**

A caching server such as Squid or Varnish can greatly improve the performance of heavily-loaded sites by caching frequently requested pages and sending the cached copies on request unless the original pages have been updated.



**Load Balancing Server (optional)**

If your site has multiple web hosts, then you may also want to install a load-balancing server to distribute requests evenly between them.

**memcached (optional)**

If your site has multiple web hosts, then you can improve caching efficiency by using this open source distributed cache manager to provide a shared cache for all your hosts.

In principle, the Content Engine can work with many different third-party components. In [section 2.1](#), however, you will find the list of officially supported software. The current version of the Content Engine has been tested with these components and is known to work satisfactorily.

---

## 1.2 Installation Types

The Content Engine can be installed in a number of different ways to satisfy different needs. Common installation types include:

**Single computer**

Everything (database, LDAP, application server, Content Engine) is installed on a single computer. This is usually the case for demonstration and developer installations.

**Multiple application servers, single database/LDAP server**

The application server and Content Engine is installed on multiple computers, but all share a single database and LDAP server. This kind of installation is common in large development environments.

**Single web server, single application server and Content Engine, single database/LDAP server**

This is a common production installation structure for small sites.

**Multiple web servers, multiple application servers/Content Engines, single LDAP, database rack**

This is a common production installation structure for large sites, and is easily scalable.

Other configurations are, of course, possible.

In large organizations the Content Engine often needs to be installed in several different configurations in order to meet the requirements of different environments:

**Development**

In a development environment there are typically several updates of several files several times a day. Depending on the number of developers working on a system and the potential for conflicting changes, developers may either:

- Share a single development installation, or

- Each use a personal Content Engine installation but share a single database and LDAP server.

**Test**

A test environment is a complete Content Engine installation on which 'current stable versions' can be tested. Code updates in this environment are typically carried out at longer intervals and in 'batches' rather than on the file level. In some organizations, the test environment may be merged with the staging environment.

**Staging**

A staging environment should be as similar to the production environment as practically possible: it should have the same software, the same configuration, and if possible the same hardware, firewall and network setup. This is the last stage before actual production, and is intended to catch "works on the test server but not in production" problems.

**Production**

A production environment is heavily optimized for high workloads and security.

---

## 1.3 Conventions Used in This Manual

This section describes various conventions used in this manual.

### 1.3.1 Identifying Host Machines

The Content Engine and the software it depends on may need to be installed on one or several host machines depending on the type of installation required (see [section 1.2](#)). In order to unambiguously identify the machines on which various installation actions must be carried out, the following special host names are used throughout this manual:

**assembly-host**

The machine used to assemble the various Content Engine components into an enterprise archive or .EAR file.

**database-host**

The machine used to host the database server.

**ldap-host**

The machine used to host the LDAP server.

**engine-host**

The machine(s) used to host application servers and Content Engine instances.

**editorial-host**

**engine-host**(s) that are used solely for (internal) editorial purposes. All your Content Studio clients communicate with these hosts.

**presentation-host**

**engine-host**(s) that are used solely for (public) presentation purposes. Readers' browsers communicate with these hosts.

**web-host**

The machine(s) used to host web servers (if needed).

**nfs-host**

The machine used to host the NFS server (if needed).

These host names always appear in a **bold typeface**. If you are installing everything on one host you can, of course, ignore them: you will be doing everything on the same machine. If you are creating a larger multi-host installation, then they should help ensure that you do things in the right places.

### 1.3.2 Shell Commands

All shell command examples given in the manual are tested on Debian Linux servers: they may need minor modifications to be used on other Linux or UNIX platforms, and it is assumed that you are able to make the necessary "conversions" to your own platform. Some of the commands should be executed as the owner of the Content Engine installation. This is signalled by use of the `$` command prompt. For example:

```
$ ls
```

Other commands must be executed as `root`. This is signalled by the use of the `#` command prompt:

```
# /etc/init.d/slaped restart
```

Many of the code and command line examples in the manual contain placeholders: words printed in an *italic* typeface, which you are expected to replace with an appropriate value. For example:

```
$ cat filename
```

The accompanying text usually specifies what kind of value any placeholder(s) represent: *filename* in this case is the name of a file you want to see the contents of.

### 1.3.3 Standard File Hierarchy

All file paths and URLs shown in the manual are based on the following standard folder structure:

Standard location	Component
<code>/opt/escenic</code>	Escenic
<code>/opt/escenic/engine</code>	Escenic Content Engine
<code>/opt/escenic/assemblytool</code>	Escenic assembly tool
<code>/etc/escenic</code>	Escenic configuration
<code>/etc/escenic/engine</code>	Escenic Content Engine configuration

<b>Standard location</b>	<b>Component</b>
/opt/java/jdk	Java
/opt/java/ant	Ant
/opt/tomcat	Tomcat

If your system is organized differently, then adjust the paths you use accordingly.



---

## 2 Product Requirements

The Escenic Content Engine is written in Java, runs on an application server, uses a database to store content, and an LDAP server to store user data. One of each of these supporting products must be installed before the Content Engine can be installed. The specific supporting products with which the Content Engine is known to work are listed in [section 2.1](#). Note, however, that these products have their own hardware and software requirements. Please make sure that your environment meets the requirements stipulated by the relevant product suppliers.

Some additional, more general requirements and recommendations are listed in [section 2.2](#).

---

### 2.1 Required Supporting Software

The Content Engine requires a combination of the following software components:

#### **Operating System**

- Ubuntu Server (Long Term Support version recommended)
- Solaris 10
- Windows Server 2008 R2 (Note however, that most of this guide describes installation on a Linux/Unix platform. For additional guidance regarding installation on Windows, see [chapter 5](#).)

#### **Application Server**

- Apache Tomcat 6.0

#### **Database Server**

- MySQL 5.0.x
- Oracle Database 10.1 or later

#### **LDAP Server**

Open LDAP 2.0 or later (using LDAP v3 protocol)

#### **Java SE Development Kit**

Oracle Java SE Development Kit 6

#### **Web Server**

Apache Web Server Version 2.2.12

---

### 2.2 Other Requirements

#### 2.2.1 Shared Filesystem

In large installations, the web server may well run on a different host to the application server. However, the web server and application server need

access to a common file system, since the Content Engine writes and updates files that are served by the web server (multimedia files, for example). On Linux systems, NFS can be used to set up the required shared folders. For further information, see [section 3.5](#).

### 2.2.2 SSH/FTP Server

The Content Engine does not need an SSH or FTP Server. Setting up an SSH or FTP server is, however, highly recommended. An SSH/FTP Server can be used by template developers, among others, to upload new and updated files to the server. If it is not possible to install an SSH or FTP server on the **engine-host** (for security reasons, perhaps), then some other way of uploading files must be provided.

### 2.2.3 Apache Ant

Ant is used to package an Enterprise Archive (EAR file) containing the Content Engine and all the web applications required to run it.

Ant is only required on the **assembly-host**. It is not necessary to install ant on all of the hosts in a cluster.



---

## 3 Installation Procedure

This chapter contains step-by-step instructions for installing the Content Engine on a single host computer or on a cluster of several host computers. The instructions use the host names listed in [section 1.3.1](#) to indicate where you should carry out various steps. If some of your hosts are "multi-purpose" (if for example, your database server is installed on the same host as your editorial Content Engine installation), then carry out all appropriate steps on that machine (for example, all **database-host** steps and all **editorial-host** steps).

Some of the steps are not required for single-host deployment. These steps are clearly marked.

---

The Content Engine can be installed in many different ways. You are, however, recommended to follow the procedure described below as it will result in a standardized, easily-maintained installation structure.

---

---

### 3.1 Install Java SE Development Kit (JDK)

On your **assembly-host** and **engine-host(s)**, while logged in as **root**:

Download and install Java SE development kit (JDK), version 6. It may be possible to download it from one of your Linux distribution's software repositories. If so, this is usually the easiest method. On Debian-based Linux machines, for example, you can install it by entering:

```
# apt-get install sun-java6-jdk
```

---

If you are installing on Ubuntu 10.04 or higher, Oracle Java has been moved from the standard repository to the partner repository. Therefore you must configure your system to use the partner repository before you can enter the above command (if you have not already done so). To add the partner repository, open `/etc/apt/sources.list` for editing and add the following line:

```
deb http://archive.canonical.com/ubuntu lucid partner
```

---

Otherwise, you can download Oracle Java 6 from <http://www.oracle.com/technetwork/java/javase/downloads/index.html>.

Verify that Java is correctly installed by entering:

```
# java -version
```

This should result in output something like this:

```
java version "1.6.0_16"  
Java(TM) SE Runtime Environment (build 1.6.0_16-b01)
```

```
Java HotSpot(TM) Server VM (build 14.2-b01, mixed mode)
```

---

## 3.2 Install Ant

On your **assembly-host**, while logged in as **root**:

Download a binary distribution of **ant** from <http://ant.apache.org/bindownload.cgi>. Unpack the distribution to an appropriate place and follow the installation instructions given here:

<http://ant.apache.org/manual/install.html#installing>.

It may also be possible to download **ant** from one of your Linux distribution's software repositories. On Debian-based Linux machines, for example, you can install it by entering:

```
# apt-get install ant
```

Verify that **ant** is correctly installed by entering:

```
# ant -version
```

This should produce a response something like this:

```
Apache Ant version 1.7.0 compiled on April 29 2008
```

---

## 3.3 Install Various Utilities

Make sure the following utilities are available on all your hosts:

- **unzip**
- **telnet**

On Debian-based Linux machines, you can easily install these by entering the following commands:

```
# apt-get install unzip  
# apt-get install telnet
```

You may also find it useful to install an SSH server on all your hosts so that you can easily switch sessions and copy files between them. On Debian-based Linux machines, you can do this as follows:

```
# apt-get install openssh-server
```

---

## 3.4 Create Escenic Users

All the Content Engine components should run under the same user, and this user should have the same UID on all hosts. So create a user (preferably called **escenic**) on each host machine as follows:

```
# adduser scenic
```





and make sure they all have the same UID. If you are installing on clean machines, and **escenic** is the first user you create, then they will all have the same UID by default.

On your **assembly-host** and **engine-host(s)**, make sure that Java is in the path of the **escenic** users:

```
# su - escenic
$ echo 'export JAVA_HOME=/usr/lib/jvm/java-6-sun' >> .bashrc
$ echo 'export PATH=$JAVA_HOME/bin:$PATH' >> .bashrc
$ source ~/.bashrc
```

---

## 3.5 Create Shared File System

---

-----  
 If you are installing everything on a single host, you can skip this section.  
 -----

If you are installing the Content Engine on multiple hosts then it is usually a good idea to create some shared folders that can be accessed from all the hosts. If your **engine-host** is not the same machine as your **web-host**, then it is more or less a requirement, since the Content Engine creates and modifies multimedia files that need to be accessed by the web server. (If you cannot create shared folders for some reason, then you need to find some other means of synchronizing files between the hosts.)

Shared folders can also be useful for other reasons, such as access to downloads or shared configuration files, and these instructions assume that you will use shared folders for these purposes.

The following instructions show how to set up three shared folders: **ece-conf**, **multimedia** and **download**.

On your **nfs-host**, while logged in as **root**:

1. Download and install the NFS server software. For example, on a Debian-based Linux distribution:

```
# apt-get install nfs-kernel-server
```

2. Create the folders to be shared and set **escenic** as their owner:

```
# mkdir -p /exports/ece-conf
# mkdir -p /exports/multimedia
# mkdir -p /exports/download
# chown escenic:escenic /exports/ece-conf /exports/multimedia /exports/download
```

3. Open the NFS exports file **/etc/exports** and add an entry for each folder that you want to share:

```
/exports/ece-conf  subnet-specification(rw, sync)
/exports/multimedia  subnet-specification(rw, sync)
/exports/download  subnet-specification(rw, sync)
```

where *subnet-specification* identifies the IP address range within which the shared folders are to be accessible. If, for example, all your hosts have IP addresses in the **192.168.71.nnn** subnet, then you might enter the subnet specification **192.168.71.0/255.255.255.0**.

4. Restart all NFS-related services:

```
# /etc/init.d/portmap restart
# /etc/init.d/nfs-kernel-server restart
# /etc/init.d/nfs-common restart
```

5. Create symbolic links to the shared folders as follows:

```
# ln -s /exports/ece-conf /mnt/ece-conf
# ln -s /exports/multimedia /mnt/multimedia
# ln -s /exports/download /mnt/download
```

This ensures that the shared folders will be accessible via the same paths on your **nfs-host** as on all the other hosts.

6. Change the owner of the links to **escenic**:

```
# chown escenic:escenic /mnt/ece-conf
# chown escenic:escenic /mnt/multimedia
# chown escenic:escenic /mnt/download
```

On all your other hosts, while logged in as **root**:

1. Download and install the NFS client software. For example, on a Debian-based Linux distribution:

```
# apt-get install nfs-common
```

2. Create mount points for the shared folders:

```
# mkdir /mnt/ece-conf
# mkdir /mnt/multimedia
# mkdir /mnt/download
```

3. Open **/etc/fstab** and add an entry for each shared folder:

```
nfs-host:/exports/ece-conf /mnt/ece-conf nfs defaults 0 0
nfs-host:/exports/multimedia /mnt/multimedia nfs defaults 0 0
nfs-host:/exports/download /mnt/download nfs defaults 0 0
```

where *nfs-host* is the host name or IP address of your **nfs-host**.

4. Mount the shared folders:

```
# mount -a
```

---

## 3.6 Download Content Engine

You will need login credentials to be able to download Content Engine and other components that you need. If you have not already received a username and password for this purpose, please contact your Vizrt representative or Vizrt support ([support@escenic.com](mailto:support@escenic.com)).

On your **assembly-host**, while logged in as **escenic**:

Download the following items from the Vizrt Technet web site (<http://technet.escenic.com>):

- The Content Engine installation set **engine-5.2.7.2.zip**
- The Content Engine assembly tool **assemblytool-2.0-1.zip**



- The Content Engine documentation package `documentation-engine-5.2.7.2-pdf.zip` (if you do not already have a full, up-to-date documentation set)

If you are installing on more than one machine and have created the shared folders described in [section 3.5](#), then it is a good idea to download these packages to the `/mnt/download` folder you created. Otherwise, download them to some temporary location of your choice.

You can download the files by entering the following commands:

```
$ cd /mnt/download
$ wget http://user:password@technet.escenic.com/downloads/release/52/engine-5.2.7.2.zip
$ wget http://user:password@technet.escenic.com/downloads/assemblytool-at-version.zip
$ wget http://user:password@technet.escenic.com/downloads/release/52/documentation-engine-5.2.7.2-pdf.zip
```

where:

- *user* and *password* are the user name(s) and password(s) you have received from Vizrt
- *at-version* is the number of the latest version of the assembly tool (2.0-1 or later)

---

## 3.7 Install LDAP Server

The following instructions describe how to install and set up an OpenLDAP server for use by the Content Engine.

On your **ldap-host**, while logged in as **root**:

1. Download and install the OpenLDAP server. For example, on a Debian-based Linux distribution:

```
# apt-get install slapd ldap-utils
```

2. Import the LDAP schemas delivered with OpenLDAP:

```
# ldapadd -Y EXTERNAL -H ldapi:/// -f /etc/ldap/schema/cosine.ldif
# ldapadd -Y EXTERNAL -H ldapi:/// -f /etc/ldap/schema/inetorgperson.ldif
# ldapadd -Y EXTERNAL -H ldapi:/// -f /etc/ldap/schema/nis.ldif
# ldapadd -Y EXTERNAL -H ldapi:/// -f /etc/ldap/schema/misc.ldif
```

3. Create a new text file and open it. Enter the following LDIF (LDAP Interchange Format) code to define your LDAP database:

```
# Load dynamic backend modules
dn: cn=module{0},cn=config
objectClass: olcModuleList
cn: module
olcModulepath: /usr/lib/ldap
olcModuleload: {0}back_hdb

# Create the database
dn: olcDatabase={1}hdb,cn=config
objectClass: olcDatabaseConfig
objectClass: olcHdbConfig
olcDatabase: {1}hdb
olcDbDirectory: /var/lib/ldap
olcSuffix: dc=your-company,dc=com
olcRootDN: cn=admin,dc=your-company,dc=com
olcRootPW: your-password
```

```

olcDbConfig: {0)set_cachesize 0 2097152 0
olcDbConfig: {1)set_lik_max_objects 1500
olcDbConfig: {2)set_lik_max_locks 1500
olcDbConfig: {3)set_lik_max_lockers 1500
olcLastMod: TRUE
olcDbCheckpoint: 512 30
olcDbIndex: cn,sn,mail pres,eq,approx,sub
olcDbIndex: objectClass eq
olcDbIndex: member eq
  
```

You can just paste the above code into the text file, but make sure that you replace the placeholders *your-company* and *your-password* with appropriate values. Note that although the password you are entering is unencrypted, it will be overwritten by an encrypted password when you create the base DN in step 9 below.

4. Save the file you have created in a temporary location (`/tmp/db.ldif`, for example).
5. Enter the following command to create your database:

```
# ldapadd -Y EXTERNAL -H ldapi:/// -f /tmp/db.ldif
```

6. Generate a user password for your database by entering:

```
# slappasswd -h {MD5}
```

**slappasswd** will then prompt you to enter a password twice:

```
New password: your-password
Re-enter new password: your-password
```

When you have entered the password twice, `slappasswd` returns the password as an encrypted string, for example:

```
{MD5}a4GK0mvxI07STJQAIzIrYw==
```

Make a copy of this string.

7. Create a new text file and open it. Enter the following LDIF (LDAP Interchange Format) code to define your base DN:

```

dn: dc=your-company,dc=com
objectClass: dcObject
objectclass: organization
o: your-company.com
dc: your-company
description: Escenic root

dn: cn=admin,dc=your-company,dc=com
objectClass: simpleSecurityObject
objectClass: organizationalRole
cn: admin
userPassword: your-encrypted-password
description: LDAP administrator
  
```

You can just paste the above code into the text file, but make sure that you replace the placeholders *your-company* and *your-encrypted-password*. Replace *your-encrypted-password* with the encrypted password you created in step 6.

-----  
 The blank line in the above code is significant. If you copy it from a PDF version of this manual, you may find that when you paste it into



your file, the blank line has been removed. Make sure you replace it, or the file will not work.

8. Save the file you have created in a temporary location (`/tmp/base.ldif`, for example):
9. Enter the following command to create your base DN:

```
# ldapadd -Y EXTERNAL -H ldapi:/// -f /tmp/base.ldif
```

10. Create a new text file and open it. Enter the following LDIF (LDAP Interchange Format) code to modify the system access control list (ACL):

```
dn: olcDatabase={0}config,cn=config
changetype: modify
add: olcRootDN
olcRootDN: cn=admin,cn=config

dn: olcDatabase={0}config,cn=config
changetype: modify
add: olcRootPW
olcRootPW: your-encrypted-password
```

You can just paste the above code into the text file, but make sure that you replace the placeholder *your-encrypted-password* with the encrypted password you created in step 6.

11. Save the file you have created in a temporary location (`/tmp/system-ac1.ldif`, for example):
12. Enter the following command to create your base DN:

```
# ldapadd -Y EXTERNAL -H ldapi:/// -f /tmp/system-ac1.ldif
```

13. Create a new text file and open it. Enter the following LDIF (LDAP Interchange Format) code to create an Escenic administrator with read/write access to the LDAP database:

```
dn: olcDatabase={1}hdb,cn=config
add: olcAccess
olcAccess: to attrs=userPassword,shadowLastChange by dn="cn=admin,dc=your-company,dc=com" write
by anonymous auth by self write by * none
olcAccess: to dn.base="" by * read
olcAccess: to * by dn="cn=admin,dc=your-company,dc=com" write by * read
```

You can just paste the above code into the text file, but make sure that you replace the placeholder *your-company*.

14. Save the file you have created in a temporary location (`/tmp/rw.ldif`, for example):
15. Enter the following command to create your administrator:

```
# ldapmodify -Y EXTERNAL -H ldapi:/// -f /tmp/rw.ldif
```

16. Create a new text file and open it. Enter the following LDIF (LDAP Interchange Format) code to create an Escenic user with read-only access to the LDAP database:

```
dn: olcDatabase={1}hdb,cn=config
add: olcAccess
olcAccess: to * by dn="cn=read,dc=your-company,dc=com" read by self write by * auth
olcAccess: to dn.base="" by * read
olcAccess: to * by dn="cn=read,dc=your-company,dc=com" write by * read
```

You can just paste the above code into the text file, but make sure that you replace the placeholder *your-company*.

17. Save the file you have created in a temporary location (`/tmp/ro.ldif`, for example):
18. Enter the following command to create your user:

```
# ldapmodify -Y EXTERNAL -H ldapi:/// -f /tmp/ro.ldif
```

You should then verify that the LDAP server is running and accessible by trying to connect to the LDAP port (389) from each of your other hosts using `telnet`:

```
$ telnet ldap-host 389
```

where `ldap-host` is the host name or IP address of the **ldap-host**. If a connection is opened, then the LDAP server is running and accessible. Press `ctrl-z` to close the connection.

## 3.8 Install Database

The following instructions describe how to install and set up MySQL for use by the Content Engine.

On your **database-host**, while logged in as `root`:

1. Install the MySQL server and client packages. For example, on a Debian-based Linux distribution:

```
# apt-get install mysql-server mysql-client
```

During the installation you will be asked to specify a root password for the database.

2. Log in to the system and create a database for the Content Engine:

```
# mysql -p
mysql> create database db-name character set utf8 collate utf8_general_ci;
Query OK, 1 row affected (0.00 sec)

mysql> grant all on db-name.* to user@%' identified by 'password';
Query OK, 0 rows affected (0.00 sec)

mysql> grant all on db-name.* to user@'localhost' identified by 'password';
Query OK, 0 rows affected (0.00 sec)

mysql> exit
```

Replace `db-name`, `user` and `password` in the above commands with a database name, user name and password of your choice.

3. If you are installing everything on a single host or if you created a shared file system as described in [section 3.5](#), then you should have direct access to the Content Engine package you downloaded earlier (see [section 3.6](#)). If not, you will need to copy it to some temporary location on the **database-host**.
4. Change user to `escenic` and unpack the Content Engine package to a temporary location:

```
# su - scenic
$ cd /tmp
$ unzip /mnt/download/engine-5.2.7.2.zip
```

5. Still as the `escenic` user, run the Content Engine's database scripts:



```
$ cd /tmp/engine-5.2.7.2/database/mysql/
$ for e1 in tables.sql indexes.sql constants.sql constraints.sql; do
> mysql -u user -ppassword db-name < $e1
> done;
```

Replace *db-name*, *user* and *password* in the above commands with the names you chose in step 2.

6. On some platforms, external access to the MySQL server is disabled by default. To enable external access, you need to bind the `mysql` process to the **database-host**'s IP address. To do this, you need to open `/etc/mysql/my.cnf` for editing (as `root` again) and set the `bind-address` parameter:

```
bind-address = database-host-ip-address
```

You should then verify that the MySQL server is running and accessible by trying to connect to port 3306 from each of your other hosts using `telnet`:

```
$ telnet database-host 3306
```

where *database-host* is the host name or IP address of the **database-host**. If a connection is opened, then the database server is running and accessible.

## 3.9 Install Application Server

The following instructions describe how to install and set up the Apache Tomcat application server for use by the Content Engine.

On your **engine-host(s)**, while logged in as **escenic**:

1. Download the latest Tomcat package from <http://tomcat.apache.org/>. If you are installing on more than one machine and have created the shared folders described in [section 3.5](#), then it is a good idea to download these packages to the `/mnt/download` folder you created. Otherwise, download them to some temporary location of your choice.
2. Change user to `root` and unpack the Tomcat package to `/opt`:

```
$ su
# cd /opt
# tar -zxvf /mnt/download/apache-tomcat-version-number.tar.gz
```

where *version-number* is the version number of the package you have downloaded.

3. Create a symbolic link from `/opt/tomcat` to the Tomcat folder, so that it will be easier to upgrade to new versions when necessary, and change the owner of the Tomcat folder to **escenic**.

```
# ln -s /opt/apache-tomcat-version-number /opt/tomcat
# chown -R scenic:escenic /opt/apache-tomcat-version-number
```

4. Open Tomcat's configuration file file (`/opt/tomcat/conf/catalina.properties`) for editing and replace the following line:

```
common.loader=${catalina.home}/lib,${catalina.home}/lib/*.jar
```

with:

```
common.loader=${catalina.home}/lib,${catalina.home}/lib/*.jar,${catalina.home}/escenic/lib/*.jar
```

5. Download a JDBC driver for the database system you are using. For MySQL, the driver is called **Connector/J** and can be downloaded from <http://dev.mysql.com/downloads/connector/>.
6. Change user back to **escenic** and install the driver by copying it to the **/opt/tomcat/lib** folder. If you downloaded the MySQL driver package to **/mnt/download**, then you can do this as follows:

```
# su - escenic
$ cd /tmp
$ tar -zxvf /mnt/download/mysql-connector-java-version-number.tar.gz
$ cp mysql-connector-java-version-number/mysql-connector-java-version-number-bin.jar /opt/tomcat/lib/
```

7. Create the **/opt/tomcat/escenic/lib** folder you added to the **common.loader** path in step 4:

```
$ mkdir -p /opt/tomcat/escenic/lib/
```

8. Set up database pooling. For MySQL, you do this by opening **/opt/tomcat/conf/context.xml** for editing and inserting the following two resource definitions as children of the root **Context** element:

```
<Resource
  name="jdbc/ECE_READ_DS"
  auth="Container"
  type="javax.sql.DataSource"
  username="user"
  password="password"
  driverClassName="com.mysql.jdbc.Driver"
  maxActive="30"
  maxIdle="10"
  maxWait="5000"
  url="jdbc:mysql://database-host/db-name?
autoReconnect=true&useUnicode=true&characterEncoding=UTF-8"
/>

<Resource
  name="jdbc/ECE_UPDATE_DS"
  auth="Container"
  type="javax.sql.DataSource"
  username="user"
  password="password"
  driverClassName="com.mysql.jdbc.Driver"
  maxActive="10"
  maxIdle="5"
  maxWait="5000"
  url="jdbc:mysql://database-host/db-name?
autoReconnect=true&useUnicode=true&characterEncoding=UTF-8"
/>
```

Replace *database-host*, *db-name*, *user* and *password* in the above definitions with the names you have defined earlier (see [section 3.8](#)).

9. Set up indexing by inserting the following elements in **/opt/tomcat/conf/context.xml** as children of the root **Context** element:

```
<Environment
  name="escenic/indexer-webservice"
  value="http://indexer-web-service-host:8080/indexer-webservice/index/"
  type="java.lang.String"
  override="false"/>

<Environment
  name="escenic/index-update-uri"
```





```

value="http://localhost:8080/solr/update/"
type="java.lang.String"
override="false"/>

<Environment
name="escenic/solr-base-uri"
value="http://localhost:8080/solr/"
type="java.lang.String"
override="false"/>

<Environment
name="escenic/head-tail-storage-file"
value="/var/lib/escenic/head-tail.index"
type="java.lang.String"
override="false"/>

<Environment
name="escenic/failing-documents-storage-file"
value="/var/lib/escenic/failures.index"
type="java.lang.String"
override="false"/>

```

where *indexer-web-service-host* is the host name or IP address of the **editorial-host** on which the Content Engine's indexer web service is to run. If you are creating a single-host installation, then you can use the host name `localhost`.

10. Open Tomcat's `web.xml` file (`/opt/tomcat/conf/web.xml`) for editing and insert links to the resources you have defined:

```

<resource-ref>
  <description>Read link</description>
  <res-ref-name>jdbc/ECE_READ_DS</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>

<resource-ref>
  <description>Update link</description>
  <res-ref-name>jdbc/ECE_UPDATE_DS</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>

```

These elements must be inserted as children of the root `web-app` element: otherwise, position is irrelevant.

11. Open `/opt/tomcat/conf/server.xml` for editing.
12. Somewhere in this file you will find a `Connector` element that configures connections on port 8080. Make sure that this element contains a `URIEncoding` attribute, and that it is set to `UTF-8`. For example:

```

<Connector port="8080" protocol="HTTP/1.1"
connectionTimeout="20000"
redirectPort="8443"
URIEncoding="UTF-8"/>

```

This ensures that Content Engine search functionality works for non-Latin characters.

### 3.10 Create Configuration Folder

While logged in as `root`, create an `escenic` folder under `/etc`, and change its owner to `escenic`:

```

# mkdir /etc/escenic
# chown escenic:escenic /etc/escenic

```

-----  
 If you are installing everything on one host, then:  
 -----

While logged in as **escenic**, create an **engine** folder under **/etc/escenic**:

```
$ mkdir /etc/escenic/engine
```

-----  
 If you are creating a multi-host installation, then:  
 -----

On one of your **engine-hosts**, while logged in as **escenic**, create a symbolic link from **/etc/escenic/engine** to **/mnt/ece-conf**:

```
$ ln -s /mnt/ece-conf /etc/escenic/engine
```

---

## 3.11 Unpack Content Engine Components

If you are installing everything on a single host or if you created a shared file system as described in [section 3.5](#), then you should have direct access to the Content Engine packages you downloaded earlier (see [section 3.6](#)). If not, you will need to copy it to some temporary location on the **assembly-host**. The following instructions assume that you have downloaded the packages to the shared folder **/mnt/download**.

On your **assembly-host**, while logged in as **root**:

1. Create a folder for the Content Engine under **/opt**:

```
# mkdir /opt/escenic
# chown -R escenic:escenic /opt/escenic
```

2. Change user to **escenic** and unpack the Content Engine package as follows:

```
# su - escenic
$ cd /opt/escenic/
$ jar -xf /mnt/download/engine-5.2.7.2.zip
```

3. Create a symbolic link from **/opt/escenic/engine** to the Content Engine folder, so that it will be easier to upgrade to new versions when necessary.

```
$ ln -s engine-5.2.7.2 engine
```

4. Create a folder and symbolic link for the assembly tool, then unpack the assembly tool into the folder:

```
$ mkdir assemblytool-2.0-1
$ ln -s assemblytool-2.0-1 assemblytool
$ cd assemblytool/
$ jar xf /mnt/download/assemblytool-2.0-1.zip
```

-----  
 Note that the **jar** command shown above is part of the JDK, and is located in the **bin** directory of the Java distribution. Some Linux distributions ship with a GNU **jar** command which has a different syntax and does not support all the same options. Make sure you use the JDK's **jar** command.  
 -----



## 3.12 Install a Daemon Script

The following instructions tell you how to install a daemon script that will automatically run the Content Engine on start-up. It won't work properly yet since the Content Engine itself is not yet deployed. You should nevertheless install it now because the script also creates a number of folders that you will need later on in the installation process.

Log in as `root` on your **assembly-host** and on each of your **engine-hosts** and do the following:

1. Copy the daemon script `ece` to `/etc/init.d/`, and the script's default settings (`ece.defaults`) to a file called `ece` in the `/etc/default` folder. On the **assembly-host** you can do this as follows:

```
# cp /opt/escenic/engine/contrib/init.d/ece /etc/init.d/
# cp /opt/escenic/engine/contrib/init.d/ece.defaults /etc/default/ece
```

On the engine-hosts, however, you will have to copy the files from the **assembly-host**. If you have installed an SSH server on the **assembly-host** and SSH clients on your **engine-hosts**, then you can do it as follows:

```
# scp escenic@assembly-host-ip-address:/opt/escenic/engine/contrib/init.d/ece /etc/init.d/
# scp escenic@assembly-host-ip-address:/opt/escenic/engine/contrib/init.d/ece.defaults /etc/default/ece
```

2. Make the daemon script executable:

```
# chmod +x /etc/init.d/ece
```

3. The next step is to run the daemon script, `ece`.

If, for any reason, you do not want to install Content Engine files in standard locations, then you should open `/etc/default/ece` in an editor now and modify the settings to meet your requirements. You are strongly advised, however, not to do so. All instructions in this manual and the **Escenic Content Engine Server Administration Guide** assume files are installed in standard locations.

Enter the following command:

```
# /etc/init.d/ece
```

-----

Executing this command will result in an error because the `ece` daemon script tries to execute the `ece` startup script, which is not yet installed. **You can ignore this error message.** The purpose of running the script now is to create the following folders:

```
/var/crash/escenic
/var/lib/escenic
/var/log/escenic
/var/run/escenic
/var/cache/escenic
/var/spool/escenic
```

The script creates these folders and assigns them to the **escenic** user.

- 
4. On your **engine-hosts** only, enter the following commands:

```
# update-rc.d ece defaults
```

This command actually installs the script as a daemon, ensuring that it will be run on system startup. Note that this command is specific to Debian-based Linux distributions: the procedure for installing daemon scripts is different on other distributions.

---

### 3.13 Logging

On one of your **engine-hosts**, while logged in as **escenic**:

1. Create the folder **/etc/escenic/engine/common**:

```
$ mkdir /etc/escenic/engine/common
```

2. Create a logging configuration file called **trace.properties** in the new folder. This file should have the following content:

```
log4j.rootCategory=ERROR
log4j.category.com.escenic=ERROR, ECELOG
log4j.category.neo=ERROR, ECELOG

log4j.appender.ECELOG=org.apache.log4j.DailyRollingFileAppender
log4j.appender.ECELOG.File=/var/log/escenic/engine/ece-messages.log
log4j.appender.ECELOG.layout=org.apache.log4j.PatternLayout
log4j.appender.ECELOG.layout.ConversionPattern=%d %5p [%t] %x (%c) %m%n
```

On all your **engine-hosts**, while logged in as **escenic**:

1. Create the output folder specified in **trace.properties**:

```
$ mkdir /var/log/escenic/engine
```

2. In order to make Tomcat use the **trace.properties** file, create a link to it from **/opt/tomcat/lib**:

```
$ cd /opt/tomcat/lib
$ ln -s /etc/escenic/engine/common/trace.properties
```

3. In order to make Tomcat output its log files to a standard Unix/Linux location, open **/opt/tomcat/conf/logging.properties** for editing and set **all** the log folder properties (they have names ending with **FileHandler.directory**) to **/var/log/tomcat**.

---

### 3.14 Copy Solr Configuration

The Solr configuration data supplied with the Content Engine installation needs to be copied to locations where it can be accessed by all the Solr instances running in your system. By default, this means that you will need a copy of the Solr configuration data in the **/var/lib/escenic/solr/default** folder on each of your **engine hosts**.



To achieve this on a single host installation, log in as **escenic** and enter the following commands:

```
$ mkdir -p /var/lib/escenic/solr/default
$ cp -r /opt/escenic/engine/solr/* /var/lib/escenic/solr/default/
```

On a multiple host installation, log in as **escenic** on your **assembly host** and copy the Solr configuration data to the required destination folder on each **engine host**. For example, enter the following commands for each **engine host**:

```
$ ssh engine-host "mkdir -p /var/lib/escenic/solr/default"
$ scp -r /opt/escenic/engine/solr/* engine-host:/var/lib/escenic/solr/default
```

### 3.15 Initialize the Assembly Tool

On your **assembly-host**, while logged in as **escenic**:

1. **cd** to the **/opt/escenic/assemblytool** folder:

```
$ cd /opt/escenic/assemblytool
```

2. Enter the following command:

```
$ ant initialize
```

This command creates the **.properties** files the assembly tool needs.

3. Open **/opt/escenic/assemblytool/assemble.properties** for editing.
4. Uncomment the **engine.root** setting near the top of the file, and set it as follows:

```
engine.root = /opt/escenic/engine
```

### 3.16 Initialize the Bootstrap Layer

On your **assembly-host**, while logged in as **escenic**:

1. **cd** to the **/opt/escenic/assemblytool** folder:

```
$ cd /opt/escenic/assemblytool
```

2. Enter the following command:

```
$ ant ear
```

This command creates the **bootstrap layer**. The bootstrap layer is a set of configuration files that "configure the configuration process". Once you have run the above command, you will find the bootstrap layer located in **/opt/escenic/assemblytool/conf**. This bootstrap layer is set up to look for three configuration layers. The first two layers (**default** and **addon**) are part of the delivered system, and do not concern you. The third layer is called **common**, and the bootstrap layer is set up to look for it in **/etc/escenic/engine/common**. You will therefore need to create a common configuration layer there, as described in [section 3.17](#).

## 3.17 Create the Common Configuration Layer

The Content Engine package installed in the [section 3.11](#) step contains a **skeleton configuration layer**. This is a folder tree containing configuration files that you can use as a basis for your configurations. All the property settings in these files are commented out. To create a configuration, you copy this tree to the location specified in your bootstrap layer (e.g. `/etc/escenic/engine/common`) and edit the contents to set the properties you need. You will find the skeleton configuration layer located in `/opt/escenic/engine/siteconfig/config-skeleton`.

On your **assembly-host**, while logged in as `escenic`:

1. Copy the skeleton configuration layer to `/etc/escenic/engine/common`:

```
$ cp -r /opt/escenic/engine/siteconfig/config-skeleton/* /etc/escenic/engine/common/
```

2. Copy the delivered security configuration files to `/etc/escenic/engine/common`

```
$ cp -r /opt/escenic/engine/security/ /etc/escenic/engine/common/
```

3. Open `/etc/escenic/engine/common/ServerConfig.properties` for editing, and make sure the following properties are uncommented and set correctly for your site:

**databaseProductName**

The name of the database you are using. Currently allowed values are:

- **MySQL**

**ldapProductName**

The name of the LDAP server you are using. Currently allowed values are:

- **OpenLdap**

**ldapProductVersion**

The version number of the LDAP server you are using. In the case of OpenLDAP, you can get the version number by entering:

```
# slapd -v
```

**filePublicationRoot**

The path of the folder in which all binary files served by the Content Engine are stored (multimedia files - video, audio, images, Word documents, PDF files and so on). You should set this to:

```
filePublicationRoot=/var/lib/escenic/publications/
```

.....  
 The path you specify **must** include a trailing slash.  
 .....

**webPublicationRoot**

The URL root you want to be used for all your publications. For example `http://my-company.com/`. The URLs of your publications are formed by appending the publication name to this string. This property is in general used only for development purposes. In production, it is often the case that each publication has its own root URL.

4. Open `/etc/escenic/engine/common/Ldap.properties` for editing, and make sure the following properties are uncommented and set correctly for your site:

**providerUrl**

Must be set as follows:

```
providerUrl=ldap://ldap-host-ip-address:389/dc=your-company,dc=com
```

where:

- *ldap-host-ip-address* is the host name or IP address of the LDAP server, as specified when you installed it (see [section 3.7](#)).
- *your-company* is the same company name you used when installing the LDAP server (see [section 3.7](#)).

**userid**

The user ID that will be used to log in to the LDAP server (often called the **principal**). It can be specified either as a distinguished name or as a search specification. The user you specify must have access to everything stored below the root node specified with **root** below. For development or test systems, you can use the root user. In the case of OpenLDAP, this means that you should enter the name you entered in response to the `LDAP account for root` prompt during installation (see [section 3.7](#)):

```
userid=cn=admin,dc=your-company,dc=com
```

**password**

The password that will be used to log in to the LDAP server. If you are using the root user for this purpose, then in the case of OpenLDAP, you should enter the password you entered in response to the `Administrator password:` prompt during installation (see [section 3.7](#)):

```
password=password
```

**basedn**

The base **distinguished name** (dn) you specified when defining the LDAP schema:

```
basedn=dc=your-company,dc=com
```

**root**

The name of the Content Engine's root node. You can set this to:

```
root=ou=ECE
```

- Open `/etc/escenic/engine/common/connector/ReadConnector.properties` for editing, and make sure the following property is uncommented and set correctly:

**dataSourceName**

Must be set as follows:

```
dataSourceName=java:comp/env/jdbc/ECE_READ_DS
```

This is the name of the read database pool that you defined when installing the application server (see [section 3.9](#)).

Do not modify any of the other settings in this file.

- Open `/etc/escenic/engine/common/connector/UpdateConnector.properties` for editing, and make sure the following property is uncommented and set correctly:

**dataSourceName**

Must be set as follows:

```
dataSourceName=java:comp/env/jdbc/ECE_UPDATE_DS
```

This is the name of the update database pool that you defined when installing the application server (see [section 3.9](#)).

Do not modify any of the other settings in this file.

- Open `/etc/escenic/engine/common/neo/io/managers/ContentManager.properties` for editing, and make sure the following properties are uncommented and set correctly:

**readConnector**

Must be set as follows:

```
readConnector=/connector/ReadConnector
```

**updateConnector**

Must be set as follows:

```
updateConnector=/connector/UpdateConnector
```

- If you are creating a multi-host installation, then you should skip this step now, and add these settings to your **editorial hosts'** configuration layer (see [section 3.18](#)).  
-----

Open `/etc/escenic/engine/common/com/escenic/webstart/StudioConfig.properties` for editing, and make sure the property `property.com.escenic.client.webservice.url` is uncommented and set as shown below:

```
property.com.escenic.client.webservice.url=http://host-name:port/webservice/index.xml
```

where:

*host-name* is the host name or IP address of your **editorial host** and *port* is the number of the port on which Tomcat is listening. This property





identifies the web service that Content Studio communicates with: if you do not set it correctly, then Content Studio will not work.

You can also set the property `vmargs` in this properties file. You can use this property to set arguments to be passed to the JVM when Content Studio is run, should this be necessary.

---

## 3.18 Create Host Configuration Layers

---

.....  
 If you are installing everything on one host, then skip this section.  
 .....

In a multi-host installation, a small number of configuration properties may need to be set differently on different hosts. This can be achieved by creating individual configuration layers for each host. The Content Engine instance on each host reads its host configuration layer **after** the common configuration layer, and any settings it finds there override both system defaults and any values set in the common configuration layer.

On your **assembly-host**, while logged in as `escenic`:

1. Create a folder called `/etc/escenic/engine/host`, and under it create a folder for each machine that is to host a Content Engine instance. If, for example, your installation has one **editorial host** (called `editorial1`) and one **presentation host** (called `presentation1`), then you should enter:

```
$ mkdir /etc/escenic/engine/host/
$ mkdir /etc/escenic/engine/host/editorial1/
$ mkdir /etc/escenic/engine/host/presentation1/
```

2. Copy `/etc/escenic/engine/common/com/escenic/webstart/StudioConfig.properties` from the skeleton configuration layer to your editorial configuration layers:

```
$ mkdir /etc/escenic/engine/host/host-name/com/escenic/webstart/
$ cd /etc/escenic/engine/host/host-name/com/escenic/webstart/
$ cp /opt/escenic/engine/siteconfig/config-skeleton/etc/escenic/engine/common/com/escenic/webstart/StudioConfig.properties .
```

Open `/etc/escenic/engine/host/host-name/com/escenic/webstart/StudioConfig.properties` for editing, and make sure the property `property.com.escenic.client.webservice.url` is uncommented and set as shown below:

```
property.com.escenic.client.webservice.url=http://host-name:port/webservice/index.xml
```

where:

`host-name` is the host name or IP address of your **editorial host** and `port` is the number of the port on which Tomcat is listening. This property identifies the web service that Content Studio communicates with: if you do not set it correctly, then Content Studio will not work.

You can also set the property `vmargs` in this properties file. You can use this property to set arguments to be passed to the JVM when Content Studio is run, should this be necessary.

If there are any other configuration settings that you know need to be set differently on a specific host, you can set them in the same way:

1. Copy the appropriate `.properties` file from the skeleton configuration layer to the same relative location in the appropriate host configuration layer.
2. Open the copied file for editing.
3. Uncomment the required property and set it to the required value.

For further information about the configuration layers, see **Escenic Content Engine Server Administration Guide, chapter 4**.

### 3.19 Install the ece Script

The `ece` shell script provides a set of commands for easily stopping, starting and assembling the Content Engine.

-----  
 The `ece` shell script is **not** the same as the daemon script (also called `ece`) which you installed earlier (see [section 3.12](#)). The daemon script manages running the Content Engine to run as a service by automatically calling this script on system startup and shutdown.  
 -----

To install the `ece` shell script:

On your **assembly-host**, while logged in as `escenic`:

1. Make `/opt/escenic/engine/bin/ece` executable:

```
$ chmod +x /opt/escenic/engine/bin/ece
```

2. Open `/home/escenic/.bashrc` for editing, and add the following line to it in order to add `/opt/escenic/engine/bin` to your `PATH`:

```
export PATH=/opt/escenic/engine/bin:$PATH
```

3. Then enter the following command to apply the change you have made in your current shell:

```
$ source ~/.bashrc
```

4. Copy the script's configuration file to `/etc/escenic/engine`:

```
$ cp /opt/escenic/engine/etc/ece.conf /etc/escenic/engine/
```

This is the configuration file that will be read when `ece` is executed on the **assembly-host**.

5. -----  
 If you are installing everything on one host, then skip this step.  
 -----



Copy the script's configuration file to each of the host configuration folders you have made. For example:

```
$ cp /opt/escenic/engine/etc/ece.conf /etc/escenic/engine/host/editorial1/
$ cp /opt/escenic/engine/etc/ece.conf /etc/escenic/engine/host/presentation1/
```

These are the configuration files that will be read when **ece** is executed on the various **engine-hosts**. (If your **assembly-host** is also an **engine-host**, then you do not need a copy in `/etc/escenic/engine`, the host-specific copy will be used.)

6. Open each of these copies of `ece.conf` for editing, and check all the settings in them; make sure that they are set correctly for your installation, and for the specific host they correspond to. If you have followed all the instructions in this guide exactly, then the following settings should work:

**java\_home**

Make sure that this is set to the path of the Java virtual machine included in the JDK you installed (see [section 3.1](#)).

**ece\_home**

Set this to `/opt/escenic/engine`.

**ece\_server\_hostname**

Set this to the host name or IP address of the RMI hub to be used at this installation (not relevant for single-host installations).

**ece\_security\_configuration**

Set this to `/etc/escenic/engine/common/security`.

**appserver**

Currently, the only valid setting for this is `tomcat`.

**tomcat\_home**

Set this to `/opt/tomcat`.

**solr\_home**

Set this to `/var/lib/escenic/solr/default`. This is the location to which you have copied the Solr configuration data (see [section 3.14](#)).

**assemblytool\_home**

Set this to `/opt/escenic/assemblytool`.

Then, on each of your **engine-hosts**, while logged in as `root`:

1. Copy the `ece` script from the **assembly-host**. If you have installed an SSH server on the **assembly-host** and SSH clients on your **engine-hosts**, then you can do this as follows:

```
# scp escenic@assembly-host-ip-address:/opt/escenic/engine/bin/ece /usr/local/bin/
```

where `assembly-host-ip-address` is the host name or IP address of your **assembly-host**.

2. Make `/usr/local/bin/ece` executable:

```
# chmod +x /usr/local/bin/ece
```

3. Verify that the script is correctly installed by entering the following while logged in as **escenic**:

```
$ ece
```

This should give the following output:

```
Usage: /opt/escenic/engine/bin/ece <command>
where command is one of:
  start      starts the ECE
  stop       stops the ECE
  kill       uses force to stop ECE
  restart    restarts the ECE
  status     ECE status
  log        the ece-messages.log
  outlog     the [ece] script log (system out log)
  applog     the app server log
  threaddump write a thread dump to standard out (system out log)
  assemble   runs the Assembly Tool
  deploy     deploys the assembled EAR
```

---

## 3.20 Assemble and Deploy

Assembly and deployment consists of the following steps:

1. On your **assembly-host**, while logged in as **escenic**, run the **ece** script to assemble a set of Content Engine applications:

```
$ ece assemble
```

This generates an **enterprise archive** (EAR file) which you can deploy on all your **engine-hosts**. The EAR file is called **engine.ear**, and is created in the `/var/cache/escenic/` folder.

2. -----  
If you are installing everything on one host, then skip this step.  
-----

On each **engine-host**, copy `/var/cache/escenic/engine.ear` from the **assembly-host**. If you have installed an SSH server on the **assembly-host** and SSH clients on your **engine-hosts**, then you can do this as follows:

```
$ scp -r escenic@assembly-host-ip-address:/var/cache/escenic/engine.ear /var/cache/escenic/
```

where *assembly-host-ip-address* is the host name or IP address of your **assembly-host**.

3. On each **engine-host**, deploy the EAR file by entering:

```
$ ece deploy
```



### 3.21 Verify the Installation

It should now be possible to start the Content Engine and verify that the major components are working. To do so, carry out the following procedure on each **engine-host**, while logged in as **escenic**:

1. Start the Content Engine by entering:

```
$ ece start
```

2. Verify that the Content Engine is running by entering:

```
$ ece status
```

This should produce something like the following output:

```
[ece] Escenic Content Engine IS running (PID 1936) up 0d 1h 33m 51s
```



3. Open a browser, and try accessing the **escenic-admin** web application on each of the **engine-hosts** by entering:

```
http://engine-host-ip-address:8080/escenic-admin/
```

in the browser address field. There may be a delay while the Content Engine initializes, but eventually, you should see the following page:

4. Click on the **Status** link to display a set of status pages containing Content Engine test results.

Property Name	Status	Value	Action
com.escenic.config	✓	/	Help
escenic.server	✓	lean-editor1	Help


5. Across the top of these pages is a menu containing links to each of the status pages. Click on each of these links in turn and verify that all the tests pass (indicated by a  icon).
6. The  icon indicates that something is not correctly configured. If you see one of these icons, then you need to check over all relevant settings. In most cases there is a help link next to the test result indicating the most likely causes of failure. Make any necessary corrections, then restart the Content Engine by entering:

```
$ ece restart
```

7. Wait a minute or so, then refresh the browser window to re-execute the tests that failed.

Verify that the Content Engine is correctly configured on **all** your **engine-hosts** before proceeding any further.

---

On some of the test pages you will see  warning icons. In most cases these are simply an indication that no publications have yet been created, and they can be ignored.

---



---

## 3.22 Create a Publication

To really verify that everything is working correctly, you need a publication in the database. The quickest way to achieve this is to add one of the demo publications included with the installation. To add the **demo-temp-dev** publication (which is a very simple, stripped down publication with almost no content):

1. Download the publication WAR file from one of your **engine-hosts** to the machine on which you are working. You will find **demo-temp-dev.war** in the `/opt/escenic/engine/contrib/wars` folder.
2. Open a browser, and access the **escenic-admin** web application on one of your **editorial-hosts** by entering:

```
http://editorial-host-ip-address:8080/escenic-admin/
```



in the browser address field. The following page is displayed:

3. Click on **New pubs**. The following page is displayed:

4. Click on the **Browse** button and locate the **demo-temp-dev.war** file you have copied to your local machine.

5. Click on **Upload**. The following page is displayed:

**Upload resources**

[Home](#)  
[List pubs](#)  
[New pubs](#)

### Uploading resources

This page allows you to upload resources to the server. Once you have uploaded the required resources, you will be able to create publication

### Create Publication

You now have enough resources to [create a publication](#)

### Available resources

You have now uploaded the following resources. Once you have uploaded enough valid resources, you can create your publication.

Type	Age	Uploaded Content Type	File name	Size
/escenic/content-type	00	application/vnd.escenic.content-type	META-INF/escenic/publication-resources/escenic/content-type	5849
/escenic/feature	00	text/plain	META-INF/escenic/publication-resources/escenic/feature	380
/escenic/image-version	00	application/vnd.escenic.image-version	META-INF/escenic/publication-resources/escenic/image-version	841
/escenic/layout	00	application/vnd.escenic.layout+xml	META-INF/escenic/publication-resources/escenic/layout	194
/escenic/layout-group	00	application/vnd.escenic.layout.group+xml	META-INF/escenic/publication-resources/escenic/layout-group	1684
/escenic/teaser-type	00	application/vnd.escenic.content-type	META-INF/escenic/publication-resources/escenic/teaser-type	221
/escenic/content	00	application/vnd.escenic.content+xml	META-INF/escenic/resources/escenic/content	5811

### Upload or overwrite a resource

Type of resource:  Publication Definition (will replace any current definitions)  
 Content Type definitions  
 Feature definitions  
 Image Version definitions  
 Layout definitions  
 Layout-group definitions  
 Content definitions  
 Other type of resource

Please specify:

File to upload:

6. Click on the **create a publication** link. The following page is displayed:

**Create publication**

[Home](#)  
[List pubs](#)  
[New pubs](#)

The publication will be created based on the publication definition files you have uploaded. If you want to use different files, you can [upload different files](#).

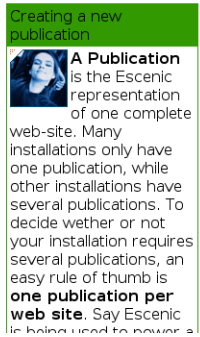
Publication Name:

Note: The name must be a symbolic name, with no spaces or punctuation. The publication name is simply an administrative name for the site. For an intranet/internet server good examples include "intranet" or "internet"; while for a hosting site, consider using a symbolic customer name, e.g. "somecompany" or "johndoe"

Administrator password:

Verify password:

Creating a new publication







7. Enter a name for the publication and an administrator password in the displayed form, then click on **Submit**. The following page is displayed:

**Menu**

[Home](#)  
[List pubs](#)  
[New pubs](#)

**Congratulations!**

You have successfully created the publication called "test". You can now...

**Log in using Studio**  
Use Escenic Content Studio to log on to this publication to start adding content immediately.

**Log in from the Web**  
Use the [Escenic Web interface](#) to log on to this publication to start administering it. Here you can create new users, create a section hierarchy. Use the user ID "test\_admin" to log on to get publication administrative rights.

**Deploy the web application in your app server**  
The publication can be deployed on this (or another) server. Use your application server administration tools to deploy this web application in order for the web application to respond to requests.

**Browse the publication**  
If you have deployed the web application, you may browse the publication at <http://localhost:8080/test/>. If the URL above does not work, try to deploy the web application to your application server.

**The Multimedia Repository**  
Multimedia files (images and attachments) are uploaded to `/opt/tomcat/webapps/test/multimedia/` on this server.

**Who is logged on?**  
[See who is currently logged in.](#)

**Export content**  
Export all contents of the publication "test" .

**Publication resources (download)**

- [/escenic/content-type](#)
- [/escenic/feature](#)
- [/escenic/image-version](#)
- [/escenic/layout](#)
- [/escenic/layout-group](#)
- [/escenic/teaser-type](#)

The publication "test" has publication ID 6.

NOTE: There is a xml import running in the background to add content to your publication.

You have now created a publication. If you click on **Home** and then **Status** to redisplay the status pages, you should see that all the warnings that were displayed earlier have now been replaced by icons. This should be the case not only on the **editorial-host** where you are currently working, but on all your **engine-hosts**.

The **demo-temp-dev** publication is now present in the database, and can be edited, but cannot be viewed from a browser until the **demo-temp-dev** web application is deployed on the application server. To deploy **demo-temp-dev**, do the following on your **assembly-host**, while logged in as **escenic**:

1. Copy **demo-temp-dev.war** to the assembly tool's **publications** folder:

```
$ cp /opt/escenic/engine/contrib/wars/demo-temp-dev.war /opt/escenic/assemblytool/publications/
```

2. Create a text file called **demo-temp-dev.properties** in the **/opt/escenic/assemblytool/publications/** folder, with the following contents:

```
name:          demo-temp-dev
source-war:    demo-temp-dev.war
context-root:  /demo-temp-dev
```

3. Reassemble the Content Engine installation:

```
$ ece assemble
```

Then, on each **engine-host**, while logged in as **escenic**:

1. -----  
If you are installing everything on one host, then skip this step.  
-----

On each **engine-host**, copy `/var/cache/escenic/engine.ear` from the **assembly-host**. If you have installed an SSH server on the **assembly-host** and SSH clients on your **engine-hosts**, then you can do this as follows:

```
$ scp -r escenic@assembly-host-ip-address:/var/cache/escenic/engine.ear /var/cache/escenic/
```

where *assembly-host-ip-address* is the host name or IP address of your **assembly-host**.

2. Redeploy the Content Engine:

```
$ ece deploy
```

3. Restart the Content Engine:

```
$ ece restart
```

4. Open a browser and enter the following URL in the address field:

```
http://engine-host-ip-address:8080/publication-name/
```

where *publication-name* is the name you gave to your **demo-temp-dev** publication. This should display the publication.

---

Publication web applications must be deployed on all hosts where the publications are to be viewed. This normally means all **engine-hosts** (including **editorial-hosts**, since writers and editors will want to preview their work). You may not want to deploy this test publication on all your hosts, however.

---

## 3.23 Test Web Studio

Now you have created a publication, you can use it to check that the Web Studio web application is correctly installed. To do this:


1. Open a browser, and access the Web Studio web application (called **escenic**) on one of your **editorial-hosts** by entering:

```
http://editorial-host-ip-address:8080/escenic/
```

in the browser address field. A login page should be displayed.

2. Enter the **demo-temp-dev** publication's administrator user name (*publication-name\_admin*) and the administrator password you specified when creating the publication.

- The following page should then be displayed:



## 3.24 Test Content Studio

To check that Content Studio is correctly installed:

- Make sure that a Java runtime environment is installed on your computer. If necessary, download the latest version from <http://www.java.com> and install it.
- Open a browser, and access the Content Studio download page on one of your **editorial-hosts** by entering:

```
http://editorial-host-ip-address:8080/studio/
```

in the browser address field. A download page should be displayed.

- Click on the **Launch Escenic Content Studio** button. After the application has downloaded, a login dialog should be displayed.
- Enter the **demo-temp-dev** publication's administrator user name (*publication-name\_admin*) and the administrator password you specified when creating the publication.
- The Content Studio application should then be downloaded, and start up after a few moments. For information on how to use Content Studio, see the **Escenic Content Studio User Guide**.





## 4 Recommended Modifications

This section contains instructions for various additions and modifications you can make to the basic installation described in [chapter 3](#). These changes are recommended for production installations.

### 4.1 Configure SSL Support

A production installation will almost always need SSL support, in order to be able to provide secure access to the Content Engine for remote users of Content Studio, Web Studio and so on.

To set up SSL support you must:

1. Obtain a certificate from a Certificate Authority (CA) such as VeriSign or Thawte. (But see [section 4.1.1](#).)
2. Install the certificate on every **engine host** for which you require HTTPS access. This means you should at least install the certificate on all your **editorial hosts**. You might, however, also want to install it on your **presentation hosts** in order to provide HTTPS access to some or all of your published content. For detailed instructions on how to install certificates for use by Tomcat, see <http://tomcat.apache.org/tomcat-5.5-doc/ssl-howto.html>.
3. Open `/opt/tomcat/conf/server.xml` for editing.
4. Somewhere in this file you will find a commented-out `Connector` element for configuring SSL connections on port 8443. Uncomment this element, and add `keystoreFile`, `keystorePass` and `URIEncoding` attributes to it as follows:

```
<Connector port="8443" protocol="HTTP/1.1" SSLEnabled="true"
  maxThreads="150" scheme="https" secure="true"
  clientAuth="false" sslProtocol="TLS"
  keystoreFile="keystore-path" keystorePass="password" URIEncoding="UTF-8"/>
```

`keystore-path` must be the path of the keystore you created in step 2. `password` must be the keystore/certificate password you created in step 2. The `URIEncoding="UTF-8"` attribute is required on all `Connector` elements in order to ensure that Content Engine search functionality works for non-Latin characters.

5. Open the configuration layer file `com/escenic/webstart/StudioConfig.properties` for editing, and change the protocol name in the web service URL used by by Content Studio. That is, change the setting of `property.com.escenic.client.webservice.url` from something like this:

```
property.com.escenic.client.webservice.url=http:host-name:port/webservice/index.xml
```

to something like this:

```
property.com.escenic.client.webservice.url=https:host-name:port/webservice/index.xml
```

If you installed everything on one host, then you will need to do this in your common configuration layer (`/etc/escenic/engine/common/com/escenic/webstart/StudioConfig.properties`). If you have a multi-host installation, then you will need to do it in one or more host configuration layers (`/etc/escenic/engine/host/host-name/com/escenic/webstart/StudioConfig.properties`).

---

The above procedure ensures that the Content Engine can support HTTPS access, but it does not enforce it in any way. Enforcement of HTTPS access to specific resources from specific locations can be achieved in a variety of ways and is outside the scope of this manual.

---

### 4.1.1 Using Self-Signed Certificates

---

The use of self-signed certificates is possible, but not recommended. Not only are self-signed certificates unverifiable, setting up Content Studio to work with self-signed certificates is complicated and error-prone.

---

A description of how to create and install self-signed certificates for Tomcat is included here: <http://tomcat.apache.org/tomcat-5.5-doc/ssl-howto.html>. In order to use Content Studio in combination with self-signed certificates, however, it is also necessary to install certificates on all the computers on which Content Studio is run. For a description of how to do this, see [http://blogs.sun.com/andreas/entry/no\\_more\\_unable\\_to\\_find](http://blogs.sun.com/andreas/entry/no_more_unable_to_find).

You will also need to set two Java system properties on each machine where Content Studio is run:

**`javax.net.ssl.trustStore`**

The path of the keystore (or trust store) on the computer.

**`javax.net.ssl.trustStorePassword`**

The keystore/certificate password you specified when adding the self-signed certificate. If you did not specify a password, the default is `changeit`.

If the certificates are identically installed on all computers from which Content Studio will be used, then you can include these settings in the `/com/escenic/webstart/StudioConfig.properties` file in your common configuration layer (see [section 3.17](#)). Otherwise they will need to be individually set on each computer.

---


## 4.2 Search Engine Deployment and Configuration

The Content Engine's search functionality is provided by the Java search engine Apache Solr, which runs as a web application. The standard Content Engine installation includes a `solr` web application and an associated `indexer` web application that indexes the content of all Escenic publications. These two



applications are deployed along with the Content Engine by the `ece` script's `deploy` action.

The result of following the basic installation procedure described in [chapter 3](#), therefore, is that a `solr` instance and `indexer` web application is deployed on every **engine host** in your installation, all with identical configurations.



This set up will work, but it is relatively inefficient and is unlikely to work well in a production environment. There are two main reasons for this:

### Solr memory usage

`solr` can at times consume large amounts of memory and trigger large garbage collection operations in the JVM, which has severe effects on Content Engine performance. It should therefore not be run in the same JVM as the Content Engine on production systems. The simplest way to achieve this separation on a single-host installation is to run `solr` and the `indexer` webapp in a separate Tomcat instance. For more about this, see **Escenic Content Engine Server Administration Guide, section 5.2.2**.

### Solr optimization issues

The default `solr` configuration is optimized for editorial purposes: it indexes all the fields needed to support the search functionality provided by Content Studio, resulting in very large indexes. This is acceptable in the editorial context, since the number of concurrent Content Studio users, even in a very large organisation, is not likely to be very large. The **presentation hosts** in a large Escenic installation, however, can be required to serve many thousands of concurrent users, and the default `solr` configuration may perform poorly in this context.

The default configuration, therefore, is fine for the **editorial hosts** in a production system, but for the **presentation hosts** you are recommended create a custom indexer configuration that only indexes

the fields actually needed to support the kinds of search required in your publications.

To do this, open `var/lib/escenic/schema.xml` for editing on each of your **presentation hosts**, and modify the index schema to meet your requirements. Editing this file is outside the scope of this manual. In order to tune the search engine you need to take account of the both the contents of your publications, your users' needs with regards to search and the limitations imposed by your particular hardware configuration. For further information and advice on tuning, see the Solr documentation on <http://lucene.apache.org/solr/>.

There are many more changes you can make to your search engine set-up in order to optimize it for your particular needs. For a discussion of the general principles involved, see **Escenic Content Engine Server Administration Guide, chapter 5**.

## 4.3 Distributed Memory Cache

A multiple-host installation in which each Content Engine instance maintains its own cache is not particularly efficient: it is better to pool the memory in a single, common cache. The most commonly-used means of doing this is to use the open source distributed cache manager **memcached**.

The instructions below tell you how to install **memcached**, configure it and then reconfigure your publications' **PresentationArticleCache** caches to use it. These caches will then share a single memory pool managed by **memcached**.

These instruction are based on the assumption that you want to install **memcached** on all your **engine-hosts**: **memcached** will then use some memory from each host. It might be the case, however, that you would prefer to concentrate the cache on a smaller number of hosts, or on one specific host. You may even wish to locate the cache on a completely different host, one that is not running the Content Engine. **memcached** is a completely free-standing application, and you can install it on any host or group of hosts in your network. The configuration described below is only one of many possible set-ups.

Note that if you are using **memcached**, the decorator must be serializable.

### 4.3.1 Install memcached

1. Install **memcached** on all your **engine hosts**.

On Debian-based Linux machines, you can do this by logging in as **root** and entering the following command:

```
# apt-get install memcached
```

2. Open `/etc/memcached.conf` for editing. Find the following line:

```
-l 127.0.0.1
```





and comment it out:

```
# -l 127.0.0.1
```

This allows **memcached** to listen on all interfaces.

3. **memcached** is installed as a daemon and will be started automatically on system reboot. However, to start it now without rebooting, enter:

```
# /etc/init.d/memcached start
```

### 4.3.2 Install the memcached Java Libraries

In order to make **memcached** accessible to Java web applications, you need to also install a Java client library on all your engine hosts. To do this:

1. Log in on your **assembly-host** as **escenic**.
2. Download the Java client library for **memcached** from <http://www.whalin.com/memcached/#download> to a temporary location.

```
$ cd /tmp/
```

```
$ wget http://img.whalin.com/memcached/jdk5/log4j/java_memcached-release_2.0.1.tar.gz
```

3. Unpack the downloaded package:

```
$ tar xzf java_memcached-release_2.0.1.tar.gz
```

4. Copy the unpacked **memcached** JAR file to the **/opt/escenic/assemblytool/lib**:

```
$ cp java_memcached-release_2.0.1/java_memcached-release_2.0.1.jar /opt/escenic/assemblytool/lib/
```

5. Assemble and redeploy the Content Engine as described in [section 3.20](#).

### 4.3.3 Configure memcached

.....

The following instructions tell you to make changes in your common configuration layer, because they are based on the assumption that you want all your Content Engines to share a common cache. If this is not the case, then the following actions should be carried out in the host configuration layers of the hosts you want to use **memcached**.

.....

To configure **memcached** you need to add a **.properties** file to your common configuration layer.

1. Log in as **escenic** on one of your **engine hosts**.
2. Create a folder for the new properties file:

```
$ mkdir -p /etc/escenic/engine/common/com/danga/
```

3. Create a file called **sockIOPool.properties** in the new folder, open it for editing and add the following content:

```
$class=com.danga.MemCached.SockIOPool
# fill in memcached servers here.
servers=host-name1:11211,host-name2:11211
```

```

# how many connections to use
initConn = 10
minConn = 5
maxConn = 100
# idle time
maxIdle = 180000
maintSleep = 5000
# socket timeout
socketTO = 3000
failover = true
# a network lookup algorithm
nagle = false

```

where *host-name1*, *host-name2*, etc. are replaced by the names of the hosts on which you have installed **memcached**.

The above parameter settings should work satisfactorily in most cases. Should you need to modify them, consult the documentation available on the memcached web site (<http://memcached.org/>).

4. Open `/etc/escenic/engine/common/Initial.properties` for editing and add the following lines:

```

# using memcached
service.0.0-memcached-socket-pool=/com/danga/SockIOPool

```

#### 4.3.4 Configure PresentationArticleCache Instances

The final step is to configure **PresentationArticleCache** in all publication web applications to use **memcached**. To do this you need to:

1. Add a `/WEB-INF/localconfig/neo/xredsys/presentation/cache/` folder to each publication web application. For example:

```
$ mkdir -p webapp/WEB-INF/localconfig/neo/xredsys/presentation/cache/
```

This will need to be done on the system(s) where your publication web applications are developed and maintained. For general information about publication development and maintenance, see the **Escenic Content Engine Template Developer Guide**.

2. Create a file called **PresentationArticleCache.properties** in each of the new folders, and add the following content to each file:

```
$class=neo.util.cache.Memcached
```

3. Redeploy the publication web applications in the usual way (see the **Escenic Content Engine Template Developer Guide** for general information about this).

---

## 5 Installing on Windows

The Content Engine is most commonly installed on hosts running Ubuntu Linux, so the detailed installation instructions in this guide are based on a Linux installation. It is, however, also possible to install and run the Content Engine on Windows-based hosts, and this possibility is of interest to some users. This chapter therefore contains additional guidance that will be of help to administrators installing the Content Engine in Windows environments.

This chapter is by no means a complete installation guide for Windows. If you are installing the Content Engine on Windows, you will need to use the description of the standard Linux installation procedure and "translate" it to your Windows environment. This chapter simply contains some hints and additional advice that will help you with that process.

---

Note that the only Windows version on which version 5.2.7.2 of the Content Engine is supported is:

**Windows Server 2008 R2**

---

---

### 5.1 Windows Software Components

Vizrt recommend use of the following components on Windows:

#### **Java Development Kit**

Sun JDK 6 update 20, downloaded from

<http://java.sun.com/javase/downloads/index.jsp>

#### **Ant**

Version 1.7.0 ( compiled December 13, 2006), downloaded from

<http://ant.apache.org/>

#### **LDAP**

Version 2.4.21, released on 23rd May 2010, downloaded from

<http://www.userbooster.de/en/download/openldap-for-windows.aspx>

#### **MySQL**

Version 5.1.47, released on 20th May 2010, downloaded from

<http://dev.mysql.com/downloads/mysql/>

#### **Tomcat**

Version 6.0.26, released on 11th March 2010, downloaded from

<http://tomcat.apache.org/download-60.cgi>

### JDBC connector

`mysql-connector-java-5.1.12`, downloaded from

<http://www.mysql.com/downloads/connector/j/>

---

## 5.2 Windows Installation Hints

This section contains Windows-specific advice regarding the installation and configuration of various components.

### 5.2.1 General Advice

You are recommended to install the Content Engine in `c:\escenic`, or in a similar location with a path that does not contain any spaces or other special characters. Do not install it in `c:\Program Files`. You should install the assembly tool in a similar location (`c:\assemblytool`, for example).

As directory structure is different in windows and linux, you should be aware of the **Standard File Hierarchy** mentioned in [section 1.3.3](#) and map the directories to convenient locations in windows.

When editing Content Engine configuration files and other `.properties` file, make sure that you either:

- Replace all the backslashes in any paths you specify with forward slashes (for example, `c:/escenic`), or else
- Escape the backslashes by entering them twice (for example, `c:\\escenic`)

### 5.2.2 LDAP on Windows

There are several issues you need to be aware of when installing and configuring LDAP on Windows, described in the following sections.

#### 5.2.2.1 Ldapadd Not Supported

The LDAP installation recommended in [section 5.2](#) does not support the `ldapadd` command. Therefore, step 5 in [section 3.7](#) will not work. Instead you must do as follows:

1. Open a terminal window.
2. Change directory to the location of `ldapmodify.exe` (probably something like `C:\Program Files (x86)\OpenLDAP\ClientTools`).
3. Enter the following command:

```
ldapmodify.exe -a -W -D "cn=admin,dc=escenic,dc=com" -x -f schema -h ip-address -p 389
```

where:



- *schema* is the path of your `schema.ldif` file
- *ip-address* is the IP address of your **ldap-host**

### 5.2.2.2 Editing slapd.conf

When editing `slapd.conf` (step 5 in [section 3.7](#)), make sure that you insert the recommended entries:

```
index objectClass,uid,member eq
index cn,mail,sn,givenname eq,subinitial
```

as **replacements** for the default `index` entry (`cn,mail,cn`) that is present in the default `slapd.conf` file.

### 5.2.2.3 Running LDAP

If you have trouble when running LDAP as a service, you can run it as an ordinary program instead:

1. Stop the `openldap` service from `run>"services.msc"`
2. Run the program `openldap-home\run\run.cmd` (where `openldap-home` is the folder in which you have installed LDAP).

### 5.2.3 MySQL on Windows

Running the Content Engine's database scripts on Windows is not quite as simple as the example shown in step 5 of [section 3.8](#). You will need to run each script individually as follows:

```
cd mysql-home\bin
mysql -u username -p database < engine\database\mysql\tables.sql
mysql -u username -p database < engine\database\mysql\indexes.sql
mysql -u username -p database < engine\database\mysql\constants.sql
mysql -u username -p database < engine\database\mysql\constraints.sql
```

where:

- *mysql-home* is the path of the folder in which you installed MySQL.
- *username* is the SQL user name.
- *database* is the name of the Content Engine database.
- *engine* is the path of the folder in which you installed the Content Engine.

### 5.2.4 Tomcat on Windows

There are a couple of issues you need to be aware of when installing and configuring Tomcat on Windows, described in the following sections.

#### 5.2.4.1 Startup Batch File

You may have security-related problems running the supplied startup batch file `tomcat-home\bin\startup.bat`. If this happens you should be able to fix the problem by creating a batch file of your own that makes the necessary security settings and then calls `tomcat-home\bin\startup.bat`. For example:

```
@echo off
```

```

rem #####
rem ### Set some initial properties ###
rem #####
rem ### Set the argument list ###
rem #####
set CATALINA_HOME=tomcat-home
set solr_home=engine-home\lib\escenic\solr
set ECE_ARGS=%ECE_ARGS% -Djava.security.auth.login.config=engine-home\security\jaas.config
set ECE_ARGS=%ECE_ARGS% -Djava.security.policy=engine-home\security\java.policy
echo ECE_ARGS=%ECE_ARGS%
set JAVA_OPTS=%JAVA_OPTS% %ECE_ARGS%

rem #####
rem ### Start Tomcat ###
rem #####
tomcat-home\bin\startup.bat
  
```

In the above batch file you must replace:

- *tomcat-home* with the path of your Tomcat installation
- *engine-home* with the path of your Content Engine installation

#### 5.2.4.2 Logging

Since Windows doesn't have an equivalent to the `ln` command, you cannot link from Tomcat to a `trace.properties` file in your Content Engine installation. You must therefore create a `trace.properties` file in your Tomcat installation's `lib` folder, and set up logging to the appropriate location. At a minimum it should contain something like:

```
log4j.appender.ECELOG.File=engine-home/log/escenic/engine/ece-messages.log
```

where *engine-home* is the location of your Content Engine installation.

#### 5.2.5 JDBC Connector on Windows

To install the JDBC connector on Windows:

1. Unzip the downloaded file.
2. Copy `mysql-connector-java-5.1.12-bin.jar` to your Tomcat installation's `lib` folder.

---

## 5.3 Assembly and Deployment on Windows

The `ece` script used for assembly and deployment on Linux platforms cannot be used in Windows. To assemble and deploy the Content Engine:

1. Open a terminal window on your **assembly host**.
2. `cd` to the folder where you have installed the assembly tool, for example:

```
CD c:\assemblytool
```

3. Enter the command below:

```
ant ear
```

4. Copy all the `.war` files in `c:\assemblytool\dist\war` to the Tomcat `webapps` folders on all your **engine hosts**.
5. Create `escenic\lib` folders in the Tomcat installation folders on all your **engine hosts**. If, for example, Tomcat is installed in `c:`



`\apache-tomcat-6.0.26`, then you need to create a `c:\apache-tomcat-6.0.26\escenic\lib` folder.

6. Copy all the `.jar` files in `c:\assemblytool\dist\work\ear\lib` on your **assembly-host** to the Tomcat `escenic\lib` folders you have just created on your **engine hosts**.