

Escenic Content Engine  
**Installation Guide**

6.3.2-1



# Table of Contents

<a href="#">1 Introduction</a>	4
<a href="#">1.1 Installation Components</a>	4
<a href="#">1.1.1 Escenic Components</a>	4
<a href="#">1.1.2 Third-Party Components</a>	5
<a href="#">1.2 Installation Types</a>	5
<a href="#">1.3 Conventions Used in This Manual</a>	6
<a href="#">1.3.1 Identifying Host Machines</a>	6
<a href="#">1.3.2 Shell Commands</a>	7
<a href="#">1.3.3 Standard File Hierarchy</a>	7
<a href="#">2 Product Requirements</a>	9
<a href="#">2.1 Required Supporting Software</a>	9
<a href="#">2.2 Other Requirements</a>	9
<a href="#">2.2.1 Shared Filesystem</a>	9
<a href="#">2.2.2 SSH/FTP Server</a>	10
<a href="#">2.2.3 Apache Ant</a>	10
<a href="#">3 Installation Procedure</a>	11
<a href="#">3.1 Preparation</a>	12
<a href="#">3.1.1 Install Java SE Development Kit (JDK)</a>	12
<a href="#">3.1.2 Install Ant</a>	12
<a href="#">3.1.3 Install Various Utilities</a>	13
<a href="#">3.1.4 Create Escenic Users</a>	13
<a href="#">3.1.5 Create Shared File System</a>	13
<a href="#">3.1.6 Install Database</a>	15
<a href="#">3.1.7 Install Application Server</a>	16
<a href="#">3.1.8 Logging</a>	19
<a href="#">3.1.9 Install Solr</a>	19
<a href="#">3.2 Package-based Installation</a>	22
<a href="#">3.2.1 Install Escenic Packages</a>	22
<a href="#">3.2.2 Initialize the Assembly Tool</a>	23
<a href="#">3.2.3 Initialize the Bootstrap Layer</a>	23
<a href="#">3.2.4 Edit the Common Configuration Layer</a>	24
<a href="#">3.2.5 Create Host Configuration Layers</a>	25
<a href="#">3.2.6 Configure the ece Script</a>	25
<a href="#">3.3 Old-style Installation</a>	26

3.3.1	<a href="#">Download Content Engine</a>	27
3.3.2	<a href="#">Create Configuration Folder</a>	27
3.3.3	<a href="#">Unpack Content Engine Components</a>	27
3.3.4	<a href="#">Initialize the Assembly Tool</a>	28
3.3.5	<a href="#">Initialize the Bootstrap Layer</a>	28
3.3.6	<a href="#">Create the Common Configuration Layer</a>	29
3.3.7	<a href="#">Create Host Configuration Layers</a>	30
3.3.8	<a href="#">Install the ece Script</a>	31
3.3.9	<a href="#">Configure the ece Script</a>	31
3.3.10	<a href="#">Copy the ece Script</a>	32
3.4	<a href="#">Finishing Up</a>	33
3.4.1	<a href="#">Assemble and Deploy</a>	33
3.4.2	<a href="#">Verify the Installation</a>	34
3.4.3	<a href="#">Install a Daemon Script</a>	35
3.4.4	<a href="#">Create a Publication</a>	36
3.4.5	<a href="#">Test Web Studio</a>	39
3.4.6	<a href="#">Test Content Studio</a>	40
4	<a href="#">Recommended Modifications</a>	41
4.1	<a href="#">Configure SSL Support</a>	41
4.1.1	<a href="#">Using Self-Signed Certificates</a>	42
4.2	<a href="#">Search Engine Deployment and Configuration</a>	42
4.3	<a href="#">Distributed Memory Cache</a>	44
4.3.1	<a href="#">Install memcached</a>	44
4.3.2	<a href="#">Configure memcached</a>	45
4.3.3	<a href="#">Configure PresentationArticleCache Instances</a>	46
4.4	<a href="#">Password Protect escenic-admin</a>	46
5	<a href="#">Upgrading</a>	48
5.1	<a href="#">Package-based Upgrade</a>	48
5.2	<a href="#">Old-style Upgrade</a>	48
5.3	<a href="#">Database Upgrades</a>	49

# 1 Introduction

This **Installation Guide** is intended to be read by the person responsible for installing the Escenic Content Engine on one or more hosts. 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 Escenic 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 Escenic 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.

### **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.

### **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, 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**

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

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

This is a common production installation structure for small sites.

**Multiple web servers, multiple application servers/Content Engines, 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.

**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 a enterprise archive or .EAR file.

**database-host**

The machine used to host the database 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).

**share-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/slapd 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. Some of the standard folders may vary according to whether you install Escenic components using the new **apt-** or **rpm-**based method or the old download archive-based method. For more about this, see [chapter 3](#).

Standard location	Component
<code>/usr/share/escenic [NEW METHOD]</code>	Escenic

<b>Standard location</b>	<b>Component</b>
<code>/opt/escenic</code> [OLD METHOD]	
<code>/usr/share/escenic/escenic-content-engine-6.3.2-1</code> [NEW METHOD] <code>/opt/escenic/engine</code> [OLD METHOD]	Escenic Content Engine
<code>/usr/share/escenic/escenic-assemblytool</code> [NEW METHOD] <code>/opt/escenic/assemblytool</code> [OLD METHOD]	Escenic assembly tool
<code>/usr/share/escenic/ece-scripts</code> [NEW METHOD] <code>/opt/escenic/engine/ece-scripts</code> [OLD METHOD]	Escenic scripts
<code>/etc/escenic</code>	Escenic configuration
<code>/etc/escenic/engine</code>	Escenic Content Engine configuration
<code>/opt/java/jdk</code>	Java
<code>/opt/java/ant</code>	Ant
<code>/opt/tomcat-engine1</code>	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, and uses a database to store content and 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)

#### **Application Server**

- Apache Tomcat version 7 (7.0.54 or later)

#### **Database Server**

One of the following:

- MySQL 5.6 (although any stable MySQL version > 5.0 should work)
- Oracle Database 10.1 or later

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

Oracle Java SE Development Kit 8. For information about upgrading an existing installation from Java 7 to Java 8, see [http://docs.escenic.com/ece-tech-notes/6.3/upgrading\\_to\\_java\\_8.html](http://docs.escenic.com/ece-tech-notes/6.3/upgrading_to_java_8.html).

#### **Web Server**

Apache Web Server (latest version)

### 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 3.1.5](#).

### 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.

### Choosing an installation method

Until recently, all Escenic-supplied components had to be installed by manually downloading and unpacking archive files. This is no longer the case. The Content Engine itself, the assembly tool and the ece scripts are now all available as **apt** packages (for installing on Ubuntu/Debian systems) and as **rpm** packages (for installation on RedHat/CentOS systems). This is now the recommended method of installing Escenic systems, although it is still possible to use the old method based on manually downloading and unpacking archives.

Using the new package installation method offers many advantages over the old manual installation procedure:

- It is faster and quicker
- It is significantly less error-prone
- It supports a fully automated upgrade process in which not only the installed packages themselves are upgraded, but also deployed EAR files. An upgrade script automatically checks the EAR files for copies of JAR files that have been updated, and replaces them with the new versions.

The old installation method will continue to be documented for the moment. Components installed using the new method are installed in `/usr/share/escenic`, whereas the old method recommends installation in `/opt/escenic`, and some other path components are slightly different. To cope with these differences, the following placeholders are used in some paths:

Placeholder	NEW METHOD path	OLD METHOD path
<i>engine-installation</i>	<code>/usr/share/escenic/ escenic-content- engine-6.3.2-1</code>	<code>/opt/escenic/engine</code>
<i>assemblytool_installation</i>	<code>/usr/share/escenic/ escenic-assemblytool</code>	<code>/opt/escenic/ assemblytool</code>

There are a number of preparatory tasks that you need to carry out whichever installation method you choose. These are described in [section 3.1](#). Once you have completed these tasks, you continue either by following the steps described in [section 3.2](#) (recommended) or [section 3.3](#). Finally, whichever installation method you choose, you will need to finish up by assembling and deploying an EAR file containing the supplied Escenic web applications, checking that everything is installed correctly and so on, as described in [section 3.4](#).

## 3.1 Preparation

First of all, it is necessary to install and configure the various software packages used by an Escenic system, and carry out a few other preparatory tasks.

### 3.1.1 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 8.

You can download the JDK from <http://www.oracle.com/technetwork/java/javase/downloads/index.html>. Follow the installation instructions supplied on the Oracle site. When you have installed it you need to make sure that it is installed as the default Java installation. You can do this by entering the following commands while logged in as **root**:

```
# update-alternatives --install "/usr/bin/java" "java" "java-install-path" 1
# update-alternatives --set java java-install-path
```

where *java-install-path* is the path of the location in which you installed Java (for example, **/usr/lib/jvm/java-8-oracle/jre/bin/java**).

Verify that Java is correctly installed by entering:

```
# java -version
```

This should result in output something like this:

```
java version "1.8.0_11"
Java(TM) SE Runtime Environment (build 1.8.0_11-b12)
Java HotSpot(TM) 64-Bit Server VM (build 25.11-b03, mixed mode)
```

### 3.1.2 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>.

You also need to install the additional library **ant-contrib** (see <http://ant-contrib.sourceforge.net/#install>).

It may also be possible to download **ant** and **ant-contrib** 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-contrib
```

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

```
# ant -version
```

This should produce a response something like this:

```
Apache Ant version 1.7.1 compiled on September 8 2010
```

### 3.1.3 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 as **root**:

```
# 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
```

### 3.1.4 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 - scenic
$ echo 'export JAVA_HOME=/usr/lib/jvm/java-8-oracle' >> .bashrc
$ echo 'export PATH=$JAVA_HOME/bin:$PATH' >> .bashrc
$ source ~/.bashrc
```

### 3.1.5 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 **share-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 **share-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/multimedia
# mkdir /mnt/download
```

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

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

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

4. Mount the shared folders:

```
# mount -a
```

5. Create a symbolic link to the multimedia shared folder from the `filePublicationRoot` path you will specify later in `ServerConfig.properties` (see [section 3.2.4](#) or [section 3.3.6](#)):

```
# ln -s /mnt/multimedia /var/lib/escenic/engine
```

### 3.1.6 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
```

Alternatively, you can install a binary distribution.

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 3.1.5](#), then you should have direct access to the Content Engine package you downloaded earlier (see [section 3.3.1](#)). 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 - escenic
$ cd /tmp
$ unzip /mnt/download/engine-6.3.2-1.zip
```

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

```
$ cd /tmp/engine-6.3.2-1/database/mysql/
$ for el in tables.sql indexes.sql constants.sql constraints.sql; do \
    mysql -u user -ppassword db-name < $el
done;
```

Replace *db-name*, *user* and *password* in the above commands with the names you chose in step 2. Your *password* must **immediately follow** the **-p** switch, with no intervening space (as shown above).

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.

On a single-host installation, you can check that MySQL is running by entering the following command:

```
$ mysqladmin -u root -p status
```

You should get a response something like this:

```
Uptime: 605 Threads: 1 Questions: 615 Slow queries: 0 Opens: 842 Flush tables: 1 Open
tables: 40 Queries per second avg: 1.16
```

If the server is not running then you will see an error reporting that it was not possible to connect to the server.

### 3.1.7 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 3.1.5](#), 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-engine1** 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-engine1
# chown -R escenic:escenic /opt/apache-tomcat-version-number
```

4. Open Tomcat's configuration file file (**/opt/tomcat-engine1/conf/catalina.properties**) for editing and add the following characters:

```
,${catalina.base}/escenic/lib/*.jar
```

to the end of the **common.loader** property setting.

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/j/>.

6. Change user back to **escenic** and install the driver by copying it to the `/opt/tomcat-engine1/lib` folder. If you downloaded the MySQL driver package to `/mnt/download`, then you can do this as follows:

```
# su - escenic
$ cd /mnt/download/
$ unzip mysql-connector-java-version-number.zip
$ cp mysql-connector-java-version-number/mysql-connector-java-version-number-
bin.jar \
> /opt/tomcat-engine1/lib/
```

7. Create the `/opt/tomcat-engine1/escenic/lib` folder you added to the `common.loader` path in step 4:

```
$ mkdir -p /opt/tomcat-engine1/escenic/lib/
```

8. Set up database pooling. For MySQL, you do this by opening `/opt/tomcat-engine1/conf/context.xml` for editing and inserting the following resource definition as a child of the root `Context` element:

```
<Resource
  name="jdbc/ECE_DS"
  auth="Container"
  type="javax.sql.DataSource"
  username="user"
  password="password"
  driverClassName="com.mysql.jdbc.Driver"
  maxTotal="200"
  maxIdle="10"
  maxWaitMillis="10000"
  removeAbandonedOnBorrow="true"
  removeAbandonedOnMaintenance="true"
  removeAbandonedTimeout="60"
  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 3.1.6](#)).

Note that the `maxTotal` and `maxWaitMillis` attribute names above are only correct if:

- You are running Tomcat 8, **and**
- Your Tomcat installation uses the default Apache Commons DBCP (see <http://commons.apache.org/proper/commons-dbcp/configuration.html> for details).

**If you are running Tomcat 7**, then you will need to change `maxTotal` to `maxActive` and `maxWaitMillis` to `maxWait`. For general information about the differences between Tomcat 7 and Tomcat 8, see <https://tomcat.apache.org/migration-8.html>.

**If you are running a Tomcat 8 that has been set up to use Tomcat's own DBCP**, then you will also need to change `maxTotal` to `maxActive` and `maxWaitMillis` to `maxWait`. For information about Tomcat's DBPC, see <http://tomcat.apache.org/tomcat-8.0-doc/jdbc-pool.html#Introduction>.

The parameter settings shown above are provided only as recommended starting values. They may need to be adjusted to provide optimum performance.

If you copy and paste the above resource definitions, make sure to remove any line breaks in the `url` values (caused by page width restrictions).

9. Set up indexing by inserting the following elements in `/opt/tomcat-engine1/conf/context.xml` as children of the root `Context` element:

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

<Environment
  name="escenic/index-update-uri"
  value="http://localhost:8983/solr/solr-core/update/"
  type="java.lang.String"
  override="false"/>

<Environment
  name="escenic/solr-base-uri"
  value="http://localhost:8983/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:

- On an **editorial-host**, `indexer-web-service-host` is the host name or IP address of the **editorial-host** on which the Content Engine's internal indexer web service is to run and `web-service-name` is **index**.
- On a **presentation-host**, `indexer-web-service-host` is the host name or IP address of the **presentation-host** on which the Content Engine's external indexer web service is to run and `web-service-name` is **presentation-index**.
- The `solr-core` is usually either **editorial** or **presentation**.

If you are creating a single-host installation, then you can use the host name `localhost` and the indexer web service name **index**.

Make sure that your `escenic/solr-base-uri` setting in `/opt/tomcat-engine1/conf/context.xml` is as shown, and **does not** include a trailing slash (`/`). If it does, then search will not work.

10. Open Tomcat's `web.xml` file (`/opt/tomcat-engine1/conf/web.xml`) for editing and insert links to the resources you have defined:

```
<resource-ref>
  <description>Escenic link</description>
  <res-ref-name>jdbc/ECE_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-engine1/conf/server.xml` for editing. 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.

### 3.1.8 Logging

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

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

```
| # mkdir -p /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-engine1/lib`:

```
| $ cd /opt/tomcat-engine1/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-engine1/conf/logging.properties` for editing and set **all** the log folder properties (they have names ending with **FileHandler.directory**) to `/var/log/tomcat`.

### 3.1.9 Install Solr

The Content Engine distribution used to include a **solr.war** file that could be deployed together with the Content Engine search. This is, however, not possible with more recent versions of Solr. Solr must now be installed as a standalone server on all your **engine** hosts.

To install Solr, log in as **root** on each **engine** host, and:

1. Download the latest stable release of Solr 6 from a suitable Apache mirror site, to a temporary location. For example:

```
# cd /tmp
# wget http://apache.uib.no/lucene/solr/solr-version/solr-solr-version.zip
```

where *solr-version* is the latest stable Solr 6 release number.

2. Install the contents of the download package in **/opt**:

```
# unzip /tmp/solr-solr-version.zip -d /opt
```

3. Create a symlink **/opt/solr** for the new installation folder:

```
# cd /opt
# ln -s solr-solr-version solr
```

4. Create the configuration and start-up files required to manage the Solr service, as described in [section 3.1.9.1](#).
5. Create a Solr **core** by copying the Solr configuration files supplied with the Content Engine to a suitable location on the host, as described in [section 3.1.9.2](#).

In Solr terminology, a **core** is the set of configuration files that define a specific application of Solr – a schema, an index and all related configuration files. A typical Escenic production installation makes use of two Solr cores – one for editorial purposes and one for presentation purposes. For further information about this, see [Customizing the Index Schema](#).

### 3.1.9.1 Create Solr Startup Files

To create the necessary Solr startup files:

1. Create a Solr core folder and **core.properties** file as follows:

```
# mkdir -p /var/lib/escenic/solr/solr-core
# cat > /var/lib/escenic/solr/solr-core/core.properties <<EOF
name=solr-core
config=solrconfig.xml
schema=schema.xml
dataDir=data
EOF
```

where *solr-core* is the name of your Solr core. Typically you might choose to use either **editorial** or **presentation** as the name of your core depending on its purpose (that is, what kind of host you are installing it on).

2. Change ownership of the Solr folder and all its contents, so that it can be used by the **escenic** user:

```
# chown -R escenic:escenic /var/lib/escenic/solr
```

3. Create a configuration file for the **/opt/solr/bin/solr** startup script. Do this by **moving** (not copying) the file as follows:

```
# mv /opt/solr/bin/solr.in.sh /etc/default/
```

4. Append settings for **SOLR\_HOME**, **SOLR\_LOGS\_DIR** and **SOLR\_PID\_DIR** to the moved script as shown below:

```
# cat >> /etc/default/solr.in.sh <<EOF
SOLR_HOME=/var/lib/escenic
SOLR_LOGS_DIR=/var/log/escenic
SOLR_PID_DIR=/var/run/escenic
EOF
```

This will ensure that Solr looks for the core in the right place (where you created it in step 1) and writes log files and runtime files to the same locations as corresponding Content Engine files.

5. Make sure that all the folders you specified in step 4 actually exist and that the **escenic** user is either the owner or has read/write permissions.
6. Add an **init.d** script for starting Solr automatically on boot:

```
# cp /opt/solr/bin/init.d/solr /etc/init.d/
# sed -i "s#RUNAS=\"solr\"#RUNAS=\"escenic\"#" /etc/init.d/solr
```

7. Add the **init.d** script to the desired run levels. For Debian- and Ubuntu-based systems, you can do this as follows:

```
# update-rc.d solr defaults 35
```

For other systems, please consult the appropriate documentation.

### 3.1.9.2 Copy Solr Configuration

You now need to copy the following Solr configuration data:

- The main Solr configuration file supplied with Solr itself
- Solr core configuration files supplied with the Content Engine

The configuration data should be copied into a folder under **/etc/escenic** and then linked to from the **/var/lib/escenic/solr/** tree as described below.

For a single host installation:

1. Log in as **root** and create a **solr** folder under **/etc/escenic**, then copy the **solr.xml** file supplied with Solr to **/etc/escenic**:

```
# mkdir -p /etc/escenic/solr
# cp /opt/solr-solr-version/server/solr/solr.xml /etc/escenic
```

where *solr-version* is the version number of your Solr installation.

2. Copy the configuration folder supplied with the Content Engine:

```
# cp -r engine-installation/solr/conf /etc/escenic/solr/solr-core
```

where *solr-core* is the name of your Solr core (most likely **editorial** on a single host installation).

3. Log in as **escenic** and create links to the copied file and folder as follows:

```
$ cd /var/lib/escenic/
$ ln -s /etc/escenic/solr.xml
$ cd solr/solr-core
$ ln -s /etc/escenic/solr/solr-core conf
```

For a multiple host installation you will need to repeat these steps on each host, and in this case, *solr-core* is likely to vary according to the type of **engine host** – either **editorial** or **presentation**. Before you start, copy the *engine-installation/solr/conf* folder from your **assembly host** to a temporary location (say **/tmp/conf**) on each engine host:

```
$ scp -r engine-installation/solr/conf engine-host:/tmp/
```

Then repeat the five steps listed above on each host, replacing step 2 with:

```
# mv /tmp/conf /etc/escenic/solr/solr-core
```

Note that the Solr configuration data supplied with the Content Engine is intended for use on an editorial host and is not ideally suited to presentation hosts, so for production installations, it will need to be modified. For more about this, see [Customizing the Index Schema](#).

## 3.2 Package-based Installation

This is the preferred installation method. You should use this method unless you have special requirements that dictate otherwise.

### 3.2.1 Install Escenic Packages

To install the required Escenic packages on your **assembly-host**, log in as **root** and perform the following steps:

1. If necessary, add the Escenic **apt** repository to your list of sources:

```
# echo "deb http://user:password@apt.escenic.com stable main non-free" >> /etc/
apt/sources.list.d/escenic.list
```

where *user* and *password* are your Escenic download credentials (the same ones you use to access the Escenic Maven repository). If you do not have any download credentials, please contact [Escenic support](#).

2. Add the repository signing key to your machine's keyring:

```
# curl --silent http://apt.escenic.com/repo.key | apt-key add -
```

3. Update your package lists and install the Content Engine and the Escenic scripts:

```
# apt-get update
# apt-get install escenic-content-engine-6.2
# apt-get install escenic-content-engine-updater-6.2
# apt-get plug-in-packages
```

where *plug-in-packages* is a list of any Escenic plug-in packages that you want to install together with the Content Engine.

4. Install the Content Engine plug-ins you need:

```
# apt-get plug-in-packages
```

5. Change the owner of the assembly tool that has been installed (as part of the `escenic-content-engine-updater-6.2` package) to **escenic**:

```
# chown escenic /usr/share/escenic/escenic-assemblytool-at-version/
```

where *at-version* is the version number of the installed assembly tool.

At multi-host installations you need to repeat this installation on all your **engine-hosts**. To do this, log in as **root** on each host and enter the following commands:

```
# echo "deb http://user:password@apt.escenic.com stable main non-free" >> /etc/apt/
sources.list.d/escenic.list
# curl --silent http://apt.escenic.com/repo.key | apt-key add -
# apt-get update
# apt-get install escenic-content-engine-6.2
# apt-get install escenic-content-engine-updater-6.2
```

```
# apt-get plug-in-packages
# chown escenic /usr/share/escenic/escenic-assemblytool-at-version/
```

The commands you need to use to install these packages on RedHat / CentOS systems are:

```
# rpm -Uvh https://user:password:yum.escenic.com/rpm/escenic-content-
engine-6.2-6.3.2-1.x86_64.rpm
# rpm -Uvh https://user:password:yum.escenic.com/rpm/escenic-content-engine-
updater-6.2-6.3.2-1.x86_64.rpm
```

where *version* is the correct version number of the package.

### 3.2.2 Initialize the Assembly Tool

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

1. **cd** to the *assemblytool\_installation* folder:

```
$ cd assemblytool_installation
```

2. Enter the following command:

```
$ ant initialize
```

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

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

```
engine.root = engine-installation
```

5. If you want the **escenic-admin** web application to be password protected, uncomment the **escenic-admin-authentication** property, and set it to **true**:

```
escenic-admin-authentication = true
```

You are strongly recommended to set this property when installing a production system. If you do so, then you must also define the required administrator username and password in Tomcat. For details of how to do this, see [section 4.4](#).

### 3.2.3 Initialize the Bootstrap Layer

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

1. **cd** to the *assemblytool\_installation* folder:

```
$ cd assemblytool_installation
```

2. Enter the following command:

```
$ ant -q ear
```

This command (which takes a few minutes to complete) 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 *assemblytool\_installation/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 3.3.6](#).

### 3.2.4 Edit the Common Configuration Layer

The Content Engine package installs a **skeleton configuration layer** in `/etc/escenic/engine/common`. This folder tree contains configuration files that you can use as a basis for your host configurations. All the property settings in these files are commented out, and you need to uncomment and set values in some of the files.

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

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

**databaseProductName**

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

- **MySQL**

**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/engine/
```

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.

**customerId**

The name of your organization, or some other name that will clearly identify your installation to Escenic support staff, plus information about the type of this Content Engine instance. This identifier is used for error and usage reporting. For production installations you **must** use the following format:

```
customerId=prod@customer-id
```

where *customer-id* is an identifier for your organization - your domain name for example:

```
customerId=prod@mycompany.com
```

You can optionally prefix the customer ID with a qualifier of some kind if you have multiple installations and separate reporting is considered useful. You might, for example, replace `prod@mycompany.com` with `prod@dailynews.mycompany.com` on some hosts and `prod@weeklomag.mycompany.com` on others.

For non-production instances, you can replace the `prod` instance type identifier with another identifier of your choice, such as `test` or `staging`.

Reporter components that are installed with the Content Engine automatically send reports about the Escenic software running at your site. The reported information includes:

- Which versions of the Content Engine and its plug-ins are running at your site.
- The number of active Content Studio users (reported once an hour)

This information is very valuable to Escenic support engineers and helps with the diagnosis of reported errors and performance problems. It is also used in the calculation of support and maintenance charges. For this reason you are **required** to set this property on all Content Engine installations and to configure your firewall so that reports can be delivered (see below).

The reporter components sends their reports to Google Analytics. In order for this to work, your firewalls must allow outgoing traffic on port 80 to **http://www.google-analytics.com**.

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

**dataConnector**

Must be set as follows:

```
dataConnector=/connector/DataConnector
```

### 3.2.5 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 **root**, create a folder called `/etc/escenic/engine/host`. Under this folder, 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 -p /etc/escenic/engine/host/{editorial1,presentation1}
```

If there are any configuration settings that you know need to be set differently on a specific host, you can set them as follows:

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

For further information about the configuration layers, see [Configuring The Content Engine](#).

### 3.2.6 Configure the ece Script

To configure the `ece` shell script:

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

- Create an empty file in the `/etc/escenic/` folder called `ece-engine1.conf`:

```
# touch /etc/escenic/ece-engine1.conf
```

- Open `/etc/escenic/ece-engine1.conf`, and copy into it the parameter settings listed below from `/etc/escenic/ece.conf`. Edit the copied parameter settings. 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 3.1.1](#)).

**ece\_home**

Set this to *engine-installation*.

**escenic.server**

Set this to the host name or IP address.

**ece\_security\_configuration\_dir**

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-engine1`.

**assemblytool\_home**

Set this to *assemblytool\_installation*.

- Create the standard installation directory structure. Enter the following commands:

```
# mkdir -p /var/{crash,lib,log,run,cache,spool}/escenic
# chown escenic:escenic /var/{crash,lib,log,run,cache,spool}/escenic -R
```

These commands will create the following folders and assign them to the **escenic** user:

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

If, for any reason, you do not want to install Content Engine files in standard locations and you have modified any of the following settings in the `ece.conf` file:

```
cache_dir
log_dir
pid_dir
heap_dump_dir
```

then you must make sure that all the referenced folders exist and that the **escenic** user has write permission to them. 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.

### 3.3 Old-style Installation

You should only use this method of installation if you have special requirements that you know make the new packaged-based method unsuitable.

### 3.3.1 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 Escenic representative or Escenic support ([support@escenic.com](mailto:support@escenic.com)).

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

Download the following items from <http://docs.escenic.com/ece-download-matrix/6.3/>:

- The Content Engine installation set **engine-6.3.2-1.zip**
- The Content Engine assembly tool **assemblytool-at-version.zip**
- The Content Engine documentation package **documentation-engine-6.3.2-1-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 3.1.5](#), 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 https://user:password@maven.escenic.com/com/escenic/engine/6.3.2-1/
engine-6.3.2-1.zip
$ wget https://user:password@maven.escenic.com/com/escenic/tools/assemblytool/at-
version/assemblytool-at-version.zip
$ wget https://user:password@maven.escenic.com/com/escenic/engine/documentation/
documentation-engine/6.3.2-1/documentation-engine-6.3.2-1-pdf.zip
```

where:

- *user* and *password* are the user name(s) and password(s) you have received from Escenic
- *at-version* is the number of the latest version of the assembly tool

### 3.3.2 Create Configuration Folder

On all your **engine-hosts**, while logged in as **root**, create an **escenic/engine** folder under **/etc**:

```
# mkdir -p /etc/escenic/engine
```

If you are installing on more than one machine, then you need to keep the contents of the **/etc/escenic** folder synchronized on all your hosts. It is recommended that you use the **rsync** utility for this purpose.

### 3.3.3 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 3.1.5](#), then you should have direct access to the Content Engine packages you downloaded earlier (see [section 3.3.1](#)). 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/
$ unzip /mnt/download/engine-6.3.2-1.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-6.3.2-1 engine
```

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

```
$ mkdir assemblytool-at-version
$ ln -s assemblytool-at-version assemblytool
$ cd assemblytool/
$ jar xf /mnt/download/assemblytool-at-version.zip
```

where *at-version* is the assembly tool version number.

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.

### 3.3.4 Initialize the Assembly Tool

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

1. `cd` to the `assemblytool_installation` folder:

```
$ cd assemblytool_installation
```

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 = engine-installation
```

5. If you want the **escenic-admin** web application to be password protected, uncomment the **escenic-admin-authentication** property, and set it to **true**:

```
escenic-admin-authentication = true
```

You are strongly recommended to set this property when installing a production system. If you do so, then you must also define the required administrator username and password in Tomcat. For details of how to do this, see [section 4.4](#).

### 3.3.5 Initialize the Bootstrap Layer

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

1. **cd** to the `assemblytool_installation` folder:

```
| $ cd assemblytool_installation
```

2. Enter the following command:

```
| $ ant -q ear
```

This command (which takes a few minutes to complete) 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 `assemblytool_installation/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 3.3.6](#).

### 3.3.6 Create the Common Configuration Layer

The Content Engine package installed in the [section 3.3.3](#) 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 `engine-installation/siteconfig/config-skeleton`.

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

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

```
| # cp -r engine-installation/siteconfig/config-skeleton/* /etc/escenic/engine/
| common/
```

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

```
| # cp -r engine-installation/security/ /etc/escenic/engine/common/
```

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

#### **databaseProductName**

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

- **MySQL**

#### **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.

**customerId**

The name of your organization, or some other name that will clearly identify your installation to Escenic support staff, plus information about the type of this Content Engine instance. This identifier is used for error and usage reporting. For production installations you **must** use the following format:

```
customerId=prod@customer-id
```

where *customer-id* is an identifier for your organization - your domain name for example:

```
customerId=prod@mycompany.com
```

You can optionally prefix the customer ID with a qualifier of some kind if you have multiple installations and separate reporting is considered useful. You might, for example, replace **prod@mycompany.com** with **prod@dailynews.mycompany.com** on some hosts and **prod@weeklomag.mycompany.com** on others.

For non-production instances, you can replace the **prod** instance type identifier with another identifier of your choice, such as **test** or **staging**.

Reporter components that are installed with the Content Engine automatically send reports about the Escenic software running at your site. The reported information includes:

- Which versions of the Content Engine and its plug-ins are running at your site.
- The number of active Content Studio users (reported once an hour)

This information is very valuable to Escenic support engineers and helps with the diagnosis of reported errors and performance problems. It is also used in the calculation of support and maintenance charges. For this reason you are **required** to set this property on all Content Engine installations and to configure your firewall so that reports can be delivered (see below).

The reporter components sends their reports to Google Analytics. In order for this to work, your firewalls must allow outgoing traffic on port 80 to **http://www.google-analytics.com**.

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

**dataConnector**

Must be set as follows:

```
dataConnector=/connector/DataConnector
```

### 3.3.7 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 **root**, create a folder called `/etc/escenic/engine/host`. Under this folder, 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 -p /etc/escenic/engine/host/{editorial1,presentation1}
```

If there are any configuration settings that you know need to be set differently on a specific host, you can set by:

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 [Configuring The Content Engine](#).

### 3.3.8 Install the ece Script

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

To install the **ece** shell script:

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

1. Make **/opt/escenic/engine/ece-scripts/usr/bin/ece** executable:
 

```
| $ chmod +x /opt/escenic/engine/ece-scripts/usr/bin/ece
```
2. Open **/home/escenic/.bashrc** for editing, and add the following line to it in order to add **/opt/escenic/engine/ece-scripts/usr/bin** to your **PATH**:
 

```
| export PATH=/opt/escenic/engine/ece-scripts/usr/bin:$PATH
```
3. Then enter the following command to apply the change you have made in your current shell:
 

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

### 3.3.9 Configure the ece Script

To configure the **ece** shell script:

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

1. Copy the script's configuration file to **/etc/escenic**:
 

```
| # cp engine-installation/ece-scripts/etc/escenic/ece.conf /etc/escenic
```
2. Create an empty file in the **/etc/escenic/** folder called **ece-engine1.conf**:
 

```
| # touch /etc/escenic/ece-engine1.conf
```
3. Open **/etc/escenic/ece-engine1.conf**, and copy into it the parameter settings listed below from **/etc/escenic/ece.conf**. Edit the copied parameter settings. 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 3.1.1](#)).

**ece\_home**

Set this to *engine-installation*.

**escenic.server**

Set this to the host name or IP address.

**ece\_security\_configuration\_dir**

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-engine1`.

**assemblytool\_home**

Set this to *assemblytool\_installation*.

4. Create the standard installation folder structure. Enter the following commands:

```
# mkdir -p /var/{crash,lib,log,run,cache,spool}/escenic
# chown escenic:escenic /var/{crash,lib,log,run,cache,spool}/escenic -R
```

These commands will create the following folders and assign them to the **escenic** user:

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

If, for any reason, you do not want to install Content Engine files in standard locations and you have modified any of the following settings in the `ece.conf` file:

```
cache_dir
log_dir
pid_dir
heap_dump_dir
```

then you must make sure that all the referenced folders exist and that the **escenic** user has write permission to them. 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.

### 3.3.10 Copy the ece Script

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/ece-scripts/usr/bin/
ece /usr/bin/
```

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

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

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

3. Create the standard installation folder structure in the same way as you did on the **assembly host** (see step 4 in [section 3.3.9](#)).
4. Verify that the script is correctly installed by entering the following while logged in as **escenic**:

```
| $ ece help
```

This should give help output describing the usage of the **ece** command.

## 3.4 Finishing Up

Whichever installation method you have used, you may now need to assemble and deploy an EAR file containing the supplied Escenic web applications, check that everything is installed correctly and so on. These tasks are described in the following sections.

### 3.4.1 Assemble and Deploy

You **may** now need to assemble and deploy your web applications. You need to do this in the following cases:

- You installed the Content Engine using the old-style installation method.
- You installed the Content Engine using the package-based installation method **and** your installation includes one or more old-style JSP-based Escenic publications.

If you installed the Content Engine using the package-based installation method and **all** your publications are based on DPRES, the new decoupled presentation layer, then you do not need to 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
```

### 3.4.2 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
```

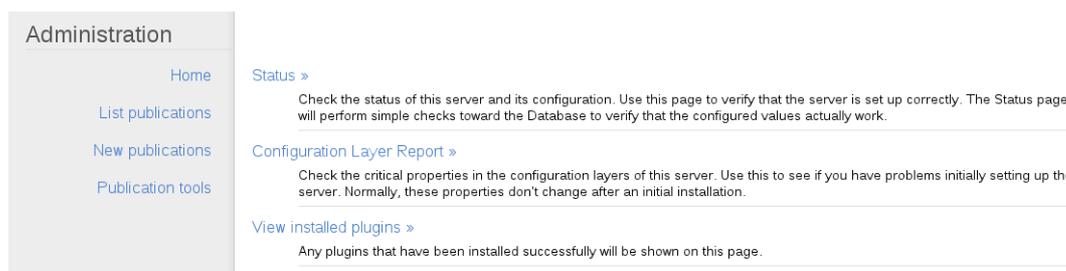
This should produce something like the following output:

```
| [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.



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.

### 3.4.3 Install a Daemon Script

At this point you have successfully installed the Content Engine and managed its instances using the **ece** script. Now you may want to install a daemon script that will automatically run the Content Engine on system start-up and shut it down on system shutdown. If you have the database server on the same machine then you have to make sure that it also starts automatically on system start-up.

The following instructions tell you how to install the daemon script.

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-scripts/etc/default/ece**) to a file called **ece** in the **/etc/default** folder. On the **assembly-host** you can do this as follows:

```
# cp engine-installation/ece-scripts/etc/init.d/ece /etc/init.d/
# cp engine-installation/ece-scripts/etc/default/ece /etc/default/
```

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:engine-installation/ece-scripts/etc/init.d/
ece /etc/init.d/
# scp escenic@assembly-host-ip-address:engine-installation/ece-scripts/etc/
default/ece /etc/default/
```

2. Make the daemon script executable:
 

```
# chmod +x /etc/init.d/ece
```
3. The next step is to run the daemon script, **ece**. But before that you may need to modify the settings for the daemon script. Open **/etc/default/ece** in an editor and modify the settings to meet your requirements.

If you have not installed the Content Engine files in standard locations, then you must modify the **dir\_list** setting according to the **/etc/escenic/ece.conf**, i.e. the directories you have created while installing the **ece** script. See [section 3.3.8](#).

Now, if you have started the Content Engine, stop it.

```
# ece stop
```

Then enter the following commands:

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

If everything is ok at this point, then you are ready to add this script as a daemon. Other wise you need to fix the problem and then try again.

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.

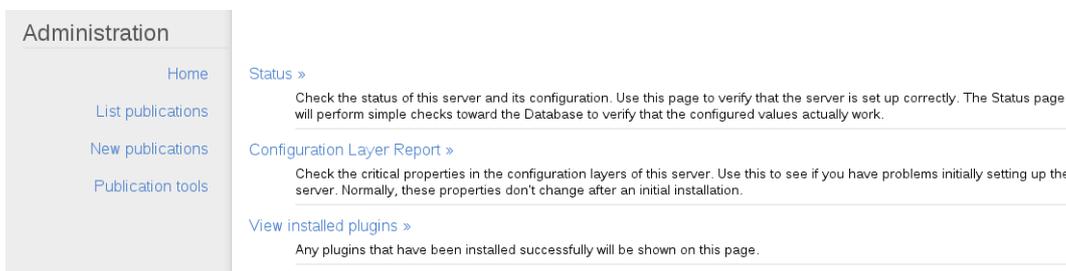
### 3.4.4 Create a Publication

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):

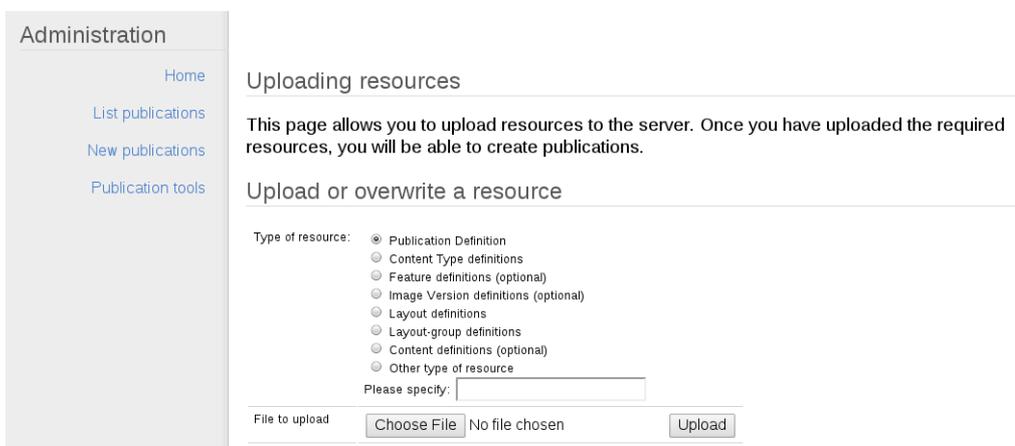
1. Download the publication WAR file from one of your **engine-hosts** to the machine on which you are working. You will find **demo-temp-dev.war** in the *engine-installation/contrib/wars* folder.
2. Open a browser, and access the **escenic-admin** web application on one of your **editorial-hosts**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

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

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



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



4. Click on the **Choose File** 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:

Administration

- [Home](#)
- [List publications](#)
- [New publications](#)
- [Publication tools](#)

### 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 publications.

### 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 (bytes)	Valid	
/escenic/content-type	00	application/vnd.escenic.content-type	META-INF/escenic/publication-resources/escenic/content-type	6332	Validated	<a href="#">Remove uploaded resource</a>
/escenic/feature	00	text/plain	META-INF/escenic/publication-resources/escenic/feature	380	Validated	<a href="#">Remove uploaded resource</a>

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

Administration

- [Home](#)
- [List publications](#)
- [New publications](#)
- [Publication tools](#)

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:

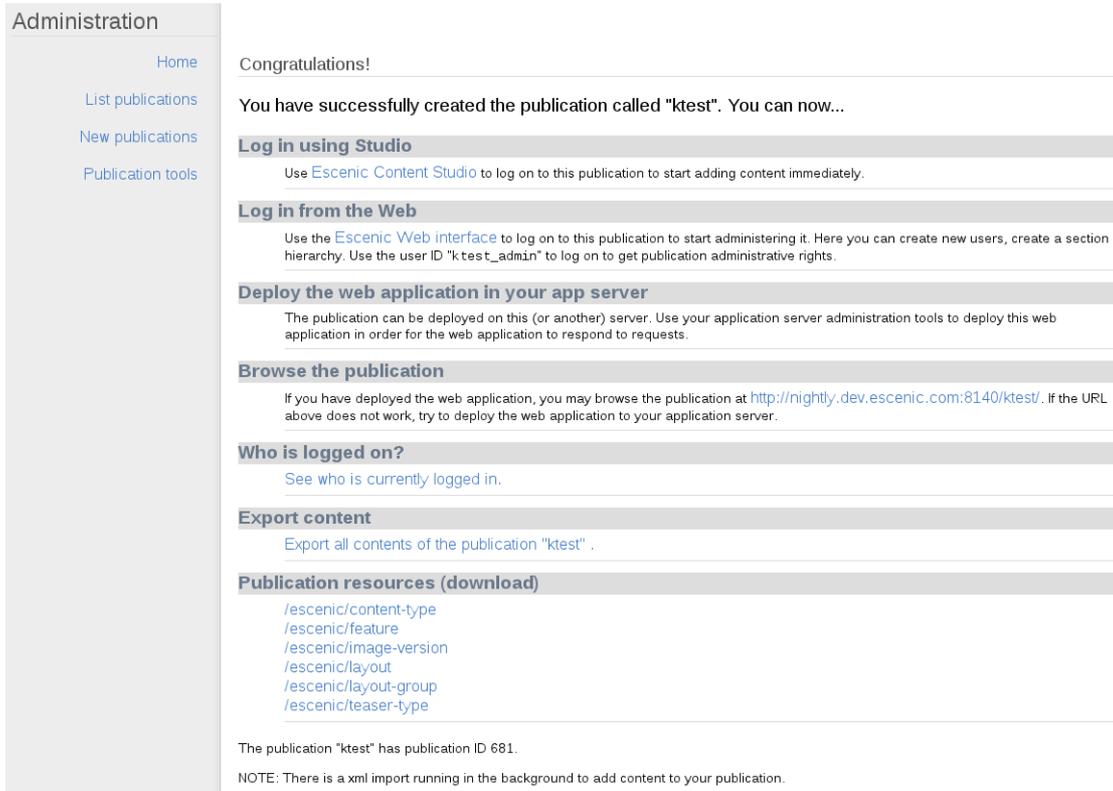
Verify password:

To access your publication after it has been created, an administrator user is created. To protect access to the publication, you can specify the administrator's password. Note: Take care of the administrator password.

Creating a new publication

**A Publication** is the Escenic representation of one complete web-site. Many installations only have one publication, while other installations have several publications. To decide whether or not your installation requires several publications, an easy rule of thumb is **one publication per web site**. Say Escenic is being used to power a company internet; you might only need one publication. If Escenic is used to power several company intranets, it is logical to create a Publication for each such web site. The same rule could apply to an installation with a Company Intranet and Extranet, where the two sites would primarily contain different data.

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



The screenshot shows the 'Administration' sidebar on the left with links for Home, List publications, New publications, and Publication tools. The main content area displays a 'Congratulations!' message: 'You have successfully created the publication called "ktest". You can now...'. Below this are several sections: 'Log in using Studio' (with a link to Escenic Content Studio), 'Log in from the Web' (with a link to the Escenic Web interface and user ID 'ktest\_admin'), 'Deploy the web application in your app server', 'Browse the publication' (with a URL: http://nightly.dev.escenic.com:8140/ktest/), 'Who is logged on?' (with a link to see who is logged in), 'Export content' (with a link to export all contents), and 'Publication resources (download)' (with links to various resource types like /escenic/content-type, /escenic/feature, etc.). At the bottom, it notes the publication ID is 681 and mentions an XML import running in the background.

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 either:

- You deploy the **demo-temp-dev** web application on the application server. This will create a viewable publication using the Content Engine's built-in presentation layer. For details of how to deploy the publication in this way, see [section 3.4.4.1](#).
- You create a viewing application using the new Escenic decoupled presentation layer, [DPRES](#). In this case you do **not** need to deploy the publication as described in [section 3.4.4.1](#).

### 3.4.4.1 Deploy The Publication

To deploy **demo-temp-dev**, do the following on your **assembly-host**, while logged in as **escenic**:

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

```
$ cp engine-installation/contrib/wars/demo-temp-dev.war assemblytool_installation/publications/
```

- Create a text file called **demo-temp-dev.properties** in the **assemblytool\_installation/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.

## 3.4.5 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.

3. The following page should then be displayed:



### 3.4.6 Test Content Studio

To check that Content Studio is correctly installed:

1. 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.
2. 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.

3. Click on the **Launch Escenic Content Studio** button. After the application has downloaded, a login dialog should be displayed.
4. Enter the **demo-temp-dev** publication's administrator user name (*publication-name\_admin*) and the administrator password you specified when creating the publication.
5. 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-engine1/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 [this blog entry](#).

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**

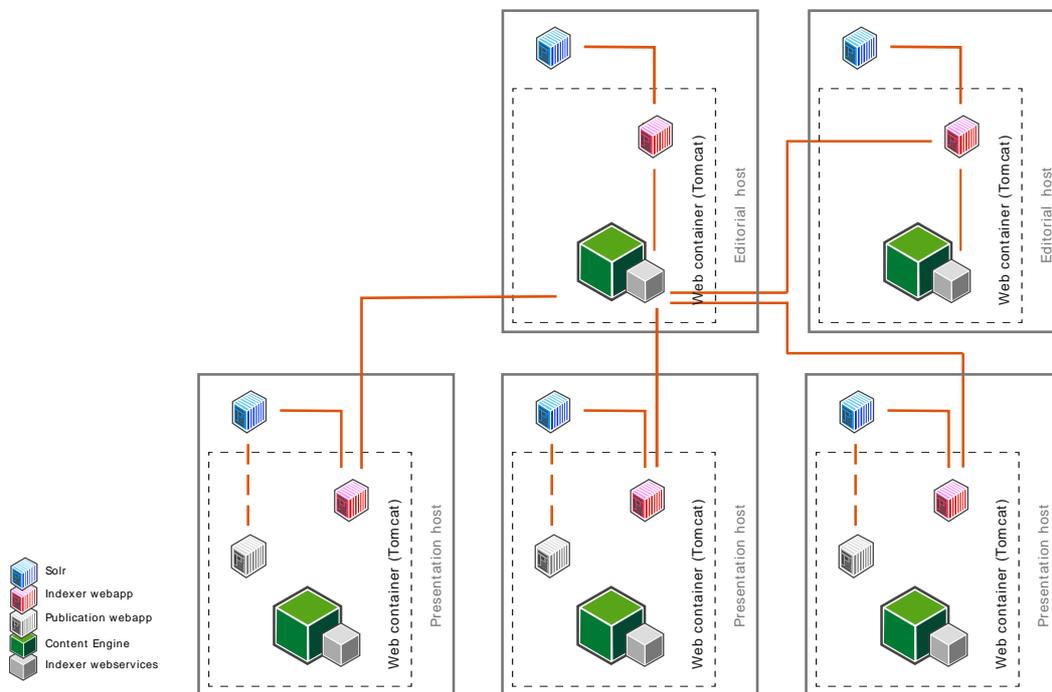
The keystore/certificate password you specified when adding the self-signed certificate. If you did not specify a password, the default is **changeit**.

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 3.3.6](#)). 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** and the **indexer** 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. They should therefore not be run in the same JVM as the Content Engine on production systems. **solr** already runs in its own webapp container (and therefore in a different JVM), but the **indexer** is deployed to the same Tomcat instance as the Content Engine. The simplest way to achieve this separation on a single-host installation is to move the **indexer** webapp to a separate Tomcat instance. For more about this, see [Isolating The Search Engine](#).

### Solr stemming

In the default **solr** configuration, English stemming is enabled by default. This means that searching non-English content might give unexpected results.

If your content is in a language other than English, you should either disable stemming or modify the configuration to suit your language.

To disable stemming, remove the following line from **schema.xml**:

```
<filter class="solr.SnowballPorterFilterFactory" protected="protowords.txt"
  language="English"/>
```

Disabling stemming will improve **solr**'s performance.

For information about how to configure stemming for other languages, see the Solr documentation on <http://lucene.apache.org/solr/>.

### 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/solr-core/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 [Search Engine Configuration and Management](#).

## 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 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 **root** on one of your **engine hosts**.
2. Create a folder for the new properties file:
 

```
| # mkdir -p /etc/escenic/engine/common/com/whalin/
```
3. Create a file called **SocketIOPool.properties** in the new folder, open it for editing and add the following content:

```
| $class=com.whalin.MemCached.SocketIOPool
| # 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/whalin/SocketIOPool
```

### 4.3.3 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).

It is also possible to configure this service so that hitting **flushCache** also flushes the **memcached** cache. This is not recommended!

```
flushMemCached=true
```

With this set, hitting **flushCache** on any of the servers connected to the same **memcached** cluster will flush the complete **memcached** cache!

## 4.4 Password Protect escenic-admin

[escenic-admin](#) is a web application for administration of the Content Engine. In a default installation it is installed without password protection. **escenic-admin** is, however, a powerful tool and should definitely not be openly accessible in a production environment. You are therefore strongly recommended to limit access to a small group of competent users.

To password protect **escenic-admin**:

1. Set **escenic-admin-authentication = true** in the **assemblytool** properties file before running the assemblytool (see [section 3.3.4](#)).
2. Define one or more user log-ins for **escenic-admin** in Tomcat's **tomcat-users.xml** file.

To define an **escenic-admin** user in Tomcat, you need to add an entry like this to **tomcat-users.xml**:

```
<user username="admin-user" password="admin-password" roles="ECEAdmin"/>
```

where:

- *admin-user* is the user name you want to use for logging in to **escenic-admin**
- *admin-password* is the password you want to use for logging in to **escenic-admin**

It is obviously **not** a good idea to enter the password in this file as plain text. For a description of how to encrypt the password, see <http://www.jdev.it/encrypting-passwords-in-tomcat/>.

## 5 Upgrading

How you upgrade the Content Engine depends upon how you installed it in the first place. If you installed it using the new package-based method as described in [section 3.2](#), then upgrading is very easy. If you installed it using the old method described in [section 3.3](#), then it requires a bit more work.

### 5.1 Package-based Upgrade

You can only use this upgrade method if you have previously installed the Content Engine as described in [section 3.2](#). Otherwise you need to follow the method described in [section 5.2](#).

All you need to do to upgrade your Content Engine (including any plug-ins that you have installed from **apt** / **rpm** packages) is:

1. Read the release notes for your planned upgrades (both Content Engine and plug-ins). Make a note of any special tasks that need to be carried out in connection with the upgrades. Some changes, for example, require you to run a database upgrade script after the upgrade.
2. Log in as **root** on your **assembly-host** and on each of your **engine-hosts**, and enter the following commands:

```
# apt-get update
# apt-get upgrade
```
3. Carry out any required upgrade tasks. For general information about how to carry out a database upgrade should that be necessary, see: [section 5.3](#).

### 5.2 Old-style Upgrade

You should only use this upgrade method if you have previously installed the Content Engine as described in [section 3.3](#). Otherwise you need to follow the method described in [section 5.1](#).

To upgrade the Content Engine:

1. Read the release notes for your planned upgrade. Make a note of any special tasks that need to be carried out in connection with the upgrade. Some changes, for example, require you to run a database upgrade script after the upgrade.
2. Download the Content Engine distribution file and the latest version of the assembly tool to your **assembly-host** as described in [section 3.3.1](#).
3. Unpack the downloaded files as described in [section 3.3.3](#).
4. Edit your assembly tool's **assemble.properties** file, setting **engine.root** to point to the new location, or If **engine.root** points to a symbolic link, update the symbolic link instead.
5. Assemble and deploy your EAR file as described in [section 3.4.1](#).
6. Restart your application server(s).
7. Carry out any required upgrade tasks. For general information about how to carry out a database upgrade should that be necessary, see: [section 5.3](#).

## 5.3 Database Upgrades

If the upgrade tasks for the upgrade you are carrying out include upgrading the database, then a script for performing the upgrade will be included in the distribution. All you need to do is log in as the **escenic** user on your **database-host** and run the appropriate script for your particular database. The scripts are located in the *engine-installation/contrib/database/sql/* folder. Here you will find a subfolder for each supported database (MySQL, Oracle and so on). In each subfolder is an **upgrade** folder and a script for upgrading to the current version (plus scripts for upgrading to a number of previous versions. To upgrade a MySQL database, for example, you would need to run:

```
| engine-installation/contrib/database/sql/upgrade/to-version
```

where *version* is the current Content Engine version number.