Red Hat Fuse 7.0

Security Guide

Making it safe for your systems to work together

Last Updated: 2018-08-27

Red Hat Fuse 7.0 Security Guide

Making it safe for your systems to work together

Legal Notice

Abstract

This guide describes how to secure the Red Hat Fuse container, the web console, message brokers, routing and integration components, web and RESTful services, and it provides a tutorial on LDAP authentication.

Table of Contents

CHAPTER 1. SECURITY ARCHITECTURE

Abstract

In the OSGi container, it is possible to deploy applications supporting a variety of security features. Currently, only the Java Authentication and Authorization Service (JAAS) is based on a common, container-wide infrastructure. Other security features are provided separately by the individual products and components deployed in the container.

1.1. OSGI CONTAINER SECURITY

Overview

Figure 1.1, “OSGi Container Security Architecture” shows an overview of the security infrastructure that is used across the container and is accessible to all bundles deployed in the container. This common security infrastructure currently consists of a mechanism for making JAAS realms (or login modules) available to all application bundles.

Figure 1.1. OSGi Container Security Architecture

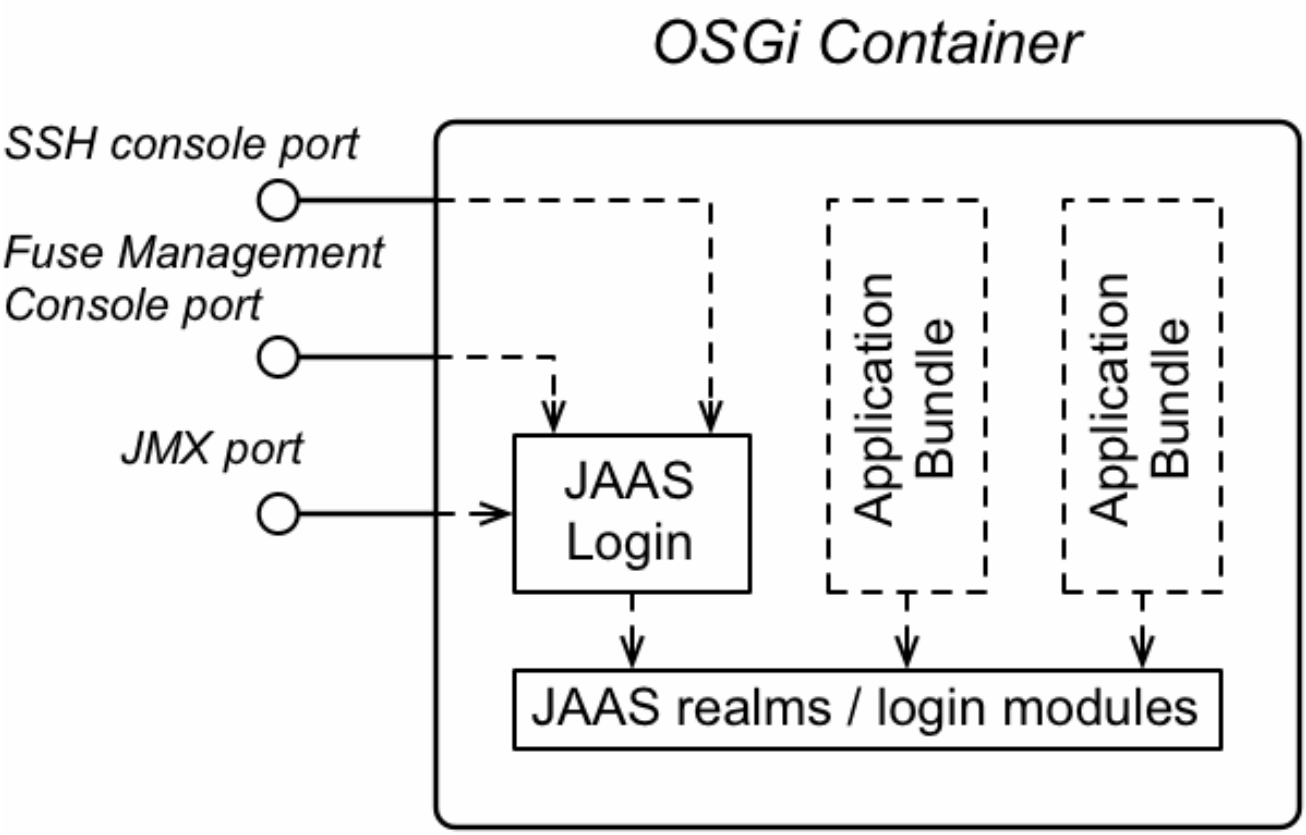

JAAS realms

A JAAS realm or login module is a plug-in module that provides authentication and authorization data to Java applications, as defined by the Java Authentication and Authorization Service (JAAS) specification.

Red Hat Fuse supports a special mechanism for defining JAAS login modules (in either a Spring or a blueprint file), which makes the login module accessible to all bundles in the container. This makes it easy for multiple applications running in the OSGi container to consolidate their security data into a single JAAS realm.

karaf realm

The OSGi container has a predefined JAAS realm, the **karaf** realm. Red Hat Fuse uses the **karaf** realm to provide authentication for remote administration of the OSGi runtime, for the Fuse Management Console, and for JMX management. The **karaf** realm uses a simple file-based repository, where authentication data is stored in the ***InstallDir*/etc/users.properties** file.

You can use the **karaf** realm in your own applications. Simply configure **karaf** as the name of the JAAS realm that you want to use. Your application then performs authentication using the data from the **users.properties** file.

Console port

You can administer the OSGi container remotely either by connecting to the console port with a Karaf client or using the Karaf **ssh:ssh** command. The console port is secured by a JAAS login feature that connects to the **karaf** realm. Users that try to connect to the console port will be prompted to enter a username and password that must match one of the accounts from the **karaf** realm.

JMX port

You can manage the OSGi container by connecting to the JMX port (for example, using Java's JConsole). The JMX port is also secured by a JAAS login feature that connects to the **karaf** realm.

Application bundles and JAAS security

Any application bundles that you deploy into the OSGi container can access the container's JAAS realms. The application bundle simply references one of the existing JAAS realms by name (which corresponds to an instance of a JAAS login module).

It is essential, however, that the JAAS realms are defined using the OSGi container's own login configuration mechanism—by default, Java provides a simple file-based login configuration implementation, but you **cannot** use this implementation in the context of the OSGi container.

1.2. APACHE CAMEL SECURITY

Overview

Figure 1.2, "Apache Camel Security Architecture" shows an overview of the basic options for securely routing messages between Apache Camel endpoints.

Figure 1.2. Apache Camel Security Architecture

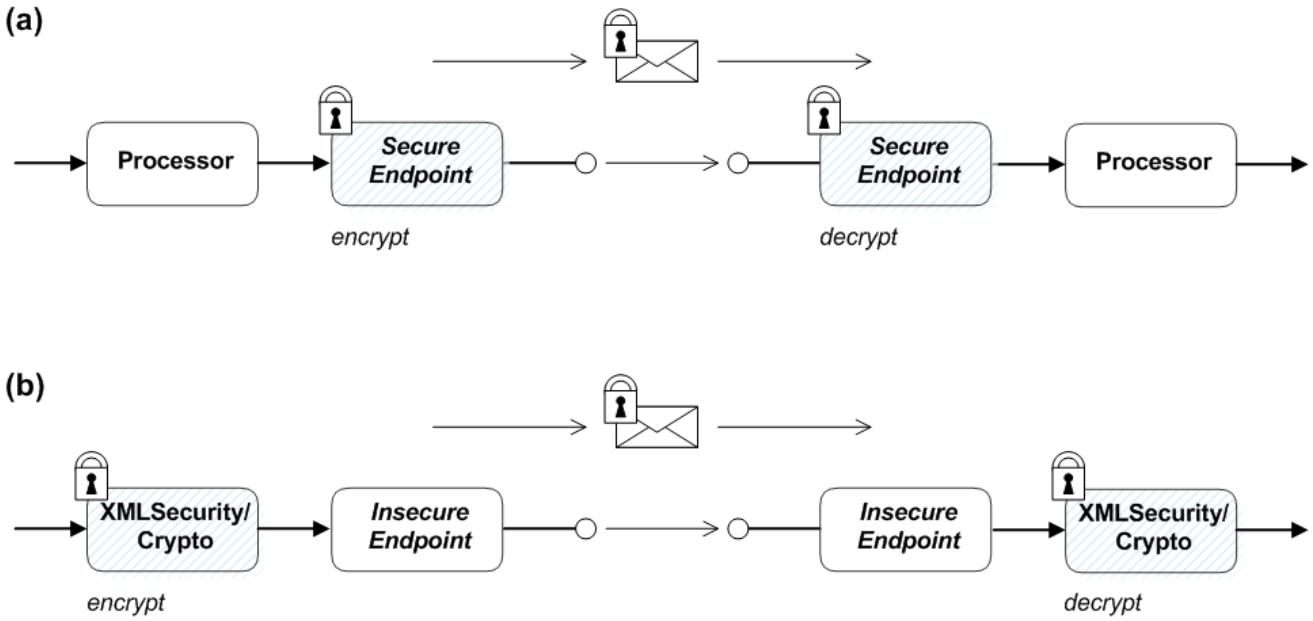

Alternatives for Apache Camel security

As shown in Figure 1.2, “Apache Camel Security Architecture”, you have the following options for securing messages:

- **Endpoint security**—part (a) shows a message sent between two routes with secure endpoints. The producer endpoint on the left opens a secure connection (typically using SSL/TLS) to the consumer endpoint on the right. Both of the endpoints support security in this scenario.
  With endpoint security, it is typically possible to perform some form of peer authentication (and sometimes authorization).

- **Payload security**—part (b) shows a message sent between two routes where the endpoints are both **insecure**. To protect the message from unauthorized snooping in this case, use a *payload processor* that encrypts the message before sending and decrypts the message after it is received.
  A limitation of payload security is that it does **not** provide any kind of authentication or authorization mechanisms.

Endpoint security

There are several Camel components that support security features. It is important to note, however, that these security features are implemented by the individual components, **not** by the Camel core. Hence, the kinds of security feature that are supported, and the details of their implementation, vary from component to component. Some of the Camel components that currently support security are, as follows:

- JMS and ActiveMQ—SSL/TLS security and JAAS security for client-to-broker and broker-to-broker communication.

- Jetty—HTTP Basic Authentication and SSL/TLS security.

- CXF—SSL/TLS security and WS-Security.

- Crypto—creates and verifies digital signatures in order to guarantee message integrity.

- Netty—SSL/TLS security.

- MINA—SSL/TLS security.
- Cometd—SSL/TLS security.
- glogin and gauth—authorization in the context of Google applications.

Payload security

Apache Camel provides the following payload security implementations, where the encryption and decryption steps are exposed as data formats on the **`marshal()`** and **`unmarshal()`** operations

- the section called “XMLSecurity data format”.
- the section called “Crypto data format”.

XMLSecurity data format

The XMLSecurity data format is specifically designed to encrypt XML payloads. When using this data format, you can specify which XML element to encrypt. The default behavior is to encrypt **all** XML elements. This feature uses a symmetric encryption algorithm.

For more details, see http://camel.apache.org/xmlsecurity-dataformat.html.

Crypto data format

The crypto data format is a general purpose encryption feature that can encrypt any kind of payload. It is based on the Java Cryptographic Extension and implements only symmetric (shared-key) encryption and decryption.

For more details, see http://camel.apache.org/crypto.html.

CHAPTER 2. SECURING THE APACHE KARAF CONTAINER

Abstract

The Apache Karaf container is secured using JAAS. By defining JAAS realms, you can configure the mechanism used to retrieve user credentials. You can also refine access to the container's administrative interfaces by changing the default roles.

2.1. JAAS AUTHENTICATION

Abstract

The Java Authentication and Authorization Service (JAAS) provides a general framework for implementing authentication in a Java application. The implementation of authentication is modular, with individual JAAS modules (or plug-ins) providing the authentication implementations.

For background information about JAAS, see the JAAS Reference Guide.

2.1.1. Default JAAS Realm

This section describes how to manage user data for the default JAAS realm in a Karaf container.

Default JAAS realm

The Karaf container has a predefined JAAS realm, the **`karaf`** realm, which is used by default to secure all aspects of the container.

How to integrate an application with JAAS

You can use the **`karaf`** realm in your own applications. Simply configure **`karaf`** as the name of the JAAS realm that you want to use.

Default JAAS login modules

When you start the Karaf container for the first time, it is configured to use the **`karaf`** default realm. In this default configuration, the **`karaf`** realm deploys five JAAS login modules, which are enabled simultaneously. To see the deployed login modules, enter the **`jaas:realms`** console command, as follows:

```
Index | Realm Name | Login Module Class Name
──────┼────────────┼──────────────────────────────────────────────────────────
──────────
1     | karaf      |
org.apache.karaf.jaas.modules.properties.PropertiesLoginModule
2     | karaf      |
org.apache.karaf.jaas.modules.publickey.PublickeyLoginModule
3     | karaf      |
org.apache.karaf.jaas.modules.audit.FileAuditLoginModule
4     | karaf      |
org.apache.karaf.jaas.modules.audit.LogAuditLoginModule
5     | karaf      |
org.apache.karaf.jaas.modules.audit.EventAdminAuditLoginModule
```

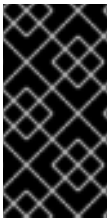

IMPORTANT

In a Karaf container, **both** the properties login module and the public key login module are enabled. When JAAS authenticates a user, it tries first of all to authenticate the user with the properties login module. If that fails, it then tries to authenticate the user with the public key login module. If that module also fails, an error is raised.

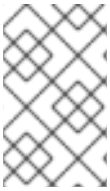

NOTE

The **`FileAuditLoginModule`** login module, the **`LogAuditLoginModule`** login module, and the **`EventAdminAuditLoginModule`** login module are used to record an audit trail of successful and failed login attempts. These login modules do **not** authenticate users.

Configuring users in the properties login module

The properties login module is used to store username/password credentials in a flat file format. To create a new user in the properties login module, open the ***`InstallDir`*`/etc/users.properties`** file using a text editor and add a line with the following syntax:

```
Username=Password[,UserGroup|Role][,UserGroup|Role]...
```

For example, to create the **`jdoe`** user with password, **`topsecret`**, and role, **`admin`**, you could create an entry like the following:

```
jdoe=topsecret,admin
```

Where the **`admin`** role gives full administrative privileges to the **`jdoe`** user.

Configuring user groups in the properties login module

Instead of (or in addition to) assigning roles directly to users, you also have the option of adding users to *user groups* in the properties login module. To create a user group in the properties login module, open the ***`InstallDir`*`/etc/users.properties`** file using a text editor and add a line with the following syntax:

```
_g_\:GroupName=Role1,Role2,...
```

For example, to create the **`admingroup`** user group with the roles, **`group`** and **`admin`**, you could create an entry like the following:

```
_g_\:admingroup=group,admin
```

You could then add the **`majorclanger`** user to the **`admingroup`**, by creating the following user entry:

```
majorclanger=secretpass,_g_:admingroup
```

Configuring the public key login module

The public key login module is used to store SSH public key credentials in a flat file format. To create a new user in the public key login module, open the ***`InstallDir`*`/etc/keys.properties`** file using a text editor and add a line with the following syntax:

```
Username=PublicKey[,UserGroup|Role][,UserGroup|Role]...
```

For example, you can create the **jdoe** user with the **admin** role by adding the following entry to the ***InstallDir*/etc/keys.properties** file (on a single line):

```
jdoe=AAAAB3NzaC1kc3MAAACBAP1/U4EddRIpUt9KnC7s5Of2EbdSPO9EAMMeP4C2USZpRV1AI
lH7WT2NWPq/xfW6MPbLm1Vs14E7gB00b/JmYLdrmVClpJ+f6AR7ECLCT7up1/63xhv4O1fnfqi
mFQ8E+4P208UewwI1VBNaFpEy9nXzrith1yrv8iIDGZ3RSAHHAAAAFQCXYFCPFSMLzLKSuYKi6
4QL8Fgc9QAAAnEA9+GghdabPd7LvKtcNrhXuXmUr7v6OuqC+VdMCz0HgmdRWVeOutRZT+ZxBxC
BgLRJFnEj6EwoFhO3zwkyjMim4TwWeotifI0o4KOuHiuzpnWRbqN/C/ohNWLx+2J6ASQ7zKTxv
qhRkImog9/hWuWfBpKLZl6Ae1UlZAFMO/7PSSoAAACBAKKSU2PFl/qOLxIwmBZPPIcJshVe7bV
UpFvyl3BbJDow8rXfskl8wO63OzP/qLmcJM0+JbcRU/53Jj7uyk31drV2qxhIOsLDC9dGCWj47
Y7TyhPdXh/0dthTRBy6bqGtRPxGa7gJov1xm/UuYYXPIUR/3x9MAZvZ5xvE0kYXO+rx,admin
```

IMPORTANT

Do not insert the entire contents of an **id_rsa.pub** file here. Insert just the block of symbols which represents the public key itself.

Configuring user groups in the public key login module

Instead of (or in addition to) assigning roles directly to users, you also have the option of adding users to *user groups* in the public key login module. To create a user group in the public key login module, open the ***InstallDir*/etc/keys.properties** file using a text editor and add a line with the following syntax:

```
_g_\:GroupName=Role1,Role2,...
```

For example, to create the **admingroup** user group with the roles, **group** and **admin**, you could create an entry like the following:

```
_g_\:admingroup=group,admin
```

You could then add the **jdoe** user to the **admingroup**, by creating the following user entry:

```
jdoe=AAAAB3NzaC1kc3MAAACBAP1/U4EddRIpUt9KnC7s5Of2EbdSPO9EAMMeP4C2USZpRV1AI
lH7WT2NWPq/xfW6MPbLm1Vs14E7gB00b/JmYLdrmVClpJ+f6AR7ECLCT7up1/63xhv4O1fnfqi
mFQ8E+4P208UewwI1VBNaFpEy9nXzrith1yrv8iIDGZ3RSAHHAAAAFQCXYFCPFSMLzLKSuYKi6
4QL8Fgc9QAAAnEA9+GghdabPd7LvKtcNrhXuXmUr7v6OuqC+VdMCz0HgmdRWVeOutRZT+ZxBxC
BgLRJFnEj6EwoFhO3zwkyjMim4TwWeotifI0o4KOuHiuzpnWRbqN/C/ohNWLx+2J6ASQ7zKTxv
qhRkImog9/hWuWfBpKLZl6Ae1UlZAFMO/7PSSoAAACBAKKSU2PFl/qOLxIwmBZPPIcJshVe7bV
UpFvyl3BbJDow8rXfskl8wO63OzP/qLmcJM0+JbcRU/53Jj7uyk31drV2qxhIOsLDC9dGCWj47
Y7TyhPdXh/0dthTRBy6bqGtRPxGa7gJov1xm/UuYYXPIUR/3x9MAZvZ5xvE0kYXO+rx,_g_:ad
mingroup
```

Encrypting the stored passwords

By default, passwords are stored in the ***InstallDir*/etc/users.properties** file in plaintext format. To protect the passwords in this file, you must set the file permissions of the **users.properties** file so that it can be read only by administrators. To provide additional protection, you can optionally encrypt the stored passwords using a message digest algorithm.

To enable the password encryption feature, edit the ***InstallDir*/etc/org.apache.karaf.jaas.cfg** file and set the encryption properties as described in the comments. For example, the following settings would enable basic encryption using the MD5 message digest algorithm:

```
encryption.enabled = true
encryption.name = basic
encryption.prefix = {CRYPT}
encryption.suffix = {CRYPT}
encryption.algorithm = MD5
encryption.encoding = hexadecimal
```

NOTE

The encryption settings in the **org.apache.karaf.jaas.cfg** file are applied **only** to the default **karaf** realm in a Karaf container. They have no effect on a custom realm.

For more details about password encryption, see Section 2.1.8, “Encrypting Stored Passwords”.

Overriding the default realm

If you want to customise the JAAS realm, the most convenient approach to take is to override the default **karaf** realm by defining a higher ranking **karaf** realm. This ensures that all of the Red Hat Fuse security components switch to use your custom realm. For details of how to define and deploy custom JAAS realms, see Section 2.1.2, “Defining JAAS Realms”.

2.1.2. Defining JAAS Realms

When defining a JAAS realm in the OSGi container, you *cannot* put the definitions in a conventional JAAS login configuration file. Instead, the OSGi container uses a special **jaas:config** element for defining JAAS realms in a blueprint configuration file. The JAAS realms defined in this way are made available to *all* of the application bundles deployed in the container, making it possible to share the JAAS security infrastructure across the whole container.

Namespace

The **jaas:config** element is defined in the **http://karaf.apache.org/xmlns/jaas/v1.0.0** namespace. When defining a JAAS realm you need to include the line shown in Example 2.1, “JAAS Blueprint Namespace”.

Example 2.1. JAAS Blueprint Namespace

```
xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
```

Configuring a JAAS realm

The syntax for the **jaas:config** element is shown in Example 2.2, “Defining a JAAS Realm in Blueprint XML”.

Example 2.2. Defining a JAAS Realm in Blueprint XML

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
           xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0">

    <jaas:config name="JaasRealmName"
                 rank="IntegerRank">
        <jaas:module className="LoginModuleClassName"
                     flags="[required|requisite|sufficient|optional]">
            Property=Value
            ...
        </jaas:module>
        ...
        <!-- Can optionally define multiple modules -->
        ...
    </jaas:config>

</blueprint>
```

The elements are used as follows:

`jaas:config`

Defines the JAAS realm. It has the following attributes:

- **`name`**—specifies the name of the JAAS realm.
- **`rank`**—specifies an optional rank for resolving naming conflicts between JAAS realms . When two or more JAAS realms are registered under the same name, the OSGi container always picks the realm instance with the highest rank. If you decide to override the default realm, **`karaf`**, you should specify a **`rank`** of **`100`** or more, so that it overrides all of the previously installed **`karaf`** realms.

`jaas:module`

Defines a JAAS login module in the current realm. **`jaas:module`** has the following attributes:

- **`className`**—the fully-qualified class name of a JAAS login module. The specified class must be available from the bundle classloader.
- **`flags`**—determines what happens upon success or failure of the login operation. Table 2.1, “Flags for Defining a JAAS Module” describes the valid values.

Table 2.1. Flags for Defining a JAAS Module

Value	Description
`required`	Authentication of this login module must succeed. Always proceed to the next login module in this entry, irrespective of success or failure.
`requisite`	Authentication of this login module must succeed. If success, proceed to the next login module; if failure, return immediately without processing the remaining login modules.

Value	Description
sufficient	Authentication of this login module is not required to succeed. If success, return immediately without processing the remaining login modules; if failure, proceed to the next login module.
optional	Authentication of this login module is not required to succeed. Always proceed to the next login module in this entry, irrespective of success or failure.

The contents of a **jaas:module** element is a space separated list of property settings, which are used to initialize the JAAS login module instance. The specific properties are determined by the JAAS login module and must be put into the proper format.

NOTE

You can define multiple login modules in a realm.

Converting standard JAAS login properties to XML

Red Hat Fuse uses the same properties as a standard Java login configuration file, however Red Hat Fuse requires that they are specified slightly differently. To see how the Red Hat Fuse approach to defining JAAS realms compares with the standard Java login configuration file approach, consider how to convert the login configuration shown in Example 2.3, “Standard JAAS Properties”, which defines the **PropertiesLogin** realm using the Red Hat Fuse properties login module class, **PropertiesLoginModule**:

Example 2.3. Standard JAAS Properties

```
PropertiesLogin {
    org.apache.activemq.jaas.PropertiesLoginModule required
        org.apache.activemq.jaas.properties.user="users.properties"
        org.apache.activemq.jaas.properties.group="groups.properties";
};
```

The equivalent JAAS realm definition, using the **jaas:config** element in a blueprint file, is shown in Example 2.4, “Blueprint JAAS Properties”.

Example 2.4. Blueprint JAAS Properties

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0">
```

```
  <jaas:config name="PropertiesLogin">
    <jaas:module flags="required"
      className="org.apache.activemq.jaas.PropertiesLoginModule">
        org.apache.activemq.jaas.properties.user=users.properties
        org.apache.activemq.jaas.properties.group=groups.properties
    </jaas:module>
  </jaas:config>

</blueprint>
```

IMPORTANT

Do not use double quotes for JAAS properties in the blueprint configuration.

Example

Red Hat Fuse also provides an adapter that enables you to store JAAS authentication data in an X.500 server. Example 2.5, “Configuring a JAAS Realm” defines the **`LDAPLogin`** realm to use Red Hat Fuse’s **`LDAPLoginModule`** class, which connects to the LDAP server located at **ldap://localhost:10389**.

Example 2.5. Configuring a JAAS Realm

```
<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0">

  <jaas:config name="LDAPLogin" rank="200">
    <jaas:module flags="required"
      className="org.apache.karaf.jaas.modules.ldap.LDAPLoginModule">
        initialContextFactory=com.sun.jndi.ldap.LdapCtxFactory
        connection.username=uid=admin,ou=system
        connection.password=secret
        connection.protocol=
        connection.url = ldap://localhost:10389
        user.base.dn = ou=users,ou=system
        user.filter = (uid=%u)
        user.search.subtree = true
        role.base.dn = ou=users,ou=system
        role.filter = (uid=%u)
        role.name.attribute = ou
        role.search.subtree = true
        authentication = simple
    </jaas:module>
  </jaas:config>
</blueprint>
```

For a detailed description and example of using the LDAP login module, see Section 2.1.7, “JAAS LDAP Login Module”.

2.1.3. JAAS Properties Login Module

The JAAS properties login module stores user data in a flat file format (where the stored passwords can optionally be encrypted using a message digest algorithm). The user data can either be edited directly, using a simple text editor, or managed using the **`jaas:*`** console commands.

For example, a Karaf container uses the JAAS properties login module by default and stores the associated user data in the **`InstallDir/etc/users.properties`** file.

Supported credentials

The JAAS properties login module authenticates username/password credentials, returning the list of roles associated with the authenticated user.

Implementation classes

The following classes implement the JAAS properties login module:

`org.apache.karaf.jaas.modules.properties.PropertiesLoginModule`

Implements the JAAS login module.

`org.apache.karaf.jaas.modules.properties.PropertiesBackingEngineFactory`

Must be exposed as an OSGi service. This service makes it possible for you to manage the user data using the **`jaas:*`** console commands from the Apache Karaf shell (see Apache Karaf Console Reference).

Options

The JAAS properties login module supports the following options:

`users`

Location of the user properties file.

Format of the user properties file

The user properties file is used to store username, password, and role data for the properties login module. Each user is represented by a single line in the user properties file, where a line has the following form:

```
Username=Password[,UserGroup|Role][,UserGroup|Role]...
```

User groups can also be defined in this file, where each user group is represented by a single line in the following format:

```
_g_\:GroupName=Role1[,Role2]...
```

For example, you can define the users, **`bigcheese`** and **`guest`**, and the user groups, **`admingroup`** and **`guestgroup`**, as follows:

```
# Users
bigcheese=cheesepass,_g_:admingroup
guest=guestpass,_g_:guestgroup
```

```
# Groups
_g_\:admingroup=group,admin
_g_\:guestgroup=viewer
```

Sample Blueprint configuration

The following Blueprint configuration shows how to define a new **karaf** realm using the properties login module, where the default **karaf** realm is overridden by setting the **rank** attribute to **200**:

```
<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.1.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0">

  <type-converters>
    <bean
class="org.apache.karaf.jaas.modules.properties.PropertiesConverter"/>
  </type-converters>

<!--Allow usage of System properties, especially the karaf.base property--
>
  <ext:property-placeholder
       placeholder-prefix="$[" placeholder-suffix="]"/>

  <jaas:config name="karaf" rank="200">
    <jaas:module flags="required"
className="org.apache.karaf.jaas.modules.properties.PropertiesLoginModule"
>
        users= $[karaf.base]/etc/users.properties
    </jaas:module>
  </jaas:config>

  <!-- The Backing Engine Factory Service for the PropertiesLoginModule --
>
  <service interface="org.apache.karaf.jaas.modules.BackingEngineFactory">
    <bean
class="org.apache.karaf.jaas.modules.properties.PropertiesBackingEngineFac
tory"/>
  </service>

</blueprint>
```

Remember to export the **BackingEngineFactory** bean as an OSGi service, so that the **jaas:*** console commands can manage the user data.

2.1.4. JAAS OSGi Config Login Module

Overview

The JAAS OSGi config login modules leverages the **OSGi Config Admin Service** to store user data. This login module is fairly similar to the JAAS properties login module (for example, the syntax of the user entries is the same), but the mechanism for retrieving user data is based on the OSGi Config Admin

Service.

The user data can be edited directly by creating a corresponding OSGi configuration file, **etc/*PersistentID*.cfg** or using any method of configuration that is supported by the OSGi Config Admin Service. The **jaas:*** console commands are not supported, however.

Supported credentials

The JAAS OSGi config login module authenticates username/password credentials, returning the list of roles associated with the authenticated user.

Implementation classes

The following classes implement the JAAS OSGi config login module:

org.apache.karaf.jaas.modules.osgi.OsgiConfigLoginModule

Implements the JAAS login module.

NOTE

There is no backing engine factory for the OSGi config login module, which means that this module cannot be managed using the **jaas:*** console commands.

Options

The JAAS OSGi config login module supports the following options:

pid

The *persistent ID* of the OSGi configuration containing the user data. In the OSGi Config Admin standard, a persistent ID references a set of related configuration properties.

Location of the configuration file

The location of the configuration file follows the usual convention where the configuration for the persistent ID, ***PersistentID***, is stored in the following file:

```
InstallDir/etc/PersistentID.cfg
```

Format of the configuration file

The ***PersistentID*.cfg** configuration file is used to store username, password, and role data for the OSGi config login module. Each user is represented by a single line in the configuration file, where a line has the following form:

```
Username=Password[,Role][,Role]...
```

NOTE

User groups are **not** supported in the JAAS OSGi config login module.

Sample Blueprint configuration

The following Blueprint configuration shows how to define a new **karaf** realm using the OSGi config login module, where the default **karaf** realm is overridden by setting the **rank** attribute to **200**:

```
<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.1.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0">

  <jaas:config name="karaf" rank="200">
    <jaas:module flags="required"
className="org.apache.karaf.jaas.modules.osgi.OsgiConfigLoginModule">
        pid = org.jboss.example.osgiconfigloginmodule
    </jaas:module>
  </jaas:config>

</blueprint>
```

In this example, the user data will be stored in the file, ***InstallDir*/etc/org.jboss.example.osgiconfigloginmodule.cfg**, and it is not possible to edit the configuration using the **jaas:*** console commands.

2.1.5. JAAS Public Key Login Module

The JAAS public key login module stores user data in a flat file format, which can be edited directly using a simple text editor. The **jaas:*** console commands are not supported, however.

For example, a Karaf container uses the JAAS public key login module by default and stores the associated user data in the **InstallDir/etc/keys.properties** file.

Supported credentials

The JAAS public key login module authenticates SSH key credentials. When a user tries to log in, the SSH protocol uses the stored public key to challenge the user. The user must possess the corresponding private key in order to answer the challenge. If login is successful, the login module returns the list of roles associated with the user.

Implementation classes

The following classes implement the JAAS public key login module:

org.apache.karaf.jaas.modules.publickey.PublickeyLoginModule

Implements the JAAS login module.

NOTE

There is no backing engine factory for the public key login module, which means that this module cannot be managed using the **jaas:*** console commands.

Options

The JAAS public key login module supports the following options:

users

Location of the user properties file for the public key login module.

Format of the keys properties file

The **keys.properties** file is used to store username, public key, and role data for the public key login module. Each user is represented by a single line in the keys properties file, where a line has the following form:

```
Username=PublicKey[,UserGroup|Role][,UserGroup|Role]...
```

Where the *PublicKey* is the public key part of an SSH key pair (typically found in a user's home directory in **~/.ssh/id_rsa.pub** in a UNIX system).

For example, to create the user **jdoe** with the **admin** role, you would create an entry like the following:

```
jdoe=AAAAB3NzaC1kc3MAAACBAP1/U4EddRIpUt9KnC7s5Of2EbdSPO9EAMMeP4C2USZpRV1AI
lH7WT2NWPq/xfW6MPbLm1Vs14E7gB00b/JmYLdrmVClpJ+f6AR7ECLCT7up1/63xhv4O1fnfqi
mFQ8E+4P208UewwI1VBNaFpEy9nXzrith1yrv8iIDGZ3RSAHHAAAAFQCXYFCPFSMLzLKSuYKi6
4QL8Fgc9QAAAnEA9+GghdabPd7LvKtcNrhXuXmUr7v6OuqC+VdMCz0HgmdRWVeOutRZT+ZxBxC
BgLRJFnEj6EwoFhO3zwkyjMim4TwWeotifI0o4KOuHiuzpnWRbqN/C/ohNWLx+2J6ASQ7zKTxv
qhRkImog9/hWuWfBpKLZl6Ae1UlZAFMO/7PSSoAAACBAKKSU2PFl/qOLxIwmBZPPIcJshVe7bV
UpFvyl3BbJDow8rXfskl8wO63OzP/qLmcJM0+JbcRU/53Jj7uyk31drV2qxhIOsLDC9dGCWj47
Y7TyhPdXh/0dthTRBy6bqGtRPxGa7gJov1xm/UuYYXPIUR/3x9MAZvZ5xvE0kYXO+rx,admin
```

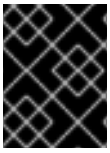

IMPORTANT

Do not insert the entire contents of the **id_rsa.pub** file here. Insert just the block of symbols which represents the public key itself.

User groups can also be defined in this file, where each user group is represented by a single line in the following format:

```
_g_\:GroupName=Role1[,Role2]...
```

Sample Blueprint configuration

The following Blueprint configuration shows how to define a new **karaf** realm using the public key login module, where the default **karaf** realm is overridden by setting the **rank** attribute to **200**:

```
<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.1.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0">

<!--Allow usage of System properties, especially the karaf.base property--
>
  <ext:property-placeholder
       placeholder-prefix="$[" placeholder-suffix="]"/>
```

```
  <jaas:config name="karaf" rank="200">
    <jaas:module flags="required"
classname="org.apache.karaf.jaas.modules.publickey.PublickeyLoginModule">
        users = $[karaf.base]/etc/keys.properties
    </jaas:module>
  </jaas:config>

</blueprint>
```

In this example, the user data will be stored in the file, ***InstallDir*/etc/keys.properties**, and it is not possible to edit the configuration using the **jaas:*** console commands.

2.1.6. JAAS JDBC Login Module

Overview

The JAAS JDBC login module enables you to store user data in a database back-end, using Java Database Connectivity (JDBC) to connect to the database. Hence, you can use any database that supports JDBC to store your user data. To manage the user data, you can use either the native database client tools or the **jaas:*** console commands (where the backing engine uses configured SQL queries to perform the relevant database updates).

You can combine multiple login modules with each login module providing both the authentication and authorization components. For example, you can combine default **PropertiesLoginModule** with **JDBCLoginModule** to ensure access to the system.

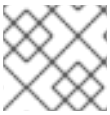

NOTE

User groups are **not** supported in the JAAS JDBC login module.

Supported credentials

The JAAS JDBC Login Module authenticates username/password credentials, returning the list of roles associated with the authenticated user.

Implementation classes

The following classes implement the JAAS JDBC Login Module:

org.apache.karaf.jaas.modules.jdbc.JDBCLoginModule

Implements the JAAS login module.

org.apache.karaf.jaas.modules.jdbc.JDBCBackingEngineFactory

Must be exposed as an OSGi service. This service makes it possible for you to manage the user data using the **jaas:*** console commands from the Apache Karaf shell (see olink:FMQCommandRef/Consolejaas).

Options

The JAAS JDBC login module supports the following options:

datasource

The JDBC data source, specified either as an OSGi service or as a JNDI name. You can specify a data source's OSGi service using the following syntax:

```
osgi:ServiceInterfaceName[/ServicePropertiesFilter]
```

The *ServiceInterfaceName* is the interface or class that is exported by the data source's OSGi service (usually **`javax.sql.DataSource`**).

Because multiple data sources can be exported as OSGi services in a Karaf container, it is usually necessary to specify a filter, *ServicePropertiesFilter*, to select the particular data source that you want. Filters on OSGi services are applied to the service property settings and follow a syntax that is borrowed from LDAP filter syntax.

query.password

The SQL query that retrieves the user's password. The query can contain a single question mark character, **?**, which is substituted by the username at run time.

query.role

The SQL query that retrieves the user's roles. The query can contain a single question mark character, **?**, which is substituted by the username at run time.

insert.user

The SQL query that creates a new user entry. The query can contain two question marks, **?**, characters: the first question mark is substituted by the username and the second question mark is substituted by the password at run time.

insert.role

The SQL query that adds a role to a user entry. The query can contain two question marks, **?**, characters: the first question mark is substituted by the username and the second question mark is substituted by the role at run time.

delete.user

The SQL query that deletes a user entry. The query can contain a single question mark character, **?**, which is substituted by the username at run time.

delete.role

The SQL query that deletes a role from a user entry. The query can contain two question marks, **?**, characters: the first question mark is substituted by the username and the second question mark is substituted by the role at run time.

delete.roles

The SQL query that deletes multiple roles from a user entry. The query can contain a single question mark character, **?**, which is substituted by the username at run time.

Example of setting up a JDBC login module

To set up a JDBC login module, perform the following main steps:

1. the section called "Create the database tables"
2. the section called "Create the data source"
3. the section called "Specify the data source as an OSGi service"

Create the database tables

Before you can set up the JDBC login module, you must set up a users table and a roles table in the backing database to store the user data. For example, the following SQL commands show how to create a suitable **users** table and **roles** table:

```
CREATE TABLE users (
  username VARCHAR(255) NOT NULL,
  password VARCHAR(255) NOT NULL,
  PRIMARY KEY (username)
);
CREATE TABLE roles (
  username VARCHAR(255) NOT NULL,
  role VARCHAR(255) NOT NULL,
  PRIMARY KEY (username,role)
);
```

The **users** table stores username/password data and the **roles** table associates a username with one or more roles.

Create the data source

To use a JDBC datasource with the JDBC login module, the correct approach to take is to create a data source instance and export the data source as an OSGi service. The JDBC login module can then access the data source by referencing the exported OSGi service. For example, you could create a MySQL data source instance and expose it as an OSGi service (of **javax.sql.DataSource** type) using code like the following in a Blueprint file:

```
<blueprint xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
           xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0">
  <bean id="mysqlDatasource"
        class="com.mysql.jdbc.jdbc2.optional.MysqlDataSource">
    <property name="serverName" value="localhost"></property>
    <property name="databaseName" value="DBName"></property>
    <property name="port" value="3306"></property>
    <property name="user" value="DBUser"></property>
    <property name="password" value="DBPassword"></property>
  </bean>

  <service id="mysqlDS" interface="javax.sql.DataSource"
        ref="mysqlDatasource">
    <service-properties>
        <entry key="osgi.jndi.service.name" value="jdbc/karafdb"/>
    </service-properties>
  </service>
</blueprint>
```

The preceding Blueprint configuration should be packaged and installed in the Karaf container as an OSGi bundle.

Specify the data source as an OSGi service

After the data source has been instantiated and exported as an OSGi service, you are ready to configure the JDBC login module. In particular, the **datasource** option of the JDBC login module can reference the data source's OSGi service using the following syntax:

```
osgi:javax.sql.DataSource/(osgi.jndi.service.name=jdbc/karafdb)
```

Where **javax.sql.DataSource** is the interface type of the exported OSGi service and the filter, **(osgi.jndi.service.name=jdbc/karafdb)**, selects the particular **javax.sql.DataSource** instance whose **osgi.jndi.service.name** service property has the value, **jdbc/karafdb**.

For example, you can use the following Blueprint configuration to override the **karaf** realm with a JDBC login module that references the sample MySQL data source:

```
<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.1.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0">

<!--Allow usage of System properties, especially the karaf.base property--
>
  <ext:property-placeholder
       placeholder-prefix="$[" placeholder-suffix="]"/>

  <jaas:config name="karaf" rank="200">
    <jaas:module flags="required"
      className="org.apache.karaf.jaas.modules.jdbc.JDBCLoginModule">
        datasource =
osgi:javax.sql.DataSource/(osgi.jndi.service.name=jdbc/karafdb)
        query.password = SELECT password FROM users WHERE username=?
        query.role = SELECT role FROM roles WHERE username=?
        insert.user = INSERT INTO users VALUES(?,?)
        insert.role = INSERT INTO roles VALUES(?,?)
        delete.user = DELETE FROM users WHERE username=?
        delete.role = DELETE FROM roles WHERE username=? AND role=?
        delete.roles = DELETE FROM roles WHERE username=?
    </jaas:module>
  </jaas:config>

  <!-- The Backing Engine Factory Service for the JDBCLoginModule -->
  <service interface="org.apache.karaf.jaas.modules.BackingEngineFactory">
    <bean
class="org.apache.karaf.jaas.modules.jdbc.JDBCBackingEngineFactory"/>
  </service>

</blueprint>
```

NOTE

The SQL statements shown in the preceding configuration are in fact the default values of these options. Hence, if you create user and role tables consistent with these SQL statements, you could omit the options settings and rely on the defaults.

In addition to creating a JDBCLoginModule, the preceding Blueprint configuration also instantiates and exports a **JDBCBackingEngineFactory** instance, which enables you to manage the user data using the **jaas:*** console commands.

2.1.7. JAAS LDAP Login Module

Overview

The JAAS LDAP login module enables you to store user data in an LDAP database. To manage the stored user data, use a standard LDAP client tool. The **`jaas:*`** console commands are **not** supported.

For more details about using LDAP with Red Hat Fuse see Chapter 8, *LDAP Authentication Tutorial*.

NOTE

User groups are **not** supported in the JAAS LDAP login module.

Supported credentials

The JAAS LDAP Login Module authenticates username/password credentials, returning the list of roles associated with the authenticated user.

Implementation classes

The following classes implement the JAAS LDAP Login Module:

`org.apache.karaf.jaas.modules.ldap.LDAPLoginModule`

Implements the JAAS login module. It is preloaded in the Karaf container, so you do not need to install its bundle.

NOTE

There is no backing engine factory for the LDAP Login Module, which means that this module cannot be managed using the **`jaas:*`** console commands.

Options

The JAAS LDAP login module supports the following options:

`authentication`

Specifies the authentication method used when binding to the LDAP server. Valid values are

- **`simple`**—bind with user name and password authentication, requiring you to set the **`connection.username`** and **`connection.password`** properties.
- **`none`**—bind anonymously. In this case the **`connection.username`** and **`connection.password`** properties can be left unassigned.

NOTE

The connection to the directory server is used only for performing searches. In this case, an anonymous bind is often preferred, because it is faster than an authenticated bind (but you would also need to ensure that the directory server is sufficiently protected, for example by deploying it behind a firewall).

`connection.url`

Specifies specify the location of the directory server using an ldap URL, **ldap://*Host*:*Port***. You can optionally qualify this URL, by adding a forward slash, **/**, followed by the DN of a particular node in the directory tree. To enable SSL security on the connection, you need to specify the **ldaps:** scheme in the URL—for example, **ldaps://*Host*:*Port***. You can also specify multiple URLs, as a space-separated list, for example:

```
connection.url=ldap://10.0.0.153:2389 ldap://10.10.178.20:389
```

connection.username

Specifies the DN of the user that opens the connection to the directory server. For example, **uid=admin,ou=system**.

connection.password

Specifies the password that matches the DN from **connection.username**. In the directory server, the password is normally stored as a **userPassword** attribute in the corresponding directory entry.

context.com.sun.jndi.ldap.connect.pool

If **true**, enables connection pooling for LDAP connections. Default is **false**.

context.com.sun.jndi.ldap.connect.timeout

Specifies the timeout for creating a TCP connection to the LDAP server, in units of milliseconds. We recommend that you set this property explicitly, because the default value is infinite, which can result in a hung connection attempt.

context.com.sun.jndi.ldap.read.timeout

Specifies the read timeout for an LDAP operation, in units of milliseconds. We recommend that you set this property explicitly, because the default value is infinite.

context.java.naming.referral

An *LDAP referral* is a form of indirection supported by some LDAP servers. The LDAP referral is an entry in the LDAP server which contains one or more URLs (usually referencing a node or nodes in another LDAP server). The **context.java.naming.referral** property can be used to enable or disable referral following. It can be set to one of the following values:

- **follow** to follow the referrals (assuming it is supported by the LDAP server),
- **ignore** to silently ignore all referrals,
- **throw** to throw a **PartialResultException** whenever a referral is encountered.

disableCache

The user and role caches can be disabled by setting this property to **true**. Default is **false**.

initial.context.factory

Specifies the class of the context factory used to connect to the LDAP server. This must always be set to **com.sun.jndi.ldap.LdapCtxFactory**.

role.base.dn

Specifies the DN of the subtree of the DIT to search for role entries. For example, **ou=groups,ou=system**.

role.filter

Specifies the LDAP search filter used to locate roles. It is applied to the subtree selected by **role.base.dn**. For example, **(member=uid=%u)**. Before being passed to the LDAP search operation, the value is subjected to string substitution, as follows:

- **%u** is replaced by the user name extracted from the incoming credentials, and
- **%dn** is replaced by the RDN of the corresponding user in the LDAP server (which was found by matching against the **user.filter** filter).
- **%fqdn** is replaced by the DN of the corresponding user in the LDAP server (which was found by matching against the **user.filter** filter).

role.mapping

Specifies the mapping between LDAP groups and JAAS roles. If no mapping is specified, the default mapping is for each LDAP group to map to the corresponding JAAS role of the same name. The role mapping is specified with the following syntax:

```
ldap-group=jaas-role(,jaas-role)*(;ldap-group=jaas-role(,jaas-role)*)*
```

Where each LDAP group, ***ldap-group***, is specified by its Common Name (CN).

For example, given the LDAP groups, **admin**, **devop**, and **tester**, you could map them to JAAS roles, as follows:

```
role.mapping=admin=admin;devop=admin,manager;tester=viewer
```

role.name.attribute

Specifies the attribute type of the role entry that contains the name of the role/group. If you omit this option, the role search feature is effectively disabled. For example, **cn**.

role.search.subtree

Specifies whether the role entry search scope includes the subtrees of the tree selected by **role.base.dn**. If **true**, the role lookup is recursive (**SUBTREE**). If **false**, the role lookup is performed only at the first level (**ONELEVEL**).

ssl

Specifies whether the connection to the LDAP server is secured using SSL. If **connection.url** starts with **ldaps://** SSL is used regardless of this property.

ssl.provider

Specifies the SSL provider to use for the LDAP connection. If not specified, the default SSL provider is used.

ssl.protocol

Specifies the protocol to use for the SSL connection. You **must** set this property to **TLSv1**, in order to prevent the SSLv3 protocol from being used (POODLE vulnerability).

ssl.algorithm

Specifies the algorithm used by the trust store manager. For example, **PKIX**.

ssl.keystore

The ID of the keystore that stores the LDAP client's own X.509 certificate (required only if SSL client authentication is enabled on the LDAP server). The keystore must be deployed using a **jaas:keystore** element (see the section called "Sample configuration for Apache DS").

ssl.keyalias

The keystore alias of the LDAP client's own X.509 certificate (required only if there is more than one certificate stored in the keystore specified by **ssl.keystore**).

`ssl.truststore`

The ID of the keystore that stores trusted CA certificates, which are used to verify the LDAP server's certificate (the LDAP server's certificate chain must be signed by one of the certificates in the truststore). The keystore must be deployed using a **`jaas:keystore`** element.

`user.base.dn`

Specifies the DN of the subtree of the DIT to search for user entries. For example, **`ou=users,ou=system`**.

`user.filter`

Specifies the LDAP search filter used to locate user credentials. It is applied to the subtree selected by **`user.base.dn`**. For example, **`(uid=%u)`**. Before being passed to the LDAP search operation, the value is subjected to string substitution, as follows:

- **`%u`** is replaced by the user name extracted from the incoming credentials.

`user.search.subtree`

Specifies whether the user entry search scope includes the subtrees of the tree selected by **`user.base.dn`**. If **`true`**, the user lookup is recursive (**`SUBTREE`**). If **`false`**, the user lookup is performed only at the first level (**`ONELEVEL`**).

Sample configuration for Apache DS

The following Blueprint configuration shows how to define a new **`karaf`** realm using the LDAP login module, where the default **`karaf`** realm is overridden by setting the **`rank`** attribute to **`200`**, and the LDAP login module connects to an Apache Directory Server:

```
<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.1.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0">

  <jaas:config name="karaf" rank="100">

    <jaas:module
className="org.apache.karaf.jaas.modules.ldap.LDAPLoginModule"
flags="sufficient">
      debug=true

      <!-- LDAP Configuration -->
      initialContextFactory=com.sun.jndi.ldap.LdapCtxFactory
<!--  multiple LDAP servers can be specified as a space separated list of
URLs -->
      connection.url=ldap://10.0.0.153:2389 ldap://10.10.178.20:389

<!--  authentication=none -->
      authentication=simple
      connection.username=cn=Directory Manager
      connection.password=directory

      <!-- User Info -->
      user.base.dn=dc=redhat,dc=com
      user.filter=(&amp;(objectClass=InetOrgPerson)(uid=%u))
```

```
      user.search.subtree=true

      <!-- Role/Group Info-->
      role.base.dn=dc=redhat,dc=com
      role.name.attribute=cn
<!--
      The 'dc=redhat,dc=com' used in the role.filter
      below is the user.base.dn.
-->
<!--      role.filter=(uniquemember=%dn,dc=redhat,dc=com) -->
      role.filter=(&amp;(objectClass=GroupOfUniqueNames)
(UniqueMember=%fqdn))
      role.search.subtree=true

<!-- role mappings - a ';' separated list -->
      role.mapping=JBossAdmin=admin;JBossMonitor=viewer

<!-- LDAP context properties -->
      context.com.sun.jndi.ldap.connect.timeout=5000
      context.com.sun.jndi.ldap.read.timeout=5000

<!-- LDAP connection pooling -->
<!-- http://docs.oracle.com/javase/jndi/tutorial/ldap/connect/pool.html --
>
<!-- http://docs.oracle.com/javase/jndi/tutorial/ldap/connect/config.html
-->
     context.com.sun.jndi.ldap.connect.pool=true

<!-- How are LDAP referrals handled?

     Can be `follow`, `ignore` or `throw`.  Configuring `follow` may not
work on all LDAP servers, `ignore` will
     silently ignore all referrals, while `throw` will throw a partial
results exception if there is a referral.
-->
     context.java.naming.referral=ignore

<!-- SSL configuration -->
     ssl=false
     ssl.protocol=SSL
<!-- matches the keystore/truststore configured below -->
     ssl.truststore=ks
     ssl.algorithm=PKIX
<!-- The User and Role caches can be disabled - 6.3.0 179 and later -->
     disableCache=true
    </jaas:module>
  </jaas:config>

  <!-- Location of the SSL truststore/keystore
  <jaas:keystore name="ks"
path="file:///${karaf.home}/etc/ldap.truststore" keystorePassword="XXXXXX"
/>
-->
</blueprint>
```

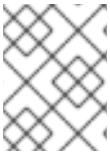

NOTE

In order to enable SSL, you must remember to use the **`ldaps`** scheme in the **`connection.url`** setting.

IMPORTANT

You must set **`ssl.protocol`** to **`TLSv1`** (or later), in order to protect against the Poodle vulnerability (CVE-2014-3566)

Filter settings for different directory servers

The most significant differences between directory servers arise in connection with setting the filter options in the LDAP login module. The precise settings depend ultimately on the organisation of your DIT, but the following table gives an idea of the typical role filter settings required for different directory servers:

Directory Server	Typical Filter Settings

389-DS
 	
Red Hat DS
	`user.filter=(&`
	`(objectClass=InetOrgPerson)`
	`(uid=%u))`
	`role.filter=(uniquemember=%fqdn)`

MS Active Directory	`user.filter=(&`
	`(objectCategory=person)`
	`(samAccountName=%u))`
	`role.filter=(uniquemember=%fqdn)`

Apache DS	`user.filter=(uid=%u)` `role.filter=(member=uid=%u)`
OpenLDAP	`user.filter=(uid=%u)` `role.filter=(member:=uid=%u)`

NOTE

In the preceding table, the **&** symbol (representing the logical **And** operator) is escaped as **`&`** because the option settings will be embedded in a Blueprint XML file.

2.1.8. Encrypting Stored Passwords

By default, the JAAS login modules store passwords in plaintext format. Although you can (and should) protect such data by setting file permissions appropriately, you can provide additional protection to passwords by storing them in an obscured format (using a *message digest* algorithm).

Red Hat Fuse provides a set of options for enabling password encryption, which can be combined with **any** of the JAAS login modules (except the public key login module, where it is not needed).

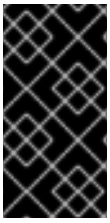

IMPORTANT

Although message digest algorithms are difficult to crack, they are not invulnerable to attack (for example, see the Wikipedia article on cryptographic hash functions). Always use file permissions to protect files containing passwords, in addition to using password encryption.

Options

You can optionally enable password encryption for JAAS login modules by setting the following login module properties. To do so, either edit the *InstallDir*`/etc/org.apache.karaf.jaas.cfg` file or deploy your own blueprint file as described in the section called “Example of a login module with Jasypt encryption”.

`encryption.enabled`

Set to **`true`**, to enable password encryption.

`encryption.name`

Name of the encryption service, which has been registered as an OSGi service.

`encryption.prefix`

Prefix for encrypted passwords.

`encryption.suffix`

Suffix for encrypted passwords.

`encryption.algorithm`

Specifies the name of the encryption algorithm—for example, **`MD5`** or **`SHA-1`**. You can specify one of the following encryption algorithms:

- **`MD2`**
- **`MD5`**
- **`SHA-1`**
- **`SHA-256`**
- **`SHA-384`**
- **`SHA-512`**

`encryption.encoding`

Encrypted passwords encoding: **`hexadecimal`** or **`base64`**.

`encryption.providerName` (Jasypt only)

Name of the **`java.security.Provider`** instance that is to provide the digest algorithm.

`encryption.providerClassName` (Jasypt only)

Class name of the security provider that is to provide the digest algorithm

`encryption.iterations` (Jasypt only)

Number of times to apply the hash function recursively.

`encryption.saltSizeBytes` (Jasypt only)

Size of the salt used to compute the digest.

`encryption.saltGeneratorClassName` (Jasypt only)

Class name of the salt generator.

`role.policy`

Specifies the policy for identifying role principals. Can have the values, **`prefix`** or **`group`**.

`role.discriminator`

Specifies the discriminator value to be used by the role policy.

Encryption services

There are two encryption services provided by Fuse:

- **`encryption.name = basic`**, described in the section called “Basic encryption service”,
- **`encryption.name = jasypt`**, described in the section called “Jasypt encryption”.

You can also create your own encryption service. To do so, you need to:

- Implement the **`org.apache.karaf.jaas.modules.EncryptionService`** interface, and
- Expose your implementation as OSGI service.

The following listing shows how to expose a custom encryption service to the OSGI container:

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0">

    <service interface="org.apache.karaf.jaas.modules.EncryptionService">
        <service-properties>
            <entry key="name" value="jasypt" />
        </service-properties>
        <bean
class="org.apache.karaf.jaas.jasypt.impl.JasyptEncryptionService"/>
    </service>
    ...
</blueprint>
```

Basic encryption service

The basic encryption service is installed in the Karaf container by default and you can reference it by setting the **`encryption.name`** property to the value, **`basic`**. In the basic encryption service, the message digest algorithms are provided by the SUN security provider (the default security provider in the Oracle JDK).

Jasypt encryption

The Jasypt encryption service is normally installed by default on Karaf. If necessary, you can install it explicitly by installing the **`jasypt-encryption`** feature, as follows:

```
JBossA-MQ:karaf@root> features:install jasypt-encryption
```

This command installs the requisite Jasypt bundles and exports Jasypt encryption as an OSGi service, so that it is available for use by JAAS login modules. To access the Jasypt encryption service, set the **`encryption.name`** property to the value, **`jasypt`**.

For more information about Jasypt encryption, see the Jasypt documentation.

Example of a login module with Jasypt encryption

Assuming that you have already installed the **`jasypt-encryption`** feature, you could deploy a properties login module with Jasypt encryption using the following Blueprint configuration:

```
<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.1.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0">

  <type-converters>
      <bean
class="org.apache.karaf.jaas.modules.properties.PropertiesConverter"/>
  </type-converters>

<!--Allow usage of System properties, especially the karaf.base property--
>
  <ext:property-placeholder
       placeholder-prefix="$[" placeholder-suffix="]"/>

  <jaas:config name="karaf" rank="200">
      <jaas:module flags="required"
className="org.apache.karaf.jaas.modules.properties.PropertiesLoginModule"
>
        users = $[karaf.base]/etc/users.properties
        encryption.enabled = true
        encryption.name = jasypt
        encryption.algorithm = SHA-256
        encryption.encoding = base64
        encryption.iterations = 100000
        encryption.saltSizeBytes = 16
        encryption.prefix = {CRYPT}
        encryption.suffix = {CRYPT}
      </jaas:module>
  </jaas:config>

  <!-- The Backing Engine Factory Service for the PropertiesLoginModule --
>
  <service interface="org.apache.karaf.jaas.modules.BackingEngineFactory">
      <bean
class="org.apache.karaf.jaas.modules.properties.PropertiesBackingEngineFac
tory"/>
  </service>

  <!-- Enable automatic encryption of all user passwords
    in InstallDir/etc/users.properties file.
    No login required to activate.
    Encrypted passwords appear in the
```

```
    InstallDir/etc/users.properties file as values enclosed
    by {CRYPT}...{CRYPT} prefix/suffix pairs -->

  <bean init-method="init" destroy-method="destroy"
class="org.apache.karaf.jaas.modules.properties.AutoEncryptionSupport">
    <argument>
      <map>
        <entry key="org.osgi.framework.BundleContext"
              value-ref="blueprintBundleContext"/>
        <entry key="users" value="$[karaf.base]/etc/users.properties"/>
        <entry key="encryption.name" value="jasypt"/>
        <entry key="encryption.enabled" value="true"/>
        <entry key="encryption.prefix" value="{CRYPT}"/>
        <entry key="encryption.suffix" value="{CRYPT}"/>
        <entry key="encryption.algorithm" value="SHA-256"/>
        <entry key="encryption.encoding" value="base64"/>
        <entry key="encryption.iterations" value="100000"/>
        <entry key="encryption.saltSizeBytes" value="16"/>
      </map>
    </argument>
  </bean>

</blueprint>
```

2.2. ROLE-BASED ACCESS CONTROL

Abstract

This section describes the role-based access control (RBAC) feature, which is enabled by default in the Karaf container. You can immediately start taking advantage of the RBAC feature, simply by adding one of the standard roles (such as **`manager`** or **`admin`**) to a user's credentials. For more advanced usage, you have the option of customizing the access control lists, in order to control exactly what each role can do. Finally, you have the option of applying custom ACLs to your own OSGi services.

2.2.1. Overview of Role-Based Access Control

By default, the Fuse role-based access control protects access through the Fuse Management Console, JMX connections, and the Karaf command console. To use the default levels of access control, simply add any of the standard roles to your user authentication data (for example, by editing the **`users.properties`** file). You also have the option of customizing access control, by editing the relevant Access Control List (ACL) files.

Mechanisms

Role-based access control in Karaf is based on the following mechanisms:

JMX Guard

The Karaf container is configured with a JMX guard, which intercepts every incoming JMX invocation and filters the invocation through the configured JMX access control lists. The JMX guard is configured at the JVM level, so it intercepts **every** JMX invocation, without exception.

OSGi Service Guard

For any OSGi service, it is possible to configure an OSGi service guard. The OSGi service guard is implemented as a proxy object, which interposes itself between the client and the original OSGi

service. An OSGi service guard must be explicitly configured for each OSGi service: it is not installed by default (except for the OSGi services that represent Karaf console commands, which are preconfigured for you).

Types of protection

The Fuse implementation of role-based access control is capable of providing the following types of protection:

Fuse Console (Hawtio)

Container access through the Fuse Console (Hawtio) is controlled by the JMX ACL files. The REST/HTTP service that provides the Fuse Console is implemented using Jolokia technology, which is layered above JMX. Hence, ultimately, all Fuse Console invocations pass through JMX and are regulated by JMX ACLs.

JMX

Direct access to the Karaf container's JMX port is regulated by the JMX ACLs. Moreover, any additional JMX ports opened by an application running in the Karaf container would also be regulated by the JMX ACLs, because the JMX guard is set at the JVM level.

Karaf command console

Access to the Karaf command console is regulated by the command console ACL files. Access control is applied no matter how the Karaf console is accessed. Whether accessing the command console through the Fuse Console or through the SSH protocol, access control is applied in both cases.

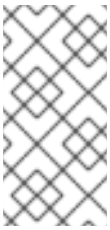

NOTE

In the special case where you start up the Karaf container directly at the command line (for example, using the **`./bin/fuse`** script) and no user authentication is performed, you automatically get the roles specified by the **`karaf.local.roles`** property in the **`etc/system.properties`** file.

OSGi services

For any OSGi service deployed in the Karaf container, you can optionally enable an ACL file, which restricts method invocations to specific roles.

Adding roles to users

In the system of role-based access control, you can give users permissions by adding roles to their user authentication data. For example, the following entry in the **`etc/users.properties`** file defines the **`admin`** user and grants the **`admin`** role.

```
admin = secretpass,group,admin,manager,viewer,systembundles,ssh
```

You also have the option of defining user groups and then assigning users to a particular user group. For example, you could define and use an **`admingroup`** user group as follows:

```
admin = secretpass, _g_:admingroup

_g_\:admingroup = group,admin,manager,viewer,systembundles,ssh
```

NOTE

User groups are not supported by every type of JAAS login module.

Standard roles

Table 2.2, “Standard Roles for Access Control” lists and describes the standard roles that are used throughout the JMX ACLs and the command console ACLs.

Table 2.2. Standard Roles for Access Control

Roles	Description
`viewer`	Grants read-only access to the Karaf container.
`manager`	Grants read-write access at the appropriate level for ordinary users, who want to deploy and run applications. But blocks access to sensitive Karaf container configuration settings.
`admin`	Grants unrestricted access to the Karaf container.
`ssh`	Grants users permission to connect to the Karaf command console (through the ssh port).

ACL files

The standard set of ACL files are located under the **`etc/auth/`** directory of the Fuse installation, as follows:

`etc/auth/jmx.acl[.*].cfg`

JMX ACL files.

`etc/auth/org.apache.karaf.command.acl.*.cfg`

Command console ACL files.

Customizing role-based access control

A complete set of JMX ACL files and command console ACL files are provided by default. You are free to customize these ACLs as required to suit the requirements of your system. Details of how to do this are given in the following sections.

Additional properties for controlling access

The **`system.properties`** file under the **`etc`** directory provides the following additional properties for controlling access through the Karaf command console and the Fuse Console (Hawtio):

`karaf.local.roles`

Specifies the roles that apply when a user starts up the Karaf container console **locally** (for example, by running the script).

`hawtio.roles`

Specifies the roles that are allowed to access the Karaf container through the Fuse Console. This constraint is applied **in addition to** the access control defined by the JMX ACL files.

`karaf.secured.command.compulsory.roles`

Specifies the default roles required to invoke a Karaf console command, in case the console command is not configured explicitly by a command ACL file, **`etc/auth/org.apache.karaf.command.acl.*.cfg`**. A user must be configured with at least one of the roles from the list in order to invoke the command. The value is specified as a comma-separated list of roles.

2.2.2. Customizing the JMX ACLs

The JMX ACLs are stored in the OSGi Config Admin Service and are normally accessible as the files, **`etc/auth/jmx.acl.*.cfg`**. This section explains how you can customize the JMX ACLs by editing these files yourself.

Architecture

Figure 2.1, “Access Control Mechanism for JMX” shows an overview of the role-based access control mechanism for JMX connections to the Karaf container.

Figure 2.1. Access Control Mechanism for JMX

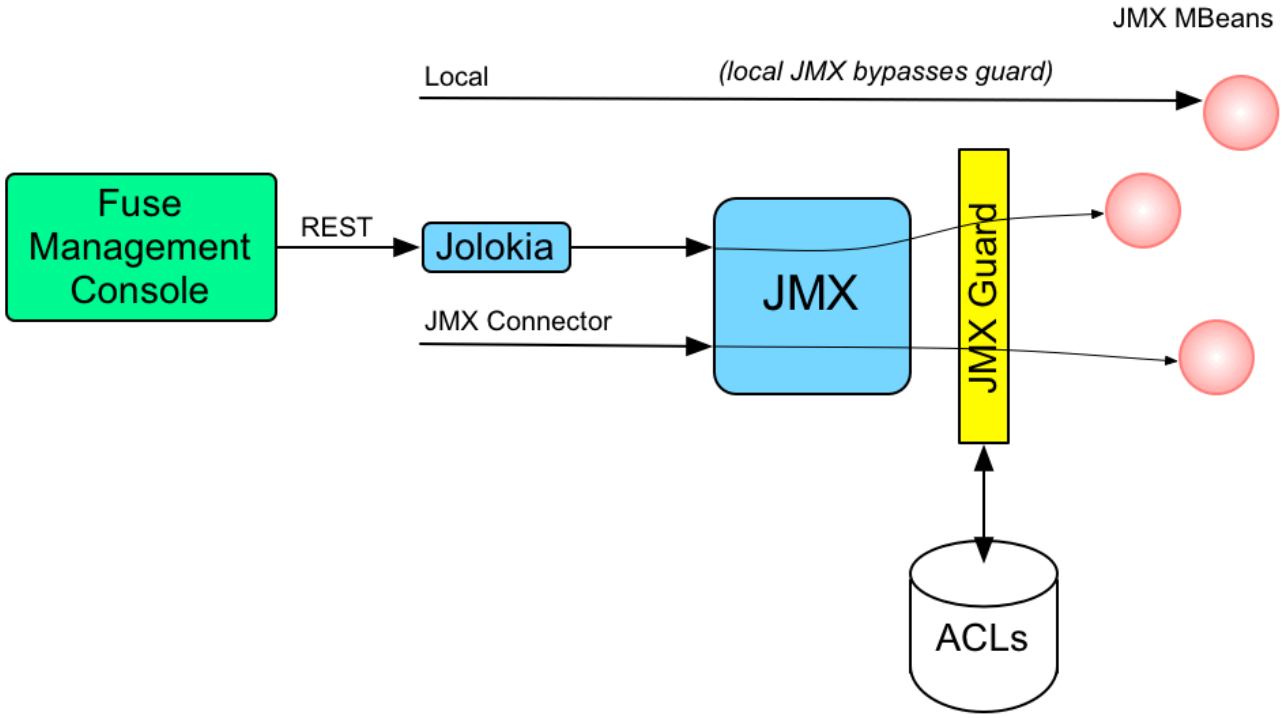

How it works

JMX access control works by inserting a *JMX Guard*, which is configured through a JVM-wide **`MBeanServerBuilder`** object. The Apache Karaf launching scripts have been modified to include the following setting:

```
-
Djavax.management.builder.initial=org.apache.karaf.management.boot.KarafMB
eanServerBuilder
```

JMX access control is then applied as follows:

1. For every non-local JMX invocation, the JVM-wide **`MBeanServerBuilder`** calls into an OSGi bundle that contains the JMX Guard.
2. The JMX Guard looks up the relevant ACL for the MBean the user is trying to access (where the ACLs are stored in the OSGi Config Admin service).
3. The ACL returns the list of roles that are allowed to make this particular invocation on the MBean.
4. The JMX Guard checks the list of roles against the current security subject (the user that is making the JMX invocation), to see whether the current user has any of the required roles.
5. If no matching role is found, the JMX invocation is blocked and a **`SecurityException`** is raised.

Location of JMX ACL files

The JMX ACL files are located in the ***`InstallDir`*`/etc/auth`** directory, where the ACL file names obey the following convention:

```
etc/auth/jmx.acl[.*].cfg
```

Technically, the ACLs are mapped to OSGi persistent IDs (PIDs), matching the pattern, **`jmx.acl[.*]`**. It just so happens that the Karaf container stores OSGi PIDs as files, **`PID.cfg`**, under the **`etc/`** directory by default.

Mapping MBeans to ACL file names

The JMX Guard applies access control to **every** MBean class that is accessed through JMX (including any MBeans you define in your own application code). The ACL file for a specific MBean class is derived from the MBean's Object Name, by prefixing it with **`jmx.acl`**. For example, given the MBean whose Object Name is given by **`org.apache.activemq:type=Broker`**, the corresponding PID would be:

```
jmx.acl.org.apache.activemq.Broker
```

The OSGi Config Admin service stores this PID data in the following file:

```
etc/auth/jmx.acl.org.apache.activemq.Broker.cfg
```

ACL file format

Each line of a JMX ACL file is an entry in the following format:

```
Pattern = Role1[,Role2][,Role3]...
```

Where ***`Pattern`*** is a pattern that matches a method invocation on an MBean, and the right-hand side of the equals sign is a comma-separated list of roles that give a user permission to make that invocation. In the simplest cases, the ***`Pattern`*** is simply a method name. For example, as in the following settings for the **`jmx.acl.hawtio.OSGiTools`** MBean (from the **`jmx.acl.hawtio.OSGiTools.cfg`** file):

```
getResourceURL = admin, manager, viewer
getLoadClassOrigin = admin, manager, viewer
```

It is also possible to use the wildcard character, **`*`**, to match multiple method names. For example, the

following entry gives permission to invoke all method names starting with **set**:

```
set* = admin, manager, viewer
```

But the ACL syntax is also capable of defining much more fine-grained control of method invocations. You can define patterns to match methods invoked with specific arguments or even arguments that match a regular expression. For example, the ACL for the **org.apache.karaf.config** MBean package exploits this capability to prevent ordinary users from modifying sensitive configuration settings. The **create** method from this package is restricted, as follows:

```
create(java.lang.String)[/jmx[.]acl.*/] = admin
create(java.lang.String)[/org[.]apache[.]karaf[.]command[.]acl.+/] = admin
create(java.lang.String)[/org[.]apache[.]karaf[.]service[.]acl.+/] = admin
create(java.lang.String) = admin, manager
```

In this case, the **manager** role generally has permission to invoke the **create** method, but only the **admin** role has permission to invoke **create** with a PID argument matching **jmx.acl.***, **org.apache.karaf.command.acl.***, or **org.apache.karaf.service.***.

For complete details of the ACL file format, please see the comments in the **etc/auth/jmx.acl.cfg** file.

ACL file hierarchy

Because it is often impractical to provide an ACL file for every single MBean, you have the option of specifying an ACL file at the level of a Java package, which provides default settings for **all** of the MBeans in that package. For example, the **org.apache.cxf.Bus** MBean could be affected by ACL settings at **any** of the following PID levels:

```
jmx.acl.org.apache.cxf.Bus
jmx.acl.org.apache.cxf
jmx.acl.org.apache
jmx.acl.org
jmx.acl
```

Where the most specific PID (top of the list) takes precedence over the least specific PID (bottom of the list).

Root ACL definitions

The root ACL file, **jmx.acl.cfg**, is a special case, because it supplies the default ACL settings for **all** MBeans. The root ACL has the following settings by default:

```
list* = admin, manager, viewer
get* = admin, manager, viewer
is* = admin, manager, viewer
set* = admin
* = admin
```

This implies that the typical **read** method patterns (**list***, **get***, **is***) are accessible to all standard roles, but the typical **write** method patterns and other methods (**set*** and *****) are accessible only to the admin role, **admin**.

Package ACL definitions

Many of the standard JMX ACL files provided in **`etc/auth/jmx.acl[.*].cfg`** apply to MBean packages. For example, the ACL for the **`org.apache.camel.endpoints`** MBean package is defined with the following permissions:

```
is* = admin, manager, viewer
get* = admin, manager, viewer
set* = admin, manager
```

ACL for custom MBeans

If you define custom MBeans in your own application, these custom MBeans are automatically integrated with the ACL mechanism and protected by the JMX Guard when you deploy them into the Karaf container. By default, however, your MBeans are typically protected only by the default root ACL file, **`jmx.acl.cfg`**. If you want to define a more fine-grained ACL for your MBean, create a new ACL file under **`etc/auth`**, using the standard JMX ACL file naming convention.

For example, if your custom MBean class has the JMX Object Name, **`org.example:type=MyMBean`**, create a new ACL file under the **`etc/auth`** directory called:

```
jmx.acl.org.example.MyMBean.cfg
```

Dynamic configuration at run time

Because the OSGi Config Admin service is dynamic, you can change ACL settings while the system is running, and even while a particular user is logged on. Hence, if you discover a security breach while the system is running, you can immediately restrict access to certain parts of the system by editing the relevant ACL file, without having to restart the Karaf container.

2.2.3. Customizing the Command Console ACLs

The command console ACLs are stored in the OSGi Config Admin Service and are normally accessible as the files, **`etc/auth/org.apache.karaf.command.acl.*.cfg`**. This section explains how you can customize the command console ACLs by editing these files yourself.

Architecture

Figure 2.2, “Access Control Mechanism for OSGi Services” shows an overview of the role-based access control mechanism for OSGi services in the Karaf container.

Figure 2.2. Access Control Mechanism for OSGi Services

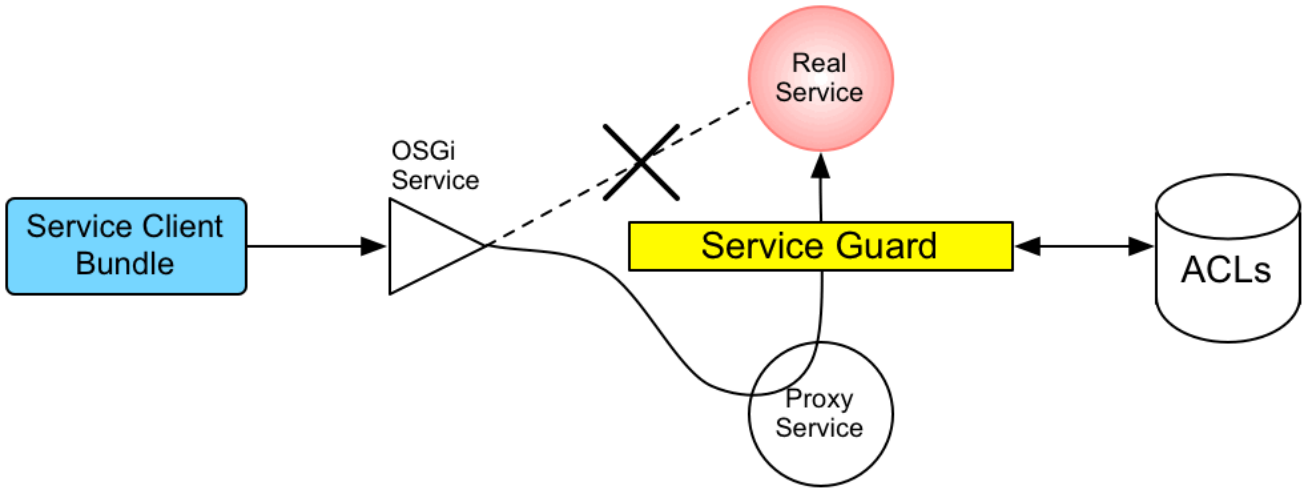

How it works

The mechanism for command console access control is, in fact, based on the generic access control mechanism for OSGi services. It so happens that console commands are implemented and exposed as OSGi services. The Karaf console itself discovers the available commands through the OSGi service registry and accesses the commands as OSGi services. Hence, the access control mechanism for OSGi services can be used to control access to console commands.

The mechanism for securing OSGi services is based on OSGi Service Registry Hooks. This is an advanced OSGi feature that makes it possible to hide OSGi services from certain consumers and to replace an OSGi service with a proxy service.

When a service guard is in place for a particular OSGi service, a client invocation on the OSGi service proceeds as follows:

1. The invocation does **not** go directly to the requested OSGi service. Instead, the request is routed to a replacement proxy service, which has the same service properties as the original service (and some extra ones).

2. The service guard looks up the relevant ACL for the target OSGi service (where the ACLs are stored in the OSGi Config Admin service).

3. The ACL returns the list of roles that are allowed to make this particular method invocation on the service.

4. If no ACL is found for this command, the service guard defaults to the list of roles specified in the **`karaf.secured.command.compulsory.roles`** property in the **`etc/system.properties`** file.

5. The service guard checks the list of roles against the current security subject (the user that is making the method invocation), to see whether the current user has any of the required roles.

6. If no matching role is found, the method invocation is blocked and a **`SecurityException`** is raised.

7. Alternatively, if a matching role is found, the method invocation is delegated to the original OSGi service.

Configuring default security roles

For any commands that do not have a corresponding ACL file, you specify a default list of security roles by setting the **karaf.secured.command.compulsory.roles** property in the **etc/system.properties** file (specified as a comma-separated list of roles).

Location of command console ACL files

The command console ACL files are located in the ***InstallDir*/etc/auth** directory, with the prefix, **org.apache.karaf.command.acl**.

Mapping command scopes to ACL file names

The command console ACL file names obey the following convention:

```
etc/auth/org.apache.karaf.command.acl.CommandScope.cfg
```

Where the ***CommandScope*** corresponds to the prefix for a particular group of Karaf console commands. For example, the **feature:install** and **features:uninstall** commands belong to the **feature** command scope, which has the corresponding ACL file, **org.apache.karaf.command.acl.features.cfg**.

ACL file format

Each line of a command console ACL file is an entry in the following format:

```
Pattern = Role1[,Role2][,Role3]...
```

Where ***Pattern*** is a pattern that matches a Karaf console command from the current command scope, and the right-hand side of the equals sign is a comma-separated list of roles that give a user permission to make that invocation. In the simplest cases, the ***Pattern*** is simply an unscoped command name. For example, the **org.apache.karaf.command.acl.feature.cfg** ACL file includes the following rules for the **feature** commands:

```
list = admin, manager, viewer
repo-list = admin, manager, viewer
info = admin, manager, viewer
version-list = admin, manager, viewer
repo-refresh = admin, manager
repo-add = admin, manager
repo-remove = admin, manager
install = admin
uninstall = admin
```

IMPORTANT

If no match is found for a specific command name, it is assumed that no role is required for this command and it can be invoked by any user.

You can also define patterns to match commands invoked with specific arguments or even arguments that match a regular expression. For example, the **org.apache.karaf.command.acl.bundle.cfg** ACL file exploits this capability to prevent ordinary users from invoking the **bundle:start** and **bundle:stop** commands with the **-f** (force) flag (which must be specified to manage system bundles). This restriction is coded as follows in the ACL file:

```
start[/.*[-][f].*/] = admin
start = admin, manager
stop[/.*[-][f].*/] = admin
stop = admin, manager
```

In this case, the **manager** role generally has permission to invoke the **bundle:start** and **bundle:stop** commands, but only the **admin** role has permission to invoke these commands with the force option, **-f**.

For complete details of the ACL file format, please see the comments in the **etc/auth/org.apache.karaf.command.acl.bundle.cfg** file.

Dynamic configuration at run time

The command console ACL settings are fully dynamic, which means you can change the ACL settings while the system is running and the changes will take effect within a few seconds, even for users that are already logged on.

2.2.4. Defining ACLs for OSGi Services

It is possible to define a custom ACL for any OSGi service (whether system level or application level). By default, OSGi services do not have access control enabled (with the exception of the OSGi services that expose Karaf console commands, which are pre-configured with command console ACL files). This section explains how to define a custom ACL for an OSGi service and how to invoke methods on that service using a specified role.

ACL file format

An OSGi service ACL file has one special entry, which identifies the OSGi service to which this ACL applies, as follows:

```
service.guard = (objectClass=InterfaceName)
```

Where the value of **service.guard** is an LDAP search filter that is applied to the registry of OSGi service properties in order to pick out the matching OSGi service. The simplest type of filter, **(objectClass=*InterfaceName*)**, picks out an OSGi service with the specified Java interface name, ***InterfaceName***.

The remaining entries in the ACL file are of the following form:

```
Pattern = Role1[,Role2][,Role3]...
```

Where ***Pattern*** is a pattern that matches a service method, and the right-hand side of the equals sign is a comma-separated list of roles that give a user permission to make that invocation. The syntax of these entries is essentially the same as the entries in a JMX ACL file—see the section called “ACL file format”.

How to define an ACL for a custom OSGi service

To define an ACL for a custom OSGi service, perform the following steps:

1. It is customary to define an OSGi service using a Java interface (you could use a regular Java class, but this is not recommended). For example, consider the Java interface, **MyService**, which we intend to expose as an OSGi service:

```
package org.example;

public interface MyService {
  void doit(String s);
}
```

2. To expose the Java interface as an OSGi service, you would typically add a **service** element to an OSGi Blueprint XML file (where the Blueprint XML file is typically stored under the **src/main/resources/OSGI-INF/blueprint** directory in a Maven project). For example, assuming that **MyServiceImpl** is the class that implements the **MyService** interface, you could expose the **MyService** OSGi service as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
           xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
           default-activation="lazy">

  <bean id="myserviceimpl" class="org.example.MyServiceImpl"/>

  <service id="myservice" ref="myserviceimpl"
interface="org.example.MyService"/>

</blueprint>
```

3. To define an ACL for the the OSGi service, you must create an OSGi Config Admin PID with the prefix, **org.apache.karaf.service.acl**.
   For example, in the case of a Karaf container (where the OSGi Config Admin PIDs are stored as **.cfg** files under the **etc/auth/** directory), you can create the following ACL file for the **MyService** OSGi service:

```
etc/auth/org.apache.karaf.service.acl.myservice.cfg
```

NOTE

It does not matter exactly how you name this file, as long as it starts with the required prefix, **org.apache.karaf.service.acl**. The corresponding OSGi service for this ACL file is actually specified by a property setting in this file (as you will see in the next step).

4. Specify the contents of the ACL file in a format like the following:

```
service.guard = (objectClass=InterfaceName)
Pattern = Role1[,Role2][,Role3]...
```

   The **service.guard** setting specifies the ***InterfaceName*** of the OSGi service (using the syntax of an LDAP search filter, which is applied to the OSGi service properties). The other entries in the ACL file consist of a method ***Pattern***, which associates a matching method to the specified roles. For example, you could define a simple ACL for the **MyService** OSGi service with the following settings in the **org.apache.karaf.service.acl.myservice.cfg** file:

```
service.guard = (objectClass=org.example.MyService)
doit = admin, manager, viewer
```

5. Finally, in order to enable the ACL for this OSGi service, you must edit the **karaf.secured.services** property in the **etc/system.properties** file. The value of the **karaf.secured.services** property has the syntax of an LDAP search filter (which gets applied to the OSGi service properties). In general, to enable ACLs for an OSGi service, ***ServiceInterface***, you must modify this property as follows:

```
karaf.secured.services=(|(objectClass=ServiceInterface)
(...ExistingPropValue...))
```

For example, to enable the **MyService** OSGi service:

```
karaf.secured.services=(|(objectClass=org.example.MyService)(&
(osgi.command.scope=*)(osgi.command.function=*)))
```

The initial value of the **karaf.secured.services** property has the settings to enable the command console ACLs. If you delete or corrupt these entries, the command console ACLs might stop working.

How to invoke an OSGi service secured with RBAC

If you are writing Java code to invoke methods an a custom OSGi service (that is, implementing a client of the OSGi service), you must use the Java security API to specify the role you are using to invoke the service. For example, to invoke the **MyService** OSGi service using the **manager** role, you could use code like the following:

```
// Java
import javax.security.auth.Subject;
import org.apache.karaf.jaas.boot.principal.RolePrincipal;
// ...
Subject s = new Subject();
s.getPrincipals().add(new RolePrincipal("Deployer"));
Subject.doAs(s, new PrivilegedAction() {
  public Object run() {
    svc.doit("foo"); // invoke the service
  }
}
```

NOTE

This example uses the Karaf role type, **org.apache.karaf.jaas.boot.principal.RolePrincipal**. If necessary, you could use your own custom role class instead, but in that case you would have to specify your roles using the syntax ***className*:*roleName*** in the OSGi service's ACL file.

How to discover the roles required by an OSGi service

When you are writing code against an OSGi service secured by an ACL, it can sometimes be useful to check what roles are allowed to invoke the service. For this purpose, the proxy service exports an additional OSGi property, **org.apache.karaf.service.guard.roles**. The value of this property is a **java.util.Collection** object, which contains a list of all the roles that could possibly invoke a method on that service.

2.3. USING ENCRYPTED PROPERTY PLACEHOLDERS

When securing a Karaf container, do not use plain text passwords in configuration files. One way to avoid this using plain text passwords is to use encrypted property placeholders when ever possible.

How to use encrypted property placeholders

To use encrypted property placeholders in a Blueprint XML file, perform the following steps:

1. Download and install Jasypt, to gain access to the Jasypt **`listAlgorithms.sh`**, **`encrypt.sh`** and **`decrypt.sh`** command-line tools.

   

   NOTE

   When installing the Jasypt command-line tools, you must enable execute permissions on the script files, by running **`chmod u+x ScriptName.sh`**.

2. Choose a master password and an encryption algorithm. To discover which algorithms are supported in your current Java environment, run the **`listAlgorithms.sh`** Jasypt command-line tool, as follows:

   ```
   ./listAlgorithms.sh
   DIGEST ALGORITHMS:   [MD2, MD5, SHA, SHA-256, SHA-384, SHA-512]

   PBE ALGORITHMS:      [PBEWITHMD5ANDDES, PBEWITHMD5ANDTRIPLEDES,
   PBEWITHSHA1ANDDESEDE, PBEWITHSHA1ANDRC2_40]
   ```

   On Windows platforms, the script is **`listAlgorithms.bat`**. Fuse uses **`PBEWithMD5AndDES`** by default.

3. Use the Jasypt encrypt command-line tool to encrypt your sensitive configuration values (for example, passwords for use in configuration files). For example, the following command encrypts the ***`PlaintextVal`*** value, using the specified algorithm and master password ***`MasterPass`***:

   ```
   ./encrypt.sh input="PlaintextVal" algorithm=PBEWithMD5AndDES
   password=MasterPass
   ```

4. Create a properties file with encrypted values. For example, suppose you wanted to store some LDAP credentials. You could create a file, **`etc/ldap.properties`**, with the following contents:

   Example 2.6. Property File with an Encrypted Property

   ```
   #ldap.properties
   ldap.password=ENC(amIsvdqno9iSwnd7kAlLYQ==)
   ldap.url=ldap://192.168.1.74:10389
   ```

   The encrypted property values (as generated in the previous step) are identified by wrapping in the **`ENC()`** function.

5. Add the required namespaces to your Blueprint XML file:

- Aries extensions—**http://aries.apache.org/blueprint/xmlns/blueprint-ext/v1.0.0**
- Apache Karaf Jasypt—**http://karaf.apache.org/xmlns/jasypt/v1.0.0**

Example 2.7, “Encrypted Property Namespaces” shows a Blueprint file with the requisite namespaces.

Example 2.7. Encrypted Property Namespaces

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0"
  xmlns:enc="http://karaf.apache.org/xmlns/jasypt/v1.0.0">
...
</blueprint>
```

6. Configure the location of the properties file for the property placeholder and configure the Jasypt encryption algorithm .
   Example 2.8, “Jasypt Blueprint Configuration” shows how to configure the **ext:property-placeholder** element to read properties from the **etc/ldap.properties** file. The **enc:property-placeholder** element configures Jasypt to use the **PBEWithMD5AndDES** encryption algorithm and to read the master password from the **JASYPT_ENCRYPTION_PASSWORD** environment variable.

Example 2.8. Jasypt Blueprint Configuration

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0"
  xmlns:enc="http://karaf.apache.org/xmlns/jasypt/v1.0.0">

  <ext:property-placeholder>
    <ext:location>file:etc/ldap.properties</ext:location>
  </ext:property-placeholder>

  <enc:property-placeholder>
    <enc:encryptor
class="org.jasypt.encryption.pbe.StandardPBEStringEncryptor">
      <property name="config">
        <bean
class="org.jasypt.encryption.pbe.config.EnvironmentStringPBEConfig
">
          <property name="algorithm" value="PBEWithMD5AndDES" />
          <property name="passwordEnvName"
value="JASYPT_ENCRYPTION_PASSWORD" />
        </bean>
      </property>
    </enc:encryptor>
  </enc:property-placeholder>
...
</blueprint>
```

Blueprint XML example

Example 2.9, “Jasypt Example in Blueprint XML” shows an example of an LDAP JAAS realm configured in Blueprint XML, using Jasypt encrypted property placeholders.

NOTE

When you use the process described in this topic to encrypt external properties you cannot use the **@PropertyInject** annotation to decrypt the properties. Instead, use XML to inject properties into Java objects, as shown in this Blueprint example.

Example 2.9. Jasypt Example in Blueprint XML

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0"
  xmlns:enc="http://karaf.apache.org/xmlns/jasypt/v1.0.0">

  <ext:property-placeholder>
    <location>file:etc/ldap.properties</location>
  </ext:property-placeholder>

  <enc:property-placeholder>
    <enc:encryptor
class="org.jasypt.encryption.pbe.StandardPBEStringEncryptor">
      <property name="config">
        <bean
class="org.jasypt.encryption.pbe.config.EnvironmentStringPBEConfig">
          <property name="algorithm" value="PBEWithMD5AndDES" />
          <property name="passwordEnvName"
value="JASYPT_ENCRYPTION_PASSWORD" />
        </bean>
      </property>
    </enc:encryptor>
  </enc:property-placeholder>

  <jaas:config name="karaf" rank="200">
    <jaas:module
className="org.apache.karaf.jaas.modules.ldap.LDAPLoginModule"
flags="required">
      initialContextFactory=com.sun.jndi.ldap.LdapCtxFactory
      debug=true
        connectionURL=${ldap.url}

connectionUsername=cn=mqbroker,ou=Services,ou=system,dc=jbossfuse,dc=com
        connectionPassword=${ldap.password}
        connectionProtocol=
        authentication=simple
        userRoleName=cn
        userBase = ou=User,ou=ActiveMQ,ou=system,dc=jbossfuse,dc=com
        userSearchMatching=(uid={0})
        userSearchSubtree=true
        roleBase = ou=Group,ou=ActiveMQ,ou=system,dc=jbossfuse,dc=com
        roleName=cn
        roleSearchMatching= (member:=uid={1})
```

```
        roleSearchSubtree=true
    </jaas:module>
  </jaas:config>

</blueprint>
```

The **${ldap.password}** placeholder is replaced with the decrypted value of the **ldap.password** property from the **etc/ldap.properties** properties file.

2.4. ENABLING REMOTE JMX SSL

Overview

Red Hat JBoss Fuse provides a JMX port that allows remote monitoring and management of Karaf containers using MBeans. By default, however, the credentials that you send over the JMX connection are unencrypted and vulnerable to snooping. To encrypt the JMX connection and protect against password snooping, you need to secure JMX communications by configuring JMX over SSL.

To configure JMX over SSL, perform the following steps:

1. Create the jbossweb.keystore file
2. Create and deploy the keystore.xml file
3. Add the required properties to org.apache.karaf.management.cfg
4. Restart the Fuse container

After you have configured JMX over SSL access, you should test the connection.

WARNING

If you are planning to enable SSL/TLS security, you must ensure that you explicitly disable the SSLv3 protocol, in order to safeguard against the Poodle vulnerability (CVE-2014-3566). For more details, see Disabling SSLv3 in JBoss Fuse 6.x and JBoss A-MQ 6.x.

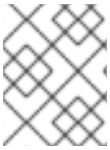

NOTE

If you configure JMX over SSL while Red Hat JBoss Fuse is running, you will need to restart it.

Prerequisites

If you haven't already done so, you need to:

- Set your **JAVA_HOME** environment variable

- Configure a Karaf user with the **admin** role
  Edit the ***InstallDir*/etc/users.properties** file and add the following entry, on a single line:

  ```
  admin=YourPassword,admin
  ```

  This creates a new user with username, **admin**, password, ***YourPassword***, and the **admin** role.

Create the jbossweb.keystore file

Open a command prompt and make sure you are in the **etc/** directory of your Karaf installation:

```
cd etc
```

At the command line, using a **-dname** value (Distinguished Name) appropriate for your application, type this command:

```
$JAVA_HOME/bin/keytool -genkey -v -alias jbossalias -keyalg RSA -keysize
1024 -keystore jbossweb.keystore -validity 3650 -keypass JbossPassword -
storepass JbossPassword -dname "CN=127.0.0.1, OU=RedHat Software Unit,
O=RedHat, L=Boston, S=Mass, C=USA"
```

IMPORTANT

Type the entire command on a single command line.

The command returns output that looks like this:

```
Generating 1,024 bit RSA key pair and self-signed certificate
(SHA256withRSA) with a validity of 3,650 days
 for: CN=127.0.0.1, OU=RedHat Software Unit, O=RedHat, L=Boston, ST=Mass,
C=USA
New certificate (self-signed):
[
[
  Version: V3
  Subject: CN=127.0.0.1, OU=RedHat Software Unit, O=RedHat, L=Boston,
ST=Mass, C=USA
  Signature Algorithm: SHA256withRSA, OID = 1.2.840.113549.1.1.11

  Key:  Sun RSA public key, 1024 bits
  modulus:
112308602579056704360496299050191816946109837286427320179534244008039380
8

159410077607500864745991099141380637280072294767016640781490175445910072027
9046

394462181373817732403106426038265948319382617744876203043766931839107261986
7218

036972335210839062722456085328301058362052369248473659880488338711351959835
357
```

```
  public exponent: 65537
  Validity: [From: Thu Jun 05 12:19:52 EDT 2014,
              To: Sun Jun 02 12:19:52 EDT 2024]
  Issuer: CN=127.0.0.1, OU=RedHat Software Unit, O=RedHat, L=Boston,
ST=Mass, C=USA
  SerialNumber: [    4666e4e6]

Certificate Extensions: 1
[1]: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: AC 44 A5 F2 E6 2F B2 5A   5F 88 FE 69 60 B4 27 7D  .D.../.Z_..i`.'.
0010: B9 81 23 9C                                        ..#.
]
]

]
  Algorithm: [SHA256withRSA]
  Signature:
0000: 01 1D 95 C0 F2 03 B0 FD   CF 3A 1A 14 F5 2E 04 E5  .........:......
0010: DD 18 DD 0E 24 60 00 54   35 AE FE 36 7B 38 69 4C  ....$`.T5..6.8iL
0020: 1E 85 0A AF AE 24 1B 40   62 C9 F4 E5 A9 02 CD D3  .....$.@b.......
0030: 91 57 60 F6 EF D6 A4 84   56 BA 5D 21 11 F7 EA 09  .W`.....V.]!....
0040: 73 D5 6B 48 4A A9 09 93   8C 05 58 91 6C D0 53 81  s.kHJ.....X.l.S.
0050: 39 D8 29 59 73 C4 61 BE   99 13 12 89 00 1C F8 38  9.)Ys.a........8
0060: E2 BF D5 3C 87 F6 3F FA   E1 75 69 DF 37 8E 37 B5  ...<..?..ui.7.7.
0070: B7 8D 10 CC 9E 70 E8 6D   C2 1A 90 FF 3C 91 84 50  .....p.m....<..P

]
[Storing jbossweb.keystore]
```

Check whether ***InstallDir*/etc** now contains the file, **jbossweb.keystore**.

Create and deploy the keystore.xml file

1. Using your favorite XML editor, create and save the **keystore.xml** file in the **<*installDir*>/jboss-fuse-7.0.0.fuse-000191-redhat-1/etc** directory.

2. Include this text in the file:

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0">
<jaas:keystore name="sample_keystore"
rank="1"
path="file:etc/jbossweb.keystore"
keystorePassword="JbossPassword"
keyPasswords="jbossalias=JbossPassword" />
</blueprint>
```

3. Deploy the **keystore.xml** file to the Karaf container, by copying it into the ***InstallDir*/deploy** directory (the hot deploy directory).

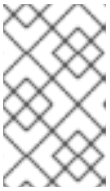

NOTE

Subsequently, if you need to undeploy the **keystore.xml** file, you can do so by deleting the **keystore.xml** file from the **deploy/** directory **while the Karaf container is running**.

Add the required properties to org.apache.karaf.management.cfg

Edit the ***InstallDir*/etc/org.apache.karaf.management.cfg** file to include these properties at the end of the file:

```
secured = true
secureProtocol = TLSv1
keyAlias = jbossalias
keyStore = sample_keystore
trustStore = sample_keystore
```

IMPORTANT

You must set **secureProtocol** to **TLSv1**, in order to protect against the Poodle vulnerability (CVE-2014-3566)

Restart the Karaf container

You must restart the Karaf container for the new JMX SSL/TLS settings to take effect.

Testing the Secure JMX connection

1. Open a command prompt and make sure you are in the **etc/** directory of your Fuse installation:

   ```
   cd <installDir>/jboss-fuse-7.0.0.fuse-000191-redhat-1/etc
   ```

2. Open a terminal, and start up JConsole by entering this command:

   ```
   jconsole -J-Djavax.net.debug=ssl -J-
   Djavax.net.ssl.trustStore=jbossweb.keystore -J-
   Djavax.net.ssl.trustStoreType=JKS -J-
   Djavax.net.ssl.trustStorePassword=JbossPassword
   ```

   Where the **-J-Djavax.net.ssl.trustStore** option specifies the location of the **jbossweb.keystore** file (make sure this location is specified correctly, or the SSL/TLS handshake will fail). The **-J-Djavax.net.debug=ssl** setting enables logging of SSL/TLS handshake messages, so you can verify that SSL/TLS has been successfully enabled.

   

   IMPORTANT

   Type the entire command on the same command line.

3. When JConsole opens, select the option **Remote Process** in the **New Connection** wizard.

4. Under the **Remote Process** option, enter the following value for the **service:jmx:<protocol>:<sap>** connection URL:

```
service:jmx:rmi://localhost:44444/jndi/rmi://localhost:1099/karaf-
root
```

And fill in the **Username**, and **Password** fields with valid JAAS credentials (as set in the **etc/users.properties** file):

```
Username: admin
Password: YourPassword
```

CHAPTER 3. SECURING THE UNDERTOW HTTP SERVER

Abstract

You can configure the built-in Undertow HTTP server to use SSL/TLS security by editing the contents of the **etc/undertow.xml** configuration file. In particular, you can add SSL/TLS security to the Fuse Console in this way.

3.1. UNDERTOW SERVER

The Fuse container is pre-configured with an Undertow server, which acts as a general-purpose HTTP server and HTTP servlet container. Through a single HTTP port (by default, **http://localhost:8181**), the Undertow container can host multiple services, for example:

- Fuse Console (by default, **http://localhost:8181/hawtio**)
- Apache CXF Web services endpoints (if the host and port are left unspecified in the endpoint configuration)
- Some Apache Camel endpoints

If you use the default Undertow server for all of your HTTP endpoints, you can conveniently add SSL/TLS security to these HTTP endpoints by following the steps described here.

3.2. CREATE X.509 CERTIFICATE AND PRIVATE KEY

Before you can enable SSL/TLS on the Undertow server, you must create an X.509 certificate and private key, where the certificate and private key must be in Java keystore format (JKS format). For details of how to create a signed certificate and private key, see Appendix A, *Managing Certificates*.

3.3. ENABLING SSL/TLS FOR UNDERTOW IN AN APACHE KARAF CONTAINER

For the following procedure, it is assumed that you have already created a signed X.509 certificate and private key pair in the keystore file, **alice.ks**, with keystore password, **StorePass**, and key password, **KeyPass**.

To enable SSL/TLS for Undertow in a Karaf container:

1. Make sure that the Pax Web server is configured to take its configuration from the **etc/undertow.xml** file. When you look at the contents of the **etc/org.ops4j.pax.web.cfg** file, you should see the following setting:

   ```
   org.ops4j.pax.web.config.file=${karaf.etc}/undertow.xml
   ```

2. Open the file, **etc/org.ops4j.pax.web.cfg**, in a text editor and add the following line:

   ```
   org.osgi.service.http.port.secure=8443
   ```

   Save and close the file, **etc/org.ops4j.pax.web.cfg**.

3. Open the file, **etc/undertow.xml**, in a text editor. The next steps assume you are working with the default **undertow.xml** file, unchanged since installation time.

4. Search for the XML elements, **http-listener** and **https-listener**. Comment out the **http-listener** element (by enclosing it between **<!--** and **-->**) and uncomment the **https-listener** element (spread over two lines). The edited fragment of XML should now look something like this:

```
<!-- HTTP(S) Listener references Socket Binding (and indirectly -
Interfaces) -->
<!-- http-listener name="http" socket-binding="http" /> -->
<!-- verify-client: org.xnio.SslClientAuthMode.NOT_REQUESTED,
org.xnio.SslClientAuthMode.REQUESTED,
org.xnio.SslClientAuthMode.REQUIRED -->
<https-listener name="https" socket-binding="https"
        security-realm="https" verify-client="NOT_REQUESTED" />
```

5. Search for the **w:keystore** element. By default, the **w:keystore** element is configured as follows:

```
<w:keystore path="${karaf.etc}/certs/server.keystore" provider="JKS"
alias="server"
            keystore-password="secret" key-password="secret"
            generate-self-signed-certificate-host="localhost" />
```

   To install the **alice** certificate as the Undertow server's certificate, modify the **w:keystore** element attributes as follows:

   - Set **path** to the absolute location of the **alice.ks** file on the file system.
   - Set **provider** to **JKS**.
   - Set **alias** to the **alice** certificate alias in the keystore.
   - Set **keystore-password** to the value of the password that unlocks the key store.
   - Set **key-password** to the value of the password that encrypts the **alice** private key.
   - Delete the **generate-self-signed-certificate-host** attribute setting.

6. For example, after installing the **alice.ks** keystore, the modified **w:keystore** element would look something like this:

```
<w:keystore path="${karaf.etc}/certs/alice.ks" provider="JKS"
alias="alice"
            keystore-password="StorePass" key-password="KeyPass" />
```

7. Search for the **<interface name="secure">** tag, which is used to specify the IP addresses the secure HTTPS port binds to. By default, this element is commented out, as follows:

```
<!--<interface name="secure">-->
    <!--<w:inet-address value="127.0.0.1" />-->
<!--</interface>-->
```

Uncomment the element and customize the **value** attribute to specify the IP address which the HTTPS port binds to. For example, the wildcard value, **0.0.0.0**, configures HTTPS to bind to all available IP addresses:

```
<interface name="secure">
    <w:inet-address value="0.0.0.0" />
</interface>
```

8. Search for and uncomment the **<socket-binding name="https"** tag. When this tag is uncommented, it should look something like this:

```
<socket-binding name="https" interface="secure"
port="${org.osgi.service.http.port.secure}" />
```

9. Save and close the file, **etc/undertow.xml**.

10. Restart the Fuse container, in order for the configuration changes to take effect.

3.4. CUSTOMIZING ALLOWED TLS PROTOCOLS AND CIPHER SUITES

You can customize the allowed TLS protocols and cipher suites by modifying the following attributes of the **w:engine** element in the **etc/undertow.xml** file:

enabled-cipher-suites

Specifies the list of allowed TLS/SSL cipher suites.

enabled-protocols

Specifies the list of allowed TLS/SSL protocols.

WARNING

Do not enable **SSL** protocol versions, as they are vulnerable to attack. Use only **TLS** protocol versions.

For full details of the available protocols and cipher suites, consult the appropriate JVM documentation and security provider documentation. For example, for Java 8, see Java Cryptography Architecture Oracle Providers Documentation for JDK 8.

3.5. CONNECT TO THE SECURE CONSOLE

After configuring SSL security for the Undertow server in the Pax Web configuration file, you should be able to open the Fuse Console by browsing to the following URL:

```
https://localhost:8443/hawtio
```

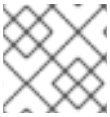

NOTE

Remember to type the **`https:`** scheme, instead of **`http:`**, in this URL.

Initially, the browser will warn you that you are using an untrusted certificate. Skip this warning and you will be presented with the login screen for the Fuse Console.

CHAPTER 4. SECURING THE CAMEL ACTIVEMQ COMPONENT

Abstract

The Camel ActiveMQ component enables you to define JMS endpoints in your routes that can connect to an Apache ActiveMQ broker. In order to make your Camel ActiveMQ endpoints secure, you must create an instance of a Camel ActiveMQ component that uses a **secure** connection factory.

4.1. SECURE ACTIVEMQ CONNECTION FACTORY

Overview

Apache Camel provides an Apache ActiveMQ component for defining Apache ActiveMQ endpoints in a route. The Apache ActiveMQ endpoints are effectively Java clients of the broker and you can either define a consumer endpoint (typically used at the start of a route to **poll for** JMS messages) or define a producer endpoint (typically used at the end or in the middle of a route to **send** JMS messages to a broker).

When the remote broker is secure (SSL security, JAAS security, or both), the Apache ActiveMQ component must be configured with the required client security settings.

Programming the security properties

Apache ActiveMQ enables you to program SSL security settings (and JAAS security settings) by creating and configuring an instance of the **`ActiveMQSslConnectionFactory`** JMS connection factory. Programming the JMS connection factory is the correct approach to use in the context of the containers such as OSGi, J2EE, Tomcat, and so on, because these settings are local to the application using the JMS connection factory instance.

NOTE

A standalone broker can configure SSL settings using **Java system properties**. For clients deployed in a container, however, this is **not** a practical approach, because the configuration must apply only to individual bundles, not the entire OSGi container. A Camel ActiveMQ endpoint is effectively a kind of Apache ActiveMQ Java client, so this restriction applies also to Camel ActiveMQ endpoints.

Defining a secure connection factory

Example 4.1, “Defining a Secure Connection Factory Bean” shows how to create a secure connection factory bean in Blueprint, enabling both SSL/TLS security **and** JAAS authentication.

Example 4.1. Defining a Secure Connection Factory Bean

```
<bean id="jmsConnectionFactory"
      class="org.apache.activemq.ActiveMQSslConnectionFactory">
  <property name="brokerURL" value="ssl://localhost:61617" />
  <property name="userName" value="Username"/>
  <property name="password" value="Password"/>
  <property name="trustStore" value="/conf/client.ts"/>
  <property name="trustStorePassword" value="password"/>
</bean>
```

The following properties are specified on the **ActiveMQSslConnectionFactory** class:

brokerURL

The URL of the remote broker to connect to, where this example connects to an SSL-enabled OpenWire port on the local host. The broker must also define a corresponding transport connector with compatible port settings.

userName and password

Any valid JAAS login credentials, ***Username*** and ***Password***.

trustStore

Location of the Java keystore file containing the certificate trust store for SSL connections. The location is specified as a classpath resource. If a relative path is specified, the resource location is relative to the **org/jbossfuse/example** directory on the classpath.

trustStorePassword

The password that unlocks the keystore file containing the trust store.

It is also possible to specify **keyStore** and **keyStorePassword** properties, but these would only be needed, if SSL mutual authentication is enabled (where the client presents an X.509 certificate to the broker during the SSL handshake).

4.2. EXAMPLE CAMEL ACTIVEMQ COMPONENT CONFIGURATION

Overview

This section describes how to initialize and configure a sample Camel ActiveMQ component instance, which you can then use to define ActiveMQ endpoints in a Camel route. This makes it possible for a Camel route to send or receive messages from a broker.

Prerequisites

The **camel-activemq** feature, which defines the bundles required for the Camel ActiveMQ component, is **not** installed by default. To install the **camel-activemq** feature, enter the following console command:

```
JBossFuse:karaf@root> features:install camel-activemq
```

Sample Camel ActiveMQ component

The following Blueprint sample shows a complete configuration of a Camel ActiveMQ component that has both SSL/TLS security and JAAS authentication enabled. The Camel ActiveMQ component instance is defined to with the **activemqssl** bean ID, which means it is associated with the **activemqssl** scheme (which you use when defining endpoints in a Camel route).

```
<?xml version="1.0" encoding="UTF-8"?>
<beans ... >
  ...
  <!--
    Configure the activemqssl component:
  -->
  <bean id="jmsConnectionFactory"
```

```
        class="org.apache.activemq.ActiveMQSslConnectionFactory">
    <property name="brokerURL" value="ssl://localhost:61617" />
    <property name="userName" value="Username"/>
    <property name="password" value="Password"/>
    <property name="trustStore" value="/conf/client.ts"/>
    <property name="trustStorePassword" value="password"/>
  </bean>

  <bean id="pooledConnectionFactory"
        class="org.apache.activemq.pool.PooledConnectionFactory">
    <property name="maxConnections" value="8" />
    <property name="maximumActive" value="500" />
    <property name="connectionFactory" ref="jmsConnectionFactory" />
  </bean>

  <bean id="jmsConfig"
class="org.apache.camel.component.jms.JmsConfiguration">
    <property name="connectionFactory" ref="pooledConnectionFactory"/>
    <property name="transacted" value="false"/>
    <property name="concurrentConsumers" value="10"/>
  </bean>

  <bean id="activemqssl"
        class="org.apache.activemq.camel.component.ActiveMQComponent">
    <property name="configuration" ref="jmsConfig"/>
  </bean>

</beans>
```

Sample Camel route

The following Camel route defines a sample endpoint that sends messages securely to the **`security.test`** queue on the broker, using the **`activemqssl`** scheme to reference the Camel ActiveMQ component defined in the preceding example:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans ...>
  ...
  <camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="timer://myTimer?fixedRate=true&period=5000"/>
      <transform><constant>Hello world!</constant></transform>
      <to uri="activemqssl:security.test"/>
    </route>
  </camelContext>
  ...
</beans>
```

CHAPTER 5. SECURING THE CAMEL CXF COMPONENT

Abstract

This chapter explains how to enable SSL/TLS security on a Camel CXF endpoint, using the Camel CXF proxy demonstration as the starting point. The Camel CXF component enables you to add Apache CXF endpoints to your Apache Camel routes. This makes it possible to simulate a Web service in Apache Camel or you could interpose a route between a WS client and a Web service to perform additional processing (which is the case considered here).

5.1. THE CAMEL CXF PROXY DEMONSTRATION

Overview

In order to explain how to secure a Camel CXF endpoint in OSGi, this tutorial builds on an example available from the standalone distribution of Apache Camel, the **Camel CXF proxy** demonstration. Figure 5.1, “Camel CXF Proxy Overview” gives an overview of how this demonstration works

Figure 5.1. Camel CXF Proxy Overview

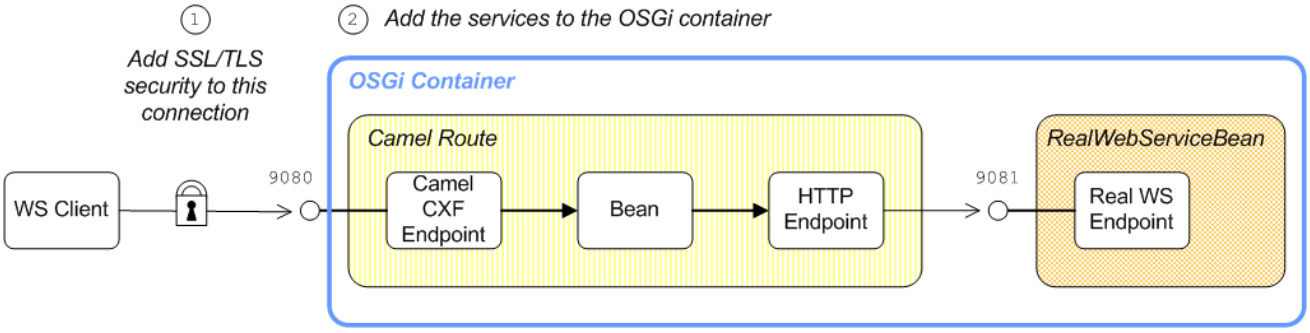

The **report incident** Web service, which is implemented by the **`RealWebServiceBean`**, receives details of an incident (for example, a traffic accident) and returns a tracking code to the client. Instead of sending its requests directly to the real Web service, however, the WS client connects to a Camel CXF endpoint, which is interposed between the WS client and the real Web service. The Apache Camel route performs some processing on the WSDL message (using the **`enrichBean`**) before forwarding it to the real Web service.

WARNING

If you enable SSL/TLS security, you must ensure that you explicitly disable the SSLv3 protocol, in order to safeguard against the Poodle vulnerability (CVE-2014-3566). For more details, see Disabling SSLv3 in JBoss Fuse 6.x and JBoss A-MQ 6.x.

Modifications

In order to demonstrate how to enable SSL/TLS on a Camel CXF endpoint in the context of OSGi, this chapter contains instructions on how to modify the basic demonstration as follows:

1. SSL/TLS security is enabled on the connection between the WS client and the Camel CXF endpoint.
2. The Apache Camel route and the **RealWebServiceBean** bean are both deployed into the OSGi container.

Obtaining the demonstration code

The Camel CXF proxy demonstration is available only from the standalone distribution of Apache Camel, which is included in the ***InstallDir*/extras** directory. Using a standard archive utility, expand the Camel archive file and extract the contents to a convenient location on your filesystem.

Assuming that you have installed Apache Camel in *CamelInstallDir*, you can find the Camel CXF proxy demonstration in the following directory:

```
CamelInstallDir/examples/camel-example-cxf-proxy
```

Obtaining the sample certificates

This demonstration needs X.509 certificates. In a real deployment, you should generate these certificates yourself using a private certificate authority. For this demonstration, however, we use some sample certificates from the Apache CXF **wsdl_first_http** example. This demonstration is available from the standalone distribution of Apache CXF, which is included in the ***InstallDir*/extras** directory. Using a standard archive utility, expand the CXF archive file and extract the contents to a convenient location on your filesystem.

Assuming that you have installed Apache CXF in *CXFInstallDir*, you can find the **wsdl_first_http** demonstration in the following directory:

```
CXFInstallDir/samples/wsdl_first_http
```

Physical part of the WSDL contract

The physical part of the WSDL contract refers to the **wsdl:service** and **wsdl:port** elements. These elements specify the transport details that are needed to connect to a specific Web services endpoint. For the purposes of this demonstration, this is the most interesting part of the contract and it is shown in Example 5.1, “The ReportIncidentEndpointService WSDL Service”.

Example 5.1. The ReportIncidentEndpointService WSDL Service

```
<wsdl:definitions xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
    ...
 xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
 targetNamespace="http://reportincident.example.camel.apache.org">
    ...
    <!-- Service definition -->
    <wsdl:service name="ReportIncidentEndpointService">
        <wsdl:port name="ReportIncidentEndpoint"
binding="tns:ReportIncidentBinding">
            <soap:address location="http://localhost:9080/camel-
example-cxf-proxy/webservices/incident"/>
        </wsdl:port>
```

```
    </wsdl:service>

</wsdl:definitions>
```

NOTE

The address URL appearing in the WSDL contract (the value of the **`soap:address`** element's **`location`** attribute) is not important here, because the application code overrides the default value of the address URL.

WSDL addressing details

A WS client needs three pieces of information to connect to a WSDL service: the *WSDL service name*, the *WSDL port name*, and the *address URL* of the Web service. The following addressing details are used to connect to the proxy Web service and to the real Web service in this example:

WSDL service name

The full QName of the WSDL service is as follows:

```
{http://reportincident.example.camel.apache.org}ReportIncidentEndpointSe
rvice
```

WSDL port name

The full QName of the WSDL port is as follows:

```
{http://reportincident.example.camel.apache.org}ReportIncidentEndpoint
```

Address URL

The address URL of the **proxy Web service** endpoint (which uses the HTTPS protocol) is as follows:

```
https://localhost:9080/camel-example-cxf-proxy/webservices/incident
```

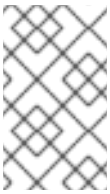

NOTE

The preceding address is specified when the **`reportIncident`** bean is created using a **`cxf:cxfEndpoint`** element in the bundle's Spring configuration file, **`src/main/resources/META-INF/spring/camel-config.xml`**.

The address URL of the **real Web service** endpoint (using the HTTP protocol) is as follows:

```
http://localhost:9081/real-webservice
```

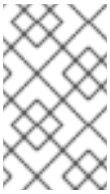

NOTE

The preceding address is specified when the **`realWebService`** bean is created in the bundle's Spring configuration file, **`src/main/resources/META-INF/spring/camel-config.xml`**.

5.2. SECURING THE WEB SERVICES PROXY

Overview

This section explains how to enable SSL/TLS security on the Camel CXF endpoint, which acts as a proxy for the real Web service. Assuming that you already have the X.509 certificates available, all that is required is to add a block of configuration data to the Spring configuration file (where the configuration data is contained in a **`httpj:engine-factory`** element). There is just one slightly subtle aspect to this, however: you need to understand how the Camel CXF endpoint gets associated with the SSL/TLS configuration details.

Implicit configuration

A WS endpoint can be configured by creating the endpoint in Spring and then configuring SSL/TLS properties on its Jetty container. The configuration can be somewhat confusing, however, for the following reason: the Jetty container (which is configured by a **`httpj:engine-factory`** element in Spring) **does not explicitly reference the WS endpoints it contains** and the WS endpoints **do not explicitly reference the Jetty container** either. The connection between the Jetty container and its contained endpoints is established implicitly, in that they are both configured to use the same TCP port, as illustrated by WS Endpoint Implicitly Configured by httpj:engine-factory.

WS Endpoint Implicitly Configured by httpj:engine-factory

```
Element
```

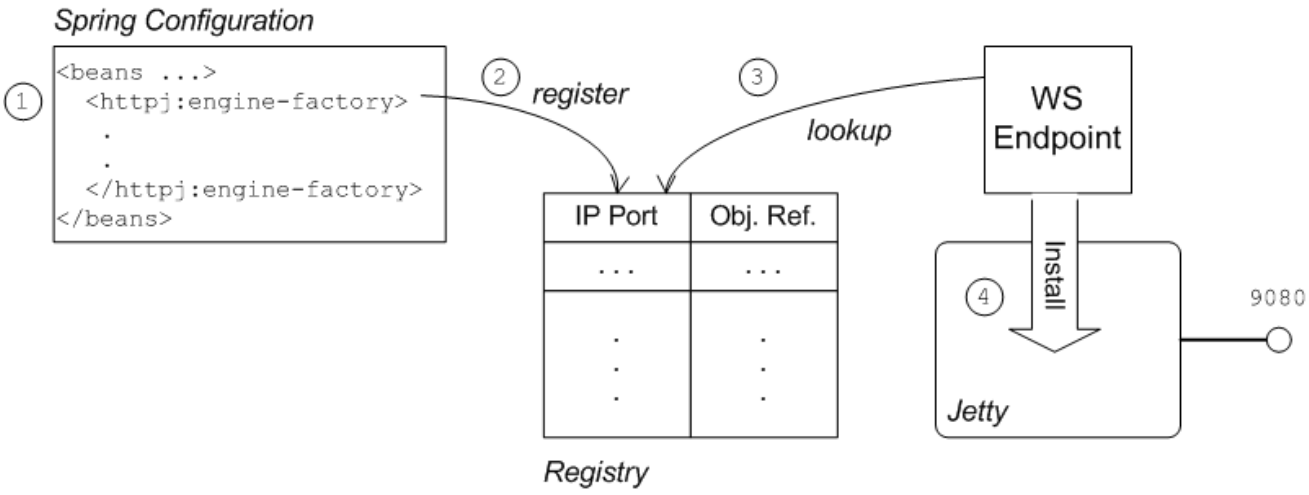

The connection between the Web service endpoint and the **`httpj:engine-factory`** element is established as follows:

1. The Spring container loads and parses the file containing the **`httpj:engine-factory`** element.
2. When the **`httpj:engine-factory`** bean is created, a corresponding entry is created in the registry, storing a reference to the bean. The **`httpj:engine-factory`** bean is also used to initialize a Jetty container that listens on the specified TCP port.
3. When the WS endpoint is created, it scans the registry to see if it can find a **`httpj:engine-factory`** bean with the same TCP port as the TCP port in the endpoint's address URL.

4. If one of the beans matches the endpoint's TCP port, the WS endpoint installs itself into the corresponding Jetty container. If the Jetty container has SSL/TLS enabled, the WS endpoint shares those security settings.

Steps to add SSL/TLS security to the Jetty container

To add SSL/TLS security to the Jetty container, thereby securing the WS proxy endpoint, perform the following steps:

1. the section called "Add certificates to the bundle resources".
2. the section called "Modify POM to switch off resource filtering".
3. the section called "Instantiate the CXF Bus".
4. the section called "Add the httpj:engine-factory element to Spring".
5. the section called "Define the cxfcore:, sec: and httpj: prefixes".
6. the section called "Modify proxy address URL to use HTTPS".

Add certificates to the bundle resources

The certificates used in this demonstration are taken from a sample in the Apache CXF 3.1.11.fuse-000243-redhat-1 product. If you install the standalone version of Apache CXF (available in the ***InstallDir*/extras/** directory), you will find the sample certificates in the ***CXFInstallDir*/samples/wsdl_first_https/src/main/config** directory.

Copy the **clientKeystore.jks** and **serviceKeystore.jks** keystores from the ***CXFInstallDir*/samples/wsdl_first_https/src/main/config** directory to the ***CamelInstallDir*/examples/camel-example-cxf-proxy/src/main/resources/certs** directory (you must first create the **certs** sub-directory).

Modify POM to switch off resource filtering

Including the certificates directly in the bundle as resource is the most convenient way to deploy them. But when you deploy certificates as resources in a Maven project, you must remember to disable Maven resource filtering, which corrupts binary files.

To disable filtering of **.jks** files in Maven, open the project POM file, ***CamelInstallDir*/examples/camel-example-cxf-proxy/pom.xml**, with a text editor and add the following **resources** element as a child of the **build** element:

```
<?xml version="1.0" encoding="UTF-8"?>
...
<project ...>
  ...
  <build>
    <plugins>
      ...
    </plugins>

    <resources> <resource> <directory>src/main/resources</directory>
<filtering>true</filtering> <excludes> <exclude>/.jks</exclude>
</excludes> </resource> <resource>
```

```
<directory>src/main/resources</directory> <filtering>false</filtering>
<includes> <include>/.jks</include> </includes> </resource> </resources>
  </build>

</project>
```

Instantiate the CXF Bus

You should instantiate the CXF bus explicitly in the Spring XML (this ensures that it will be available to the Jetty container, which is instantiated by the **`httpj:engine-factory`** element in the next step). Edit the **`camel-config.xml`** file in the **`src/main/resources/META-INF/spring`** directory, adding the **`cxfcore:bus`** element as a child of the **`beans`** element, as follows:

```
<beans ... >
    ...
    <cxfcore:bus/>
    ...
</beans>
```

NOTE

The **`cxfcore:`** namespace prefix will be defined in a later step.

Add the httpj:engine-factory element to Spring

```
configuration
```

To configure the Jetty container that listens on TCP port 9080 to use SSL/TLS security, edit the **`camel-config.xml`** file in the **`src/main/resources/META-INF/spring`** directory, adding the **`httpj:engine-factory`** element as shown in Example 5.2, “httpj:engine-factory Element with SSL/TLS Enabled”.

In this example, the **`required`** attribute of the **`sec:clientAuthentication`** element is set to **`false`**, which means that a connecting client is **not** required to present an X.509 certificate to the server during the SSL/TLS handshake (although it may do so, if it has such a certificate).

Example 5.2. httpj:engine-factory Element with SSL/TLS Enabled

```
<beans ... >
    ...
    <httpj:engine-factory bus="cxf">
      <httpj:engine port="${proxy.port}">
        <httpj:tlsServerParameters secureSocketProtocol="TLSv1">
          <sec:keyManagers keyPassword="skpass">
            <sec:keyStore resource="certs/serviceKeystore.jks"
password="sspass" type="JKS"/>
          </sec:keyManagers>
          <sec:trustManagers>
            <sec:keyStore resource="certs/serviceKeystore.jks"
password="sspass" type="JKS"/>
          </sec:trustManagers>
          <sec:cipherSuitesFilter>
            <sec:include>.*_WITH_3DES_.*</sec:include>
```

```
            <sec:include>.*_WITH_DES_.*</sec:include>
            <sec:exclude>.*_WITH_NULL_.*</sec:exclude>
            <sec:exclude>.*_DH_anon_.*</sec:exclude>
          </sec:cipherSuitesFilter>
          <sec:clientAuthentication want="true" required="false"/>
        </httpj:tlsServerParameters>
      </httpj:engine>
    </httpj:engine-factory>

</beans>
```

IMPORTANT

You must set secureSocketProtocol to **`TLSv1`** on the server side, in order to protect against the Poodle vulnerability (CVE-2014-3566)

Define the cxfcore:, sec: and httpj: prefixes

Define the **`cxfcore:`**, **`sec:`** and **`httpj:`** namespace prefixes, which appear in the definitions of the **`cxfcore:bus`** element and the **`httpj:engine-factory`** element, by adding the following highlighted lines to the **`beans`** element in the **`camel-config.xml`** file:

```
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:camel="http://camel.apache.org/schema/spring"
       xmlns:cxf="http://camel.apache.org/schema/cxf"
       xmlns:context="http://www.springframework.org/schema/context"
       xmlns:cxfcore="http://cxf.apache.org/core"
       xmlns:sec="http://cxf.apache.org/configuration/security"
       xmlns:httpj="http://cxf.apache.org/transports/http-
jetty/configuration"
       xsi:schemaLocation="
       http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd
       http://camel.apache.org/schema/spring
http://camel.apache.org/schema/spring/camel-spring.xsd
       http://camel.apache.org/schema/cxf
http://camel.apache.org/schema/cxf/camel-cxf.xsd
       http://www.springframework.org/schema/context
http://www.springframework.org/schema/context/spring-context.xsd
       http://cxf.apache.org/core http://cxf.apache.org/schemas/core.xsd
       http://cxf.apache.org/configuration/security
http://cxf.apache.org/schemas/configuration/security.xsd
       http://cxf.apache.org/transports/http-jetty/configuration
http://cxf.apache.org/schemas/configuration/http-jetty.xsd
       ">
```

NOTE

It is essential to specify the locations of the **http://cxf.apache.org/configuration/security** schema and the **http://cxf.apache.org/transports/http-jetty/configuration** schema in the **xsi:schemaLocation** attribute. These will not automatically be provided by the OSGi container.

Modify proxy address URL to use HTTPS

The proxy endpoint at the start of the Apache Camel route is configured by the **cxf:cxfEndpoint** element in the **camel-config.xml** file. By default, this proxy endpoint is configured to use the HTTP protocol. You must modify the address URL to use the secure HTTPS protocol instead, however. In the **camel-config.xml** file, edit the address attribute of the **cxf:cxfEndpoint** element, replacing the **http:** prefix by the **https:** prefix, as shown in the following fragment:

```
<beans ...>
    ...
    <cxf:cxfEndpoint id="reportIncident"
                     address="https://localhost:${proxy.port}/camel-
example-cxf-proxy/webservices/incident"
                     endpointName="s:ReportIncidentEndpoint"
                     serviceName="s:ReportIncidentEndpointService"
                     wsdlURL="etc/report_incident.wsdl"

xmlns:s="http://reportincident.example.camel.apache.org"/>
    ...
</beans>
```

Notice also that the address URL is configured to use the TCP port, **${proxy.port}** (which has the value **9080** by default). This TCP port value is the same as the value set for the Jetty container (configured by the **http:engine-factory** element), thus ensuring that this endpoint is deployed into the Jetty container. The attributes of the **cxf:cxfEndpoint** specify the WSDL addressing details as described in the section called “WSDL addressing details”:

serviceName

Specifies the WSDL service name.

endpointName

Specifies the WSDL port name.

address

Specifies the address URL of the proxy Web service.

5.3. DEPLOYING THE APACHE CAMEL ROUTE

Overview

The Maven POM file in the basic Camel CXF proxy demonstration is already configured to generate an OSGi bundle. Hence, after building the demonstration using Maven, the demonstration bundle (which contains the Apache Camel route and the **RealWebServicesBean** bean) is ready for deployment into the OSGi container.

Prerequisites

Before deploying the Apache Camel route into the OSGi container, you must configure the proxy Web service to use SSL/TLS security, as described in the previous section, Section 5.2, “Securing the Web Services Proxy”.

Steps to deploy the Camel route

To deploy the Web services proxy demonstration into the OSGi container, perform the following steps:

1. the section called “Build the demonstration”.
2. the section called “Start the OSGi container”.
3. the section called “Install the required features”.
4. the section called “Deploy the bundle”.
5. the section called “Check the console output”.

Build the demonstration

Use Maven to build and install the demonstration as an OSGi bundle. Open a command prompt, switch the current directory to ***CamelInstallDir*/examples/camel-example-cxf-proxy**, and enter the following command:

```
mvn install -Dmaven.test.skip=true
```

Start the OSGi container

If you have not already done so, start up the Karaf console (and container instance) by entering the following command in a new command prompt:

```
./fuse
```

Install the required features

The **camel-cxf** feature, which defines the bundles required for the Camel/CXF component, is**not** installed by default. To install the **camel-cxf** feature, enter the following console command:

```
JBossFuse:karaf@root> features:install camel-cxf
```

You also need the **camel-http** feature, which defines the bundles required for the Camel/HTTP component. To install the **camel-http** feature, enter the following console command:

```
JBossFuse:karaf@root> features:install camel-http
```

Deploy the bundle

Deploy the **camel-example-cxf-proxy** bundle, by entering the following console command:

```
JBossFuse:karaf@root> install -s mvn:org.apache.camel/camel-example-cxf-
proxy/2.21.0.fuse-000077-redhat-1
```

NOTE

In this case, it is preferable to deploy the bundle directly using **`install`**, rather than using hot deploy, so that you can see the bundle output on the console screen.

If you have any difficulty using the **`mvn`** URL handler, see olink:ESBOSGiGuide/UrlHandlers-Maven for details of how to set it up.

Check the console output

After the bundle is successfully deployed, you should see output like the following in the console window:

```
JBossFuse:karaf@root> Starting real web service...
Started real web service at: http://localhost:9081/real-webservice
```

5.4. SECURING THE WEB SERVICES CLIENT

Overview

In the basic Camel CXF proxy demonstration, the Web services client is actually implemented as a JUnit test under the **`src/test`** directory. This means that the client can easily be run using the Maven command, **`mvn test`**. To enable SSL/TLS security on the client, the Java implementation of the test client is completely replaced and a Spring file, containing the SSL/TLS configuration, is added to the **`src/test/resources/META-INF/spring`** directory. Before describing the steps you need to perform to set up the client, this section explains some details of the client's Java code and Spring configuration.

Implicit configuration

Apart from changing the URL scheme on the endpoint address to **`https:`**, most of the configuration to enable SSL/TLS security on a client proxy is contained in a **`http:conduit`** element in Spring configuration. The way in which this configuration is applied to the client proxy, however, is potentially confusing, for the following reason: the **`http:conduit`** element **does not explicitly reference the client proxy** and the client proxy **does not explicitly reference the `http:conduit` element**. The connection between the **`http:conduit`** element and the client proxy is established implicitly, in that they both reference the same WSDL port, as illustrated by Client Proxy Implicitly Configured by http:conduit.

Client Proxy Implicitly Configured by http:conduit

```
Element
```

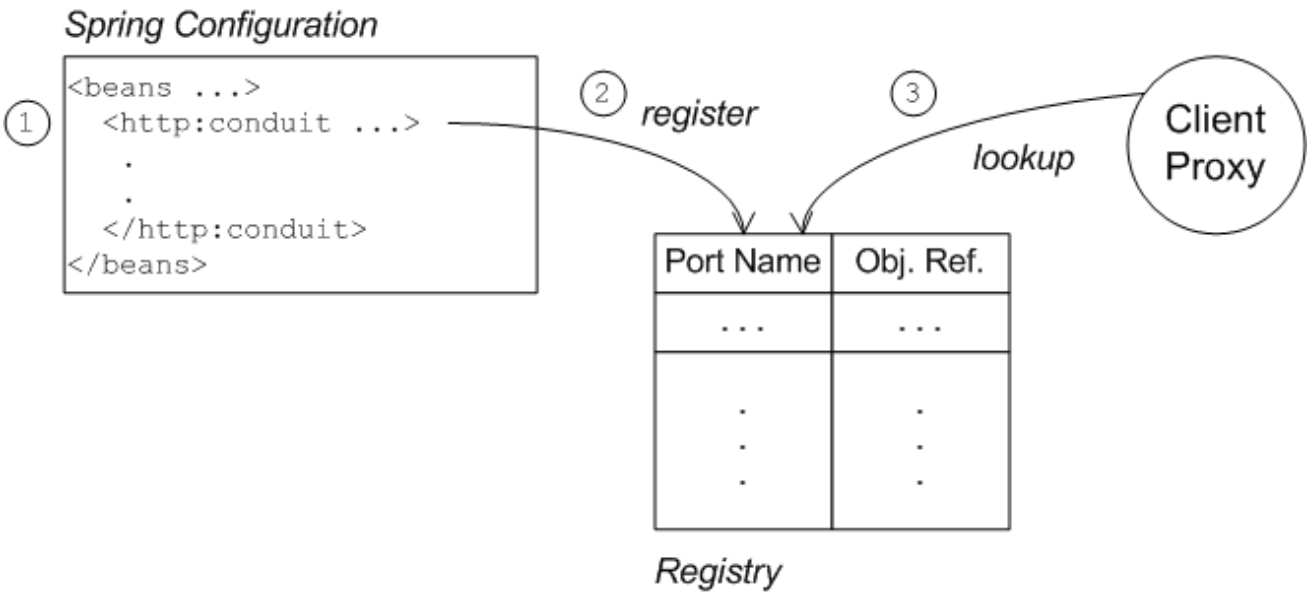

The connection between the client proxy and the **`http:conduit`** element is established as follows:

1. The client loads and parses the Spring configuration file containing the **`http:conduit`** element.
2. When the **`http:conduit`** bean is created, a corresponding entry is created in the registry, which stores a reference to the bean under the specified WSDL port name (where the name is stored in QName format).
3. When the JAX-WS client proxy is created, it scans the registry to see if it can find a **`http:conduit`** bean associated with the proxy's WSDL port name. If it finds such a bean, it automatically injects the configuration details into the proxy.

Certificates needed on the client side

The client is configured with the following **`clientKeystore.jks`** keystore file from the **`src/main/resources/certs`** directory. This keystore contains two entries, as follows:

Trusted cert entry

A trusted certificate entry containing the CA certificate that issued and signed both the server certificate and the client certificate.

Private key entry

A private key entry containing the client's own X.509 certificate and private key. In fact, this certificate is not strictly necessary to run the current example, because the server does not require the client to send a certificate during the TLS handshake (see Example 5.2, "httpj:engine-factory Element with SSL/TLS Enabled").

Loading Spring definitions into the client

The example client is not deployed directly into a Spring container, but it requires some Spring definitions in order to define a secure HTTP conduit. So how can you create the Spring definitions without a Spring container? It turns out that it is easy to read Spring definitions into a Java-based client using the **`org.apache.cxf.bus.spring.SpringBusFactory`** class.

The following code shows how to read Spring definitions from the file, **`META-INF/spring/cxf-client.xml`**, and create an Apache CXF Bus object that incorporates those definitions:

```
// Java
import org.apache.cxf.bus.spring.SpringBusFactory;
...
protected void startCxfBus() throws Exception {
    bf = new SpringBusFactory();
    Bus bus = bf.createBus("META-INF/spring/cxf-client.xml");
    bf.setDefaultBus(bus);
}
```

Creating the client proxy

In principle, there are several different ways of creating a WSDL proxy: you could use the JAX-WS API to create a proxy based on the contents of a WSDL file; you could use the JAX-WS API to create a proxy **without** a WSDL file; or you could use the Apache CXF-specific class, **`JaxWsProxyFactoryBean`**, to create a proxy.

For this SSL/TLS client, the most convenient approach is to use the JAX-WS API to create a proxy without using a WSDL file, as shown in the following Java sample:

```
// Java
import javax.xml.ws.Service;
import org.apache.camel.example.reportincident.ReportIncidentEndpoint;
...
// create the webservice client and send the request
Service s = Service.create(SERVICE_NAME);
s.addPort(
    PORT_NAME,
    "http://schemas.xmlsoap.org/soap/",
    ADDRESS_URL
  );
ReportIncidentEndpoint client =
  s.getPort(PORT_NAME, ReportIncidentEndpoint.class);
```

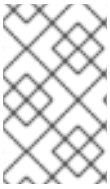

NOTE

In this example, you **cannot** use the **`JaxWsProxyFactoryBean`** approach to create a proxy, because a proxy created in this way fails to find the HTTP conduit settings specified in the Spring configuration file.

The **`SERVICE_NAME`** and **`PORT_NAME`** constants are the QNames of the WSDL service and the WSDL port respectively, as defined in Example 5.1, “The ReportIncidentEndpointService WSDL Service”. The **`ADDRESS_URL`** string has the same value as the proxy Web service address and is defined as follows:

```
private static final String ADDRESS_URL =
  "https://localhost:9080/camel-example-cxf-proxy/webservices/incident";
```

In particular, note that the address **must** be defined with the URL scheme, **`https`**, which selects HTTP over SSL/TLS.

Steps to add SSL/TLS security to the client

To define a JAX-WS client with SSL/TLS security enabled, perform the following steps:

1. the section called “Create the Java client as a test case”.
2. the section called “Add the http:conduit element to Spring configuration”.
3. the section called “Run the client”.

Create the Java client as a test case

Example 5.3, “ReportIncidentRoutesTest Java client” shows the complete code for a Java client that is implemented as a JUnit test case. This client replaces the existing test, **`ReportIncidentRoutesTest.java`**, in the **`src/test/java/org/apache/camel/example/reportincident`** sub-directory of the **`examples/camel-example-cxf-proxy`** demonstration.

To add the client to the ***`CamelInstallDir`*`/examples/camel-example-cxf-proxy`** demonstration, go to the **`src/test/java/org/apache/camel/example/reportincident`** sub-directory, move the existing **`ReportIncidentRoutesTest.java`** file to a backup location, then create a new **`ReportIncidentRoutesTest.java`** file and paste the code from Example 5.3, “ReportIncidentRoutesTest Java client” into this file.

Example 5.3. ReportIncidentRoutesTest Java client

```
// Java
package org.apache.camel.example.reportincident;

import org.apache.camel.spring.Main;
import org.apache.cxf.jaxws.JaxWsProxyFactoryBean;
import org.junit.Test;

import java.net.URL;
import javax.xml.namespace.QName;
import javax.xml.ws.Service;

import org.apache.cxf.Bus;
import org.apache.cxf.bus.spring.SpringBusFactory;
import org.apache.camel.example.reportincident.ReportIncidentEndpoint;
import
org.apache.camel.example.reportincident.ReportIncidentEndpointService;

import static org.junit.Assert.assertEquals;

/**
 * Unit test of our routes
 */
public class ReportIncidentRoutesTest {

    private static final QName SERVICE_NAME
        = new QName("http://reportincident.example.camel.apache.org",
"ReportIncidentEndpointService");

    private static final QName PORT_NAME =
        new QName("http://reportincident.example.camel.apache.org",
"ReportIncidentEndpoint");

    private static final String WSDL_URL =
```

```
"file:src/main/resources/etc/report_incident.wsdl";

    // should be the same address as we have in our route
    private static final String ADDRESS_URL =
"https://localhost:9080/camel-example-cxf-proxy/webservices/incident";

    protected SpringBusFactory bf;

    protected void startCxfBus() throws Exception {
        bf = new SpringBusFactory();
        Bus bus = bf.createBus("META-INF/spring/cxf-client.xml");
        bf.setDefaultBus(bus);
    }

    @Test
    public void testRendportIncident() throws Exception {
        startCxfBus();
        runTest();
    }

    protected void runTest() throws Exception {

        // create input parameter
        InputReportIncident input = new InputReportIncident();
        input.setIncidentId("123");
        input.setIncidentDate("2008-08-18");
        input.setGivenName("Claus");
        input.setFamilyName("Ibsen");
        input.setSummary("Bla");
        input.setDetails("Bla bla");
        input.setEmail("davsclaus@apache.org");
        input.setPhone("0045 2962 7576");

        // create the webservice client and send the request
        Service s = Service.create(SERVICE_NAME);
        s.addPort(PORT_NAME, "http://schemas.xmlsoap.org/soap/",
ADDRESS_URL);
        ReportIncidentEndpoint client = s.getPort(PORT_NAME,
ReportIncidentEndpoint.class);

        OutputReportIncident out = client.reportIncident(input);

        // assert we got a OK back
        assertEquals("OK;456", out.getCode());
    }
}
```

Add the http:conduit element to Spring configuration

Example 5.4, “http:conduit Element with SSL/TLS Enabled” shows the Spring configuration that defines a **`http:conduit`** element for the **`ReportIncidentEndpoint`** WSDL port. The **`http:conduit`** element is configured to enable SSL/TLS security for any client proxies that use the specified WSDL port.

To add the Spring configuration to the client test case, create the **src/test/resources/META-INF/spring** sub-directory, use your favorite text editor to create the file, **cxf-client.xml**, and then paste the contents of Example 5.4, “http:conduit Element with SSL/TLS Enabled” into the file.

Example 5.4. http:conduit Element with SSL/TLS Enabled

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:cxf="http://camel.apache.org/schema/cxf"
       xmlns:sec="http://cxf.apache.org/configuration/security"
       xmlns:http="http://cxf.apache.org/transports/http/configuration"
       xsi:schemaLocation="
       http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd
       http://camel.apache.org/schema/cxf
http://camel.apache.org/schema/cxf/camel-cxf.xsd
       http://cxf.apache.org/configuration/security
http://cxf.apache.org/schemas/configuration/security.xsd
       http://cxf.apache.org/transports/http/configuration
http://cxf.apache.org/schemas/configuration/http-conf.xsd
       ">

  <http:conduit name="
{http://reportincident.example.camel.apache.org}ReportIncidentEndpoint.h
ttp-conduit">
    <http:tlsClientParameters disableCNCheck="true"
secureSocketProtocol="TLSv1">
      <sec:keyManagers keyPassword="ckpass">
          <sec:keyStore password="cspass" type="JKS"
          resource="certs/clientKeystore.jks" />
      </sec:keyManagers>
      <sec:trustManagers>
          <sec:keyStore password="cspass" type="JKS"
          resource="certs/clientKeystore.jks" />
      </sec:trustManagers>
      <sec:cipherSuitesFilter>
        <sec:include>.*_WITH_3DES_.*</sec:include>
        <sec:include>.*_WITH_DES_.*</sec:include>
        <sec:exclude>.*WITH_NULL.</sec:exclude>*
        <sec:exclude>.*DH_anon.</sec:exclude>*
      </sec:cipherSuitesFilter>
    </http:tlsClientParameters>
   </http:conduit>

</beans>
```

Please note the following points about the preceding configuration:

- The **http:** and **sec:** namespace prefixes are needed to define the **http:conduit** element. In the **xsi:schemaLocation** element, it is also essential to specify the locations of the corresponding **http://cxf.apache.org/configuration/security** and **http://cxf.apache.org/transports/http/configuration** namespaces.

- The **disableCNCheck** attribute of the **http:tlsClientParameters** element is set to **true**. This means that the client does **not** check whether the Common Name in the server's X.509 certificate matches the server hostname. For more details, see Appendix A, *Managing Certificates*.

IMPORTANT

Disabling the CN check is **not** recommended in a production deployment.

- In the **sec:keystore** elements, the certificate locations are specified using the **resource** attribute, which finds the certificates on the classpath. When Maven runs the test, it automatically makes the contents of **src/main/resources** available on the classpath, so that the certificates can be read from the **src/main/resources/certs** directory.

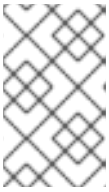

NOTE

You also have the option of specifying a certificate location using the **file** attribute, which looks in the filesystem. But the **resource** attribute is more suitable for use with applications packaged in bundles.

- The **sec:cipherSuitesFilter** element is configured to exclude cipher suites matching **.*WITH_NULL.*** and **.*DH_anon.***. These cipher suites are effectively incomplete and are **not** intended for normal use.

IMPORTANT

It is recommended that you always **exclude** the ciphers matching **.*WITH_NULL.*** and **.*DH_anon.***.

- The **secureSocketProtocol** attribute should be set to TLSv1, to match the server protocol and to ensure that the SSLv3 protocol is not used (POODLE security vulnerability (CVE-2014-3566)).

Run the client

Because the client is defined as a test case, you can run the client using the standard Maven test goal. To run the client, open a new command window, change directory to ***CamelInstallDir*/examples/camel-example-cxf-proxy**, and enter the following Maven command:

```
mvn test
```

If the test runs successfully, you should see the following output in the OSGi console window:

```
Incident was 123, changed to 456

Invoked real web service: id=456 by Claus Ibsen
```

CHAPTER 6. SECURING THE MANAGEMENT CONSOLE

Abstract

The default setting for **Access-Control-Allow-Origin** header for the Fuse Management Console permits unrestricted sharing. To restrict access to the Fuse Management Console, create an access management file which contains a list of the allowed origin URLs. To implement the restrictions, add a system property that references the access management file

6.1. CONTROLLING ACCESS TO THE FUSE MANAGEMENT CONSOLE

Create an access management file called **access-management.xml** in ***<installDir>/etc/***. The access management file must contain **<allow-origin>** sections within a **<cors>** section. The **<allow-origin>** section can contain the origin URL provided by browsers with the **Origin:** header, or a wildcard specification with *. For example:

```
<cors>
   <!-- Allow cross origin access from www.jolokia.org ... -->
   <allow-origin>http://www.jolokia.org</allow-origin>
   <!-- ... and all servers from jmx4perl.org with any protocol -->
   <allow-origin>*://*.jmx4perl.org</allow-origin>
   <!-- optionally allow access to web console from localhost -->
   <allow-origin>http://localhost:8181/*</allow-origin>
   <!-- Check for the proper origin on the server side, too -->
   <strict-checking/>
</cors>
```

Add the following line to Fuse config script **./bin/setenv**, adding the path to the access management file.

```
export EXTRA_JAVA_OPTS='-Djolokia.policyLocation=file:etc/access-
management.xml'
```

When the command **./bin/fuse** is executed, the access management file is referenced and used to restrict access to the Fuse Management Console.

CHAPTER 7. INTEGRATION WITH RED HAT SINGLE SIGN-ON

Red Hat provides a single sign-on option (Red Hat Single Sign-On) that works with JAAS to provide enterprise security for certain Web client applications and services running inside Fuse and Fuse administration services (SSH, JMX, and Fuse Management Console).

Adapters are provided for the following types of container in Red Hat Fuse:

- Section 7.1, “Adapter for Spring Boot container”
- Section 7.2, “Adapter for Apache Karaf container”
- Section 7.3, “Adapter for JBoss EAP container”

7.1. ADAPTER FOR SPRING BOOT CONTAINER

The adapter for the Spring Boot container supports the following embedded Web containers:

- Undertow
- Jetty
- Tomcat

For details on installing and using the Red Hat Single Sign-On adapter for the Spring Boot container, see Spring Boot Adapter in the Red Hat Single Sign-On *Securing Applications and Services Guide*.

7.2. ADAPTER FOR APACHE KARAF CONTAINER

The adapter for the Apache Karaf container can secure the following components:

- Classic WAR applications deployed on Fuse with Pax Web War Extender.
- Servlets deployed on Fuse as OSGI services with Pax Web Whiteboard Extender and additionally servlets registered through **`org.osgi.service.http.HttpService#registerServlet()`** which is a standard OSGi Enterprise HTTP Service.
- Apache Camel Undertow endpoints running with the Camel Undertow component.
- Apache CXF endpoints running on their own separate Undertow engine.
- Apache CXF endpoints running on the default engine provided by the CXF servlet.
- SSH and JMX admin access.
- Hawtio administration console.

For details on installing and using the Red Hat Single Sign-On adapter for the Apache Karaf container, see JBoss Fuse 7 Adapter in the Red Hat Single Sign-On *Securing Applications and Services Guide*.

7.3. ADAPTER FOR JBOSS EAP CONTAINER

The adapter for the JBoss Enterprise Application Platform (EAP) container provides security for WARs, enabling you to define role-based security constraints on your URLs.

For details on installing and using the Red Hat Single Sign-On adapter for the JBoss EAP container, see JBoss EAP Adapter in the Red Hat Single Sign-On *Securing Applications and Services Guide*.

CHAPTER 8. LDAP AUTHENTICATION TUTORIAL

Abstract

This tutorial explains how to set up an X.500 directory server and configure the OSGi container to use LDAP authentication.

8.1. TUTORIAL OVERVIEW

Goals

In this tutorial you will:

- Install 389 Directory Server
- Add user entries to the LDAP server
- Add groups to manage security roles
- Configure Fuse to use LDAP authentication
- Configure Fuse to use roles for authorization
- Configure SSL/TLS connections to the LDAP server

8.2. SET-UP A DIRECTORY SERVER AND CONSOLE

This stage of the tutorial explains how to install the X.500 directory server and the management console from the Fedora 389 Directory Server project. If you already have access to a 389 Directory Server instance, you can skip the instructions for installing the 389 Directory Server and install the 389 Management Console instead.

Prerequisites

If you are installing on a Red Hat Enterprise Linux platform, you must first install the Extra Packages for Enterprise Linux (EPEL). See the installation notes under RHEL/Cent OS/ EPEL (RHEL 6, RHEL 7, Cent OS 6, Cent OSý7) on the **`fedoraproject.org`** site.

Install 389 Directory Server

If you do not have access to an existing **389 Directory Server** instance, you can install **389 Directory Server** on your local machine, as follows:

1. On Red Hat Enterprise Linux and Fedora platforms, use the standard **`dnf`** package management utility to install **389 Directory Server**. Enter the following command at a command prompt (you must have administrator privileges on your machine):

   ```
   sudo dnf install 389-ds
   ```

NOTE

The required **389-ds** and **389-console** RPM packages are available for Fedora, RHEL6+EPEL, and CentOS7+EPEL platforms. At the time of writing, the **389-console** package is not yet available for RHEL 7.

2. After installing the 389 directory server packages, enter the following command to configure the directory server:

```
sudo setup-ds-admin.pl
```

   The script is interactive and prompts you to provide the basic configuration settings for the 389 directory server. When the script is complete, it automatically launches the 389 directory server in the background.

3. For more details about how to install **389 Directory Server**, see the Download page.

Install 389 Management Console

If you already have access to a **389 Directory Server** instance, you only need to install the 389 Management Console, which enables you to log in and manage the server remotely. You can install the 389 Management Console, as follows:

- **On Red Hat Enterprise Linux and Fedora platforms**—use the standard **dnf** package management utility to install the 389 Management Console. Enter the following command at a command prompt (you must have administrator privileges on your machine):

```
sudo dnf install 389-console
```

- **On Windows platforms**—see the Windows Console download instructions from **fedoraproject.org**.

Connect the console to the server

To connect the 389 Directory Server Console to the LDAP server:

1. Enter the following command to start up the 389 Management Console:

```
389-console
```

2. A login dialog appears. Fill in the LDAP login credentials in the **User ID** and **Password** fields, and customize the hostname in the **Administration URL** field to connect to your 389 management server instance (port **9830** is the default port for the 389 management server instance).

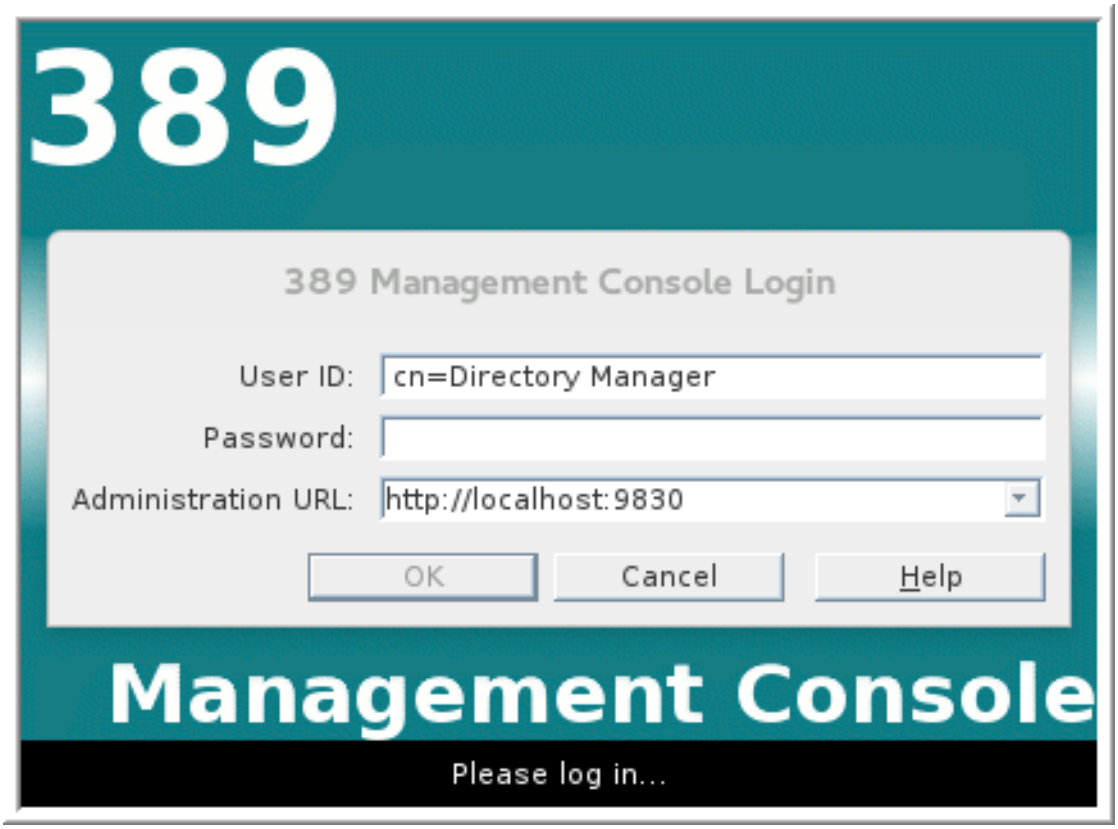

3. The 389 Management Console window appears. Select the **`Servers and Applications`** tab.

4. In the left-hand pane, drill down to the **`Directory Server`** icon.

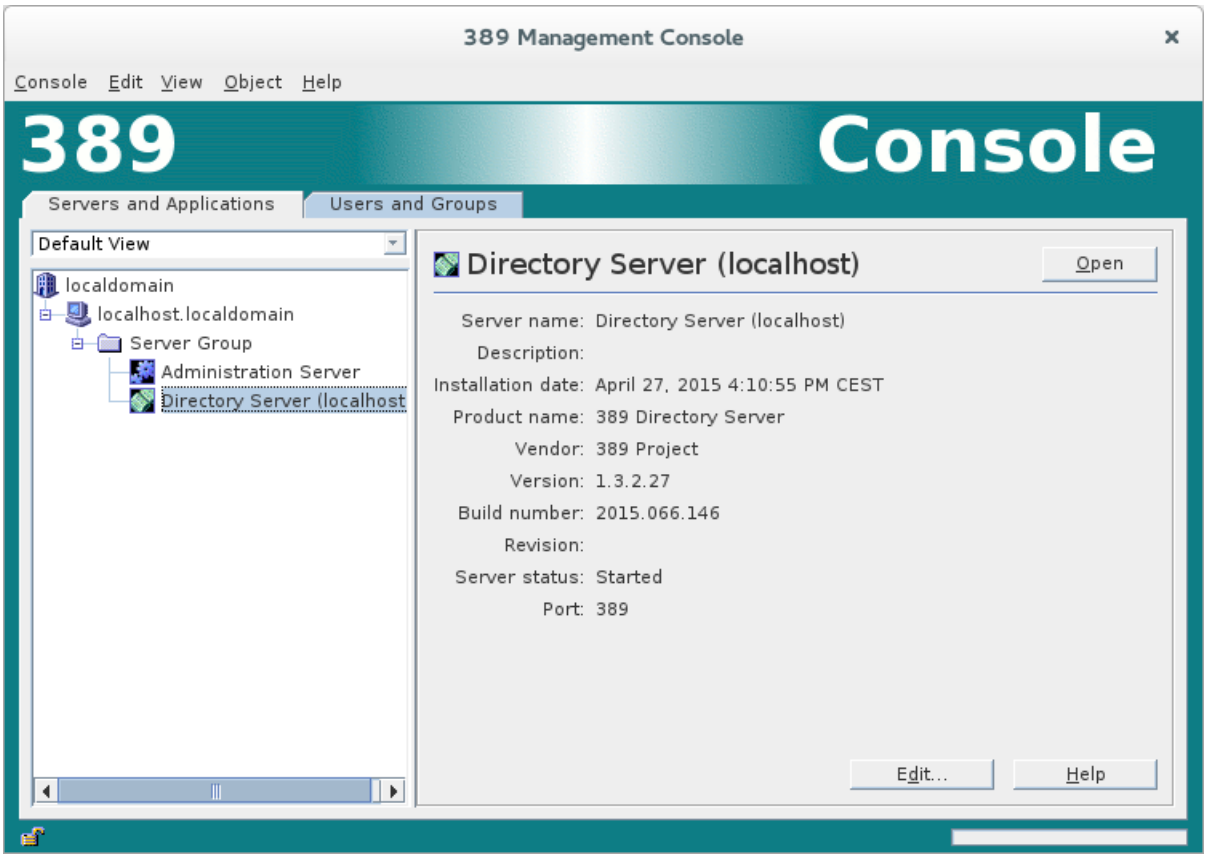

5. Select the **`Directory Server`** icon in the left-hand pane and click **`Open`**, to open the **`389 Directory Server Console`**.

6. In the **389 Directory Server Console**, click the **Directory** tab, to view the Directory Information Tree (DIT).

7. Expand the root node, ***YourDomain*** (usually named after a hostname, and shown as **localdomain** in the following screenshot), to view the DIT.

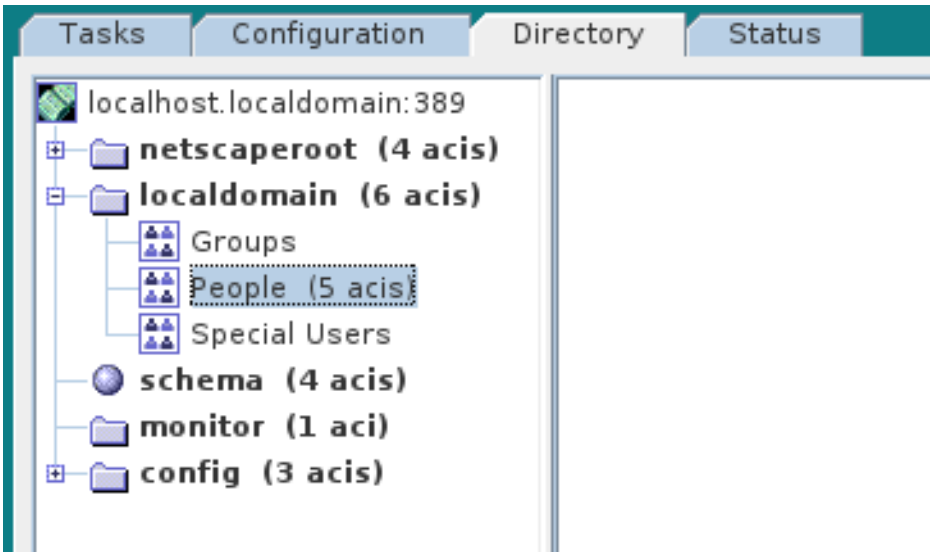

8.3. ADD USER ENTRIES TO THE DIRECTORY SERVER

The basic prerequisite for using LDAP authentication with the OSGi container is to have an X.500 directory server running and configured with a collection of user entries. For many use cases, you will also want to configure a number of groups to manage user roles.

Alternative to adding user entries

If you already have user entries and groups defined in your LDAP server, you might prefer to map the existing LDAP groups to JAAS roles using the **roles.mapping** property in the **LDAPLoginModule** configuration, instead of creating new entries. For details, see Section 2.1.7, “JAAS LDAP Login Module”.

Goals

In this portion of the tutorial you will

- Add three user entries to the LDAP server
- Add four groups to the LDAP server

Adding user entries

Perform the following steps to add user entries to the directory server:

1. Ensure that the LDAP server and console are running. See Section 8.2, “Set-up a Directory Server and Console”.

2. In the **Directory Server Console**, click on the **Directory** tab, and drill down to the **People** node, under the ***YourDomain*** node (where ***YourDomain*** is shown as **localdomain** in the following screenshots).

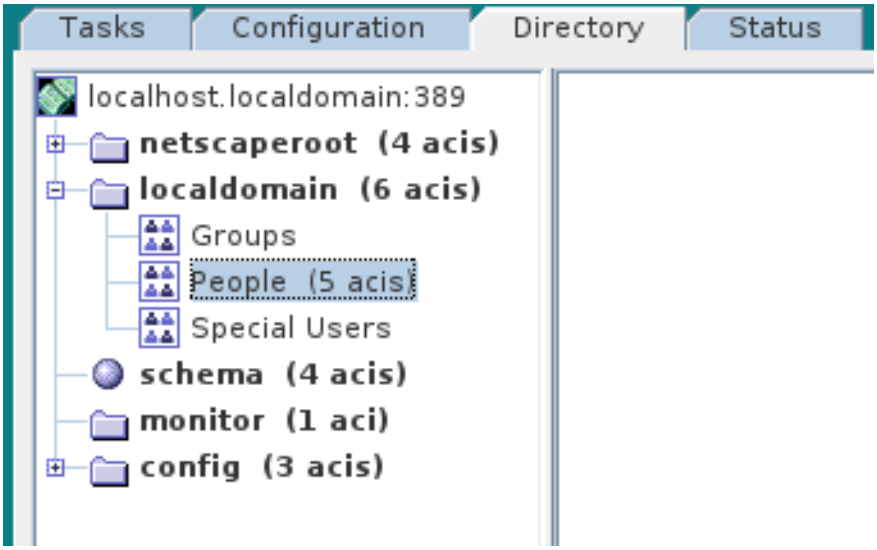

3. Right-click the **People** node, and select menu:[> New > > **User** >] from the context menu, to open the **Create New User** dialog.
4. Select the **User** tab in the left-hand pane of the **Create New User** dialog.
5. Fill in the fields of the **User** tab, as follows:
   a. Set the **First Name** field to **John**.
   b. Set the **Last Name** field to **Doe**.
   c. Set the **User ID** field to **jdoe**.
   d. Enter the password, **secret**, in the **Password** field.
   e. Enter the password, **secret**, in the **Confirm Password** field.

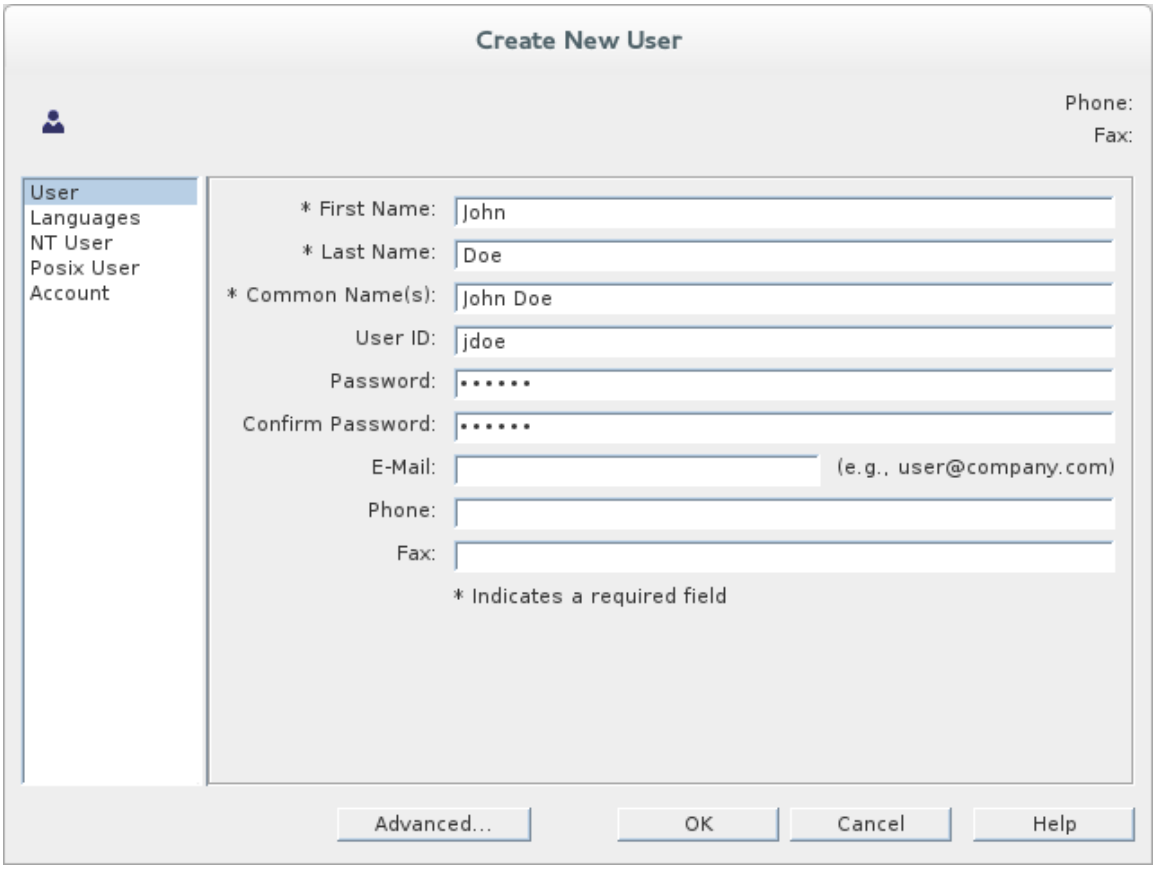

6. Click **OK**.

7. Add a user **Jane Doe** by following Step 3 to Step 6.
   In Step 5.e, use **janedoe** for the new user's **User ID** and use the password, **secret**, for the password fields.

8. Add a user **Camel Rider** by following Step 3 to Step 6.
   In Step 5.e, use **crider** for the new user's **User ID** and use the password, **secret**, for the password fields.

Adding groups for the roles

To add the groups that define the roles:

1. In the **Directory** tab of the **Directory Server Console**, drill down to the **Groups** node, under the ***YourDomain*** node.

2. Right-click the **Groups** node, and select menu:[> New > >**Group** >] from the context menu, to open the **Create New Group** dialog.

3. Select the **General** tab in the left-hand pane of the **Create New Group** dialog.

4. Fill in the fields of the **General** tab, as follows:

   a. Set the **Group Name** field to **admin**.

   b. Optionally, enter a description in the **Description** field.

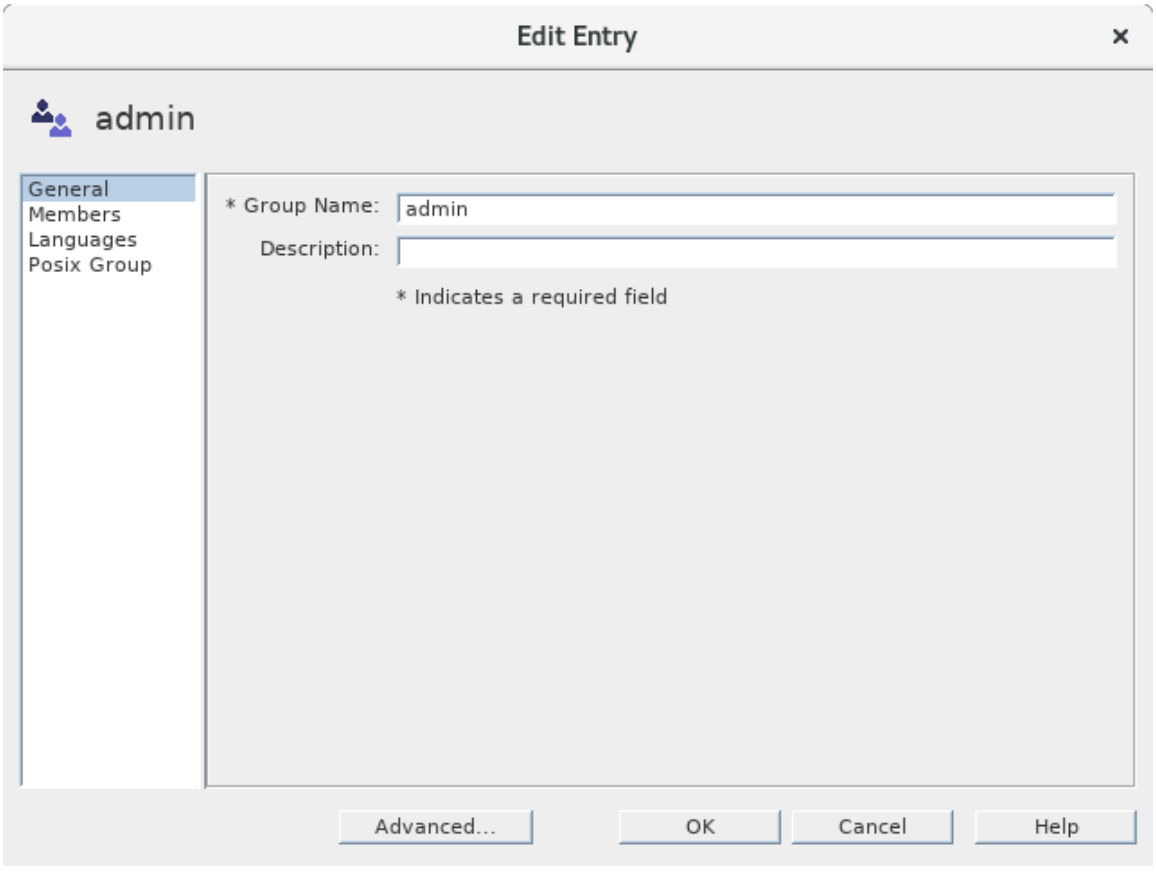

5. Select the **Members** tab in the left-hand pane of the **Create New Group** dialog.

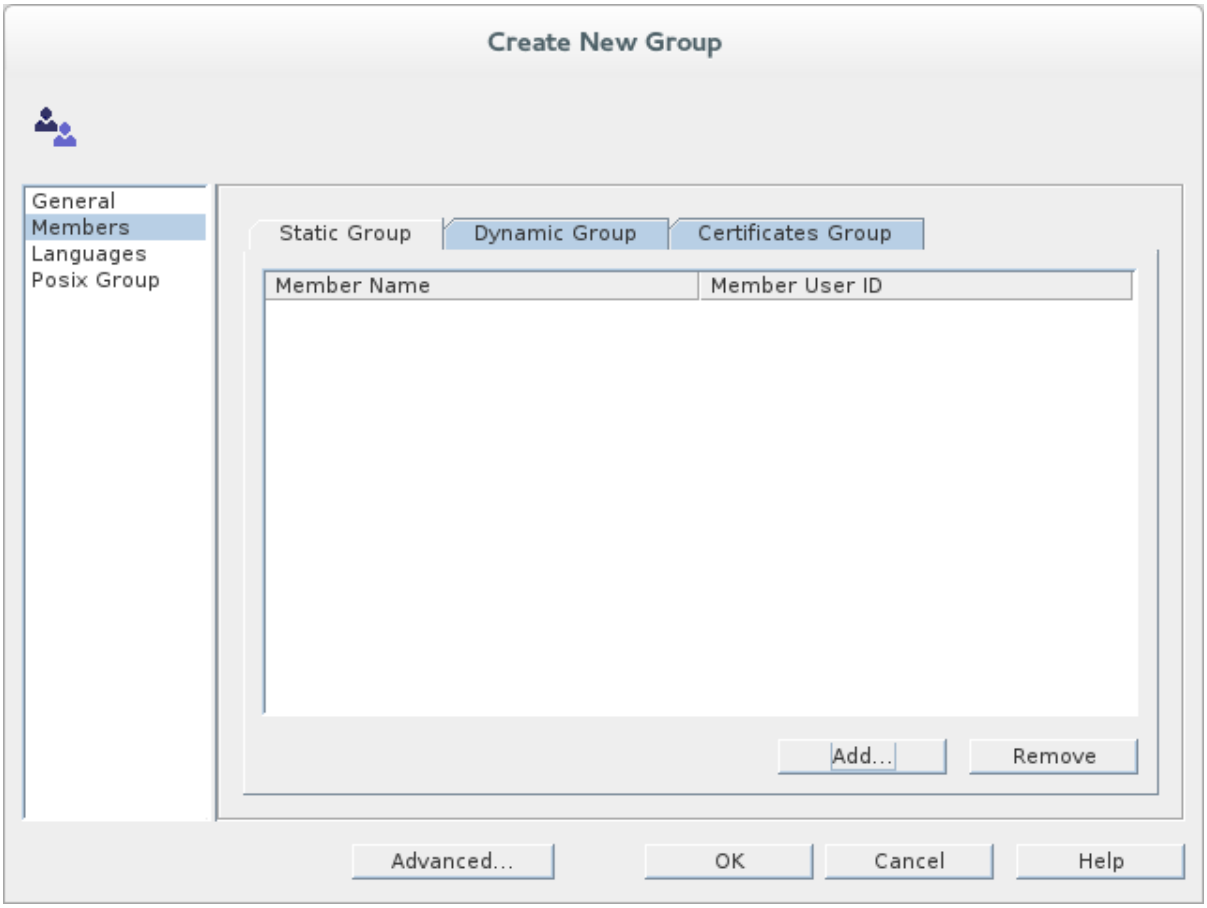

6. Click **Add** to open the **Search users and groups** dialog.
7. In the **Search** field, select **Users** from the drop-down menu, and click the **Search** button.

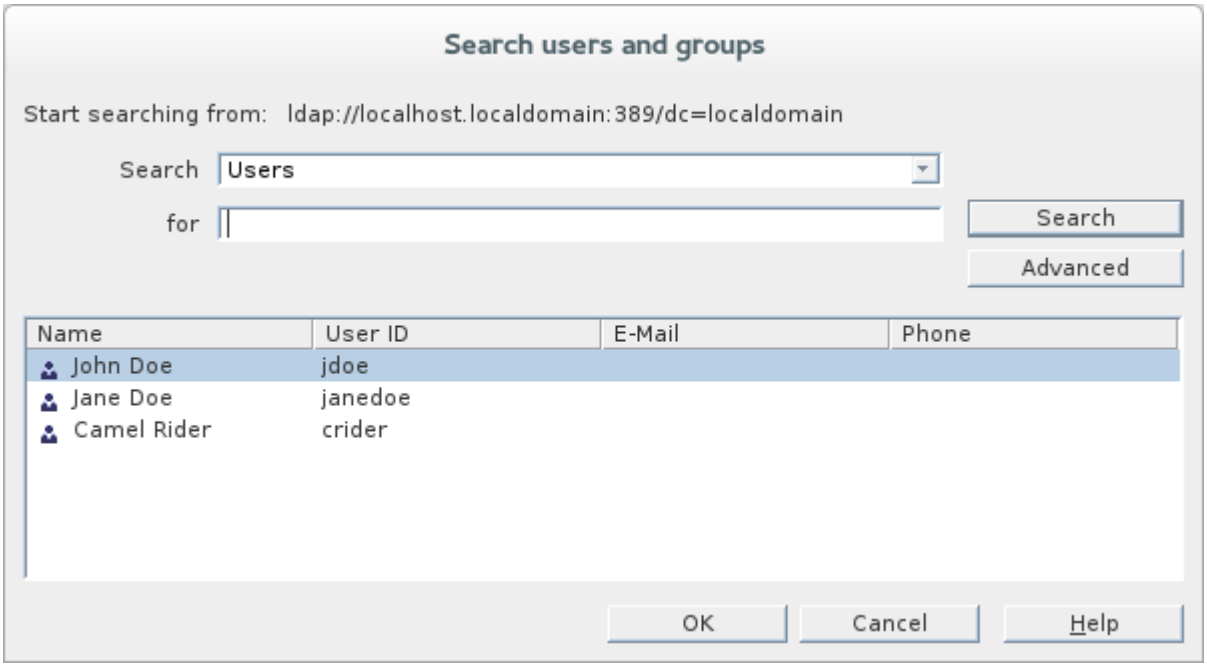

8. From the list of users that is now displayed, select **John Doe**.
9. Click **OK**, to close the **Search users and groups** dialog.
10. Click **OK**, to close the **Create New Group** dialog.
11. Add a **manager** role by following Step 2 to Step 10.

In Step 4, enter **manager** in the **Group Name** field.

In Step 8, select **Jane Doe**.

12. Add a **viewer** role by following Step 2 to Step 10.
In Step 4, enter **viewer** in the **Group Name** field.

In Step 8, select **Camel Rider**.

13. Add an **ssh** role by following Step 2 to Step 10.
In Step 4, enter **ssh** in the **Group Name** field.

In Step 8, select *all* of the users, **John Doe**, **Jane Doe**, and **Camel Rider**.

8.4. ENABLE LDAP AUTHENTICATION IN THE OSGI CONTAINER

This section explains how to configure an LDAP realm in the OSGi container. The new realm overrides the default **karaf** realm, so that the container authenticates credentials based on user entries stored in the X.500 directory server.

References

More detailed documentation is available on LDAP authentication, as follows:

- **LDAPLoginModule options**—are described in detail in Section 2.1.7, "JAAS LDAP Login Module".
- **Configurations for other directory servers**—this tutorial covers only 389-DS. For details of how to configure other directory servers, such as Microsoft Active Directory, see the section called "Filter settings for different directory servers".

Procedure for standalone OSGi container

To enable LDAP authentication in a standalone OSGi container:

1. Ensure that the X.500 directory server is running.
2. Start the Karaf container by entering the following command in a terminal window:

```
./bin/fuse
```

3. Create a file called **ldap-module.xml**.
4. Copy Example 8.1, "JAAS Realm for Standalone" into **ldap-module.xml**.

Example 8.1. JAAS Realm for Standalone

```
<?xml version="2.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0"
  xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-
ext/v1.0.0">

  <jaas:config name="karaf" rank="200">
    <jaas:module
```

```
className="org.apache.karaf.jaas.modules.ldap.LDAPLoginModule"
                flags="required">
      initialContextFactory=com.sun.jndi.ldap.LdapCtxFactory
      connection.url=ldap://localhost:389
      connection.username=cn=Directory Manager
      connection.password=DIRECTORY_MANAGER_PASSWORD
      connection.protocol=
      user.base.dn=ou=People,dc=localdomain
      user.filter=(&amp;(objectClass=inetOrgPerson)(uid=%u))
      user.search.subtree=true
      role.base.dn=ou=Groups,dc=localdomain
      role.name.attribute=cn
      role.filter=(uniquemember=%fqdn)
      role.search.subtree=true
      authentication=simple
    </jaas:module>
  </jaas:config>
</blueprint>
```

You must customize the following settings in the **`ldap-module.xml`** file:

connection.url

Set this URL to the actual location of your directory server instance. Normally, this URL has the format, **`ldap://Hostname:Port`**. For example, the default port for the 389 Directory Server is IP port **`389`**.

connection.username

Specifies the username that is used to authenticate the connection to the directory server. For 389 Directory Server, the default is usually **`cn=Directory Manager`**.

connection.password

Specifies the password part of the credentials for connecting to the directory server.

authentication

You can specify either of the following alternatives for the authentication protocol:

- **`simple`** implies that user credentials are supplied and you are obliged to set the **`connection.username`** and **`connection.password`** options in this case.
- **none** implies that authentication is **not** performed. You must not set the **`connection.username`** and **`connection.password`** options in this case.
  This login module creates a JAAS realm called **`karaf`**, which is the same name as the default JAAS realm used by Fuse. By redefining this realm with a **`rank`** attribute value greater than **`0`**, it overrides the standard **`karaf`** realm which has the rank **`0`**.

  For more details about how to configure Fuse to use LDAP, see Section 2.1.7, “JAAS LDAP Login Module”.

IMPORTANT

When setting the JAAS properties above, do **not** enclose the property values in double quotes.

5. To deploy the new LDAP module, copy the **ldap-module.xml** into the Karaf container's **deploy/** directory (hot deploy).
   The LDAP module is automatically activated.

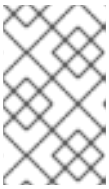

NOTE

Subsequently, if you need to undeploy the LDAP module, you can do so by deleting the **ldap-module.xml** file from the **deploy/** directory **while the Karaf container is running**.

Test the LDAP authentication

Test the new LDAP realm by connecting to the running container using the Karaf **client** utility, as follows:

1. Open a new command prompt.
2. Change directory to the Karaf ***InstallDir*/bin** directory.
3. Enter the following command to log on to the running container instance using the identity **jdoe**:

```
./client -u jdoe -p secret
```

   You should successfully log into the container's remote console. At the command console, type **jaas:** followed by the [Tab] key (to activate content completion):

```
jdoe@root()> jaas:
Display all 31 possibilities? (31 lines)?
jaas:cancel
jaas:group-add
...
jaas:whoami
```

   You should see that **jdoe** has access to all of the **jaas** commands (consistent with the **admin** role).

4. Log off the remote console by entering the **logout** command.
5. Enter the following command to log on to the running container instance using the identity **janedoe**:

```
./client -u janedoe -p secret
```

   You should successfully log into the container's remote console. At the command console, type **jaas:** followed by the [Tab] key (to activate content completion):

```
janedoe@root()> jaas:
Display all 25 possibilities? (25 lines)?
jaas:cancel
jaas:group-add
...
jaas:users
```

   You should see that **janedoe** has access to almost all of the **jaas** commands (consistent with

the **manager** role).

6. Log off the remote console by entering the **logout** command.

7. Enter the following command to log on to the running container instance using the identity **crider**:

```
./client -u crider -p secret
```

   You should successfully log into the container's remote console. At the command console, type **jaas:** followed by the [Tab] key (to activate content completion):

```
crider@root()> jaas:
jaas:manage
jaas:realm-list
jaas:realm-manage
jaas:realms
jaas:user-list
jaas:users
```

   You should see that **crider** has access to only five of the **jaas** commands (consistent with the **viewer** role).

8. Log off the remote console by entering the **logout** command.

Troubleshooting

If you run into any difficulties while testing the LDAP connection, increase the logging level to **DEBUG** to get a detailed trace of what is happening on the connection to the LDAP server.

Perform the following steps:

1. From the Karaf console, enter the following command to increase the logging level to **DEBUG**:

```
log:set DEBUG
```

2. Observe the Karaf log in real time:

```
log:tail
```

   To escape from the log listing, type Ctrl-C.

APPENDIX A. MANAGING CERTIFICATES

Abstract

TLS authentication uses X.509 certificates—a common, secure and reliable method of authenticating your application objects. You can create X.509 certificates that identify your Red Hat Fuse applications.

A.1. WHAT IS AN X.509 CERTIFICATE?

Role of certificates

An X.509 certificate binds a name to a public key value. The role of the certificate is to associate a public key with the identity contained in the X.509 certificate.

Integrity of the public key

Authentication of a secure application depends on the integrity of the public key value in the application's certificate. If an impostor replaces the public key with its own public key, it can impersonate the true application and gain access to secure data.

To prevent this type of attack, all certificates must be signed by a *certification authority* (CA). A CA is a trusted node that confirms the integrity of the public key value in a certificate.

Digital signatures

A CA signs a certificate by adding its *digital signature* to the certificate. A digital signature is a message encoded with the CA's private key. The CA's public key is made available to applications by distributing a certificate for the CA. Applications verify that certificates are validly signed by decoding the CA's digital signature with the CA's public key.

WARNING

The supplied demonstration certificates are self-signed certificates. These certificates are insecure because anyone can access their private key. To secure your system, you must create new certificates signed by a trusted CA.

Contents of an X.509 certificate

An X.509 certificate contains information about the certificate subject and the certificate issuer (the CA that issued the certificate). A certificate is encoded in Abstract Syntax Notation One (ASN.1), a standard syntax for describing messages that can be sent or received on a network.

The role of a certificate is to associate an identity with a public key value. In more detail, a certificate includes:

- A *subject distinguished name (DN)* that identifies the certificate owner.
- The *public key* associated with the subject.

- X.509 version information.
- A *serial number* that uniquely identifies the certificate.
- An *issuer DN* that identifies the CA that issued the certificate.
- The digital signature of the issuer.
- Information about the algorithm used to sign the certificate.
- Some optional X.509 v.3 extensions; for example, an extension exists that distinguishes between CA certificates and end-entity certificates.

Distinguished names

A DN is a general purpose X.500 identifier that is often used in the context of security.

See Appendix B, *ASN.1 and Distinguished Names* for more details about DNs.

A.2. CERTIFICATION AUTHORITIES

A.2.1. Introduction to Certificate Authorities

A CA consists of a set of tools for generating and managing certificates and a database that contains all of the generated certificates. When setting up a system, it is important to choose a suitable CA that is sufficiently secure for your requirements.

There are two types of CA you can use:

- commercial CAs are companies that sign certificates for many systems.
- private CAs are trusted nodes that you set up and use to sign certificates for your system only.

A.2.2. Commercial Certification Authorities

Signing certificates

There are several commercial CAs available. The mechanism for signing a certificate using a commercial CA depends on which CA you choose.

Advantages of commercial CAs

An advantage of commercial CAs is that they are often trusted by a large number of people. If your applications are designed to be available to systems external to your organization, use a commercial CA to sign your certificates. If your applications are for use within an internal network, a private CA might be appropriate.

Criteria for choosing a CA

Before choosing a commercial CA, consider the following criteria:

- What are the certificate-signing policies of the commercial CAs?
- Are your applications designed to be available on an internal network only?

- What are the potential costs of setting up a private CA compared to the costs of subscribing to a commercial CA?

A.2.3. Private Certification Authorities

Choosing a CA software package

If you want to take responsibility for signing certificates for your system, set up a private CA. To set up a private CA, you require access to a software package that provides utilities for creating and signing certificates. Several packages of this type are available.

OpenSSL software package

One software package that allows you to set up a private CA is OpenSSL, http://www.openssl.org. The OpenSSL package includes basic command line utilities for generating and signing certificates. Complete documentation for the OpenSSL command line utilities is available at http://www.openssl.org/docs.

Setting up a private CA using OpenSSL

To set up a private CA, see the instructions in Section A.5, “Creating Your Own Certificates”.

Choosing a host for a private certification authority

Choosing a host is an important step in setting up a private CA. The level of security associated with the CA host determines the level of trust associated with certificates signed by the CA.

If you are setting up a CA for use in the development and testing of Red Hat Fuse applications, use any host that the application developers can access. However, when you create the CA certificate and private key, do not make the CA private key available on any hosts where security-critical applications run.

Security precautions

If you are setting up a CA to sign certificates for applications that you are going to deploy, make the CA host as secure as possible. For example, take the following precautions to secure your CA:

- Do not connect the CA to a network.
- Restrict all access to the CA to a limited set of trusted users.
- Use an RF-shield to protect the CA from radio-frequency surveillance.

A.3. CERTIFICATE CHAINING

Certificate chain

A *certificate chain* is a sequence of certificates, where each certificate in the chain is signed by the subsequent certificate.

Figure A.1, “A Certificate Chain of Depth 2” shows an example of a simple certificate chain.

Figure A.1. A Certificate Chain of Depth 2

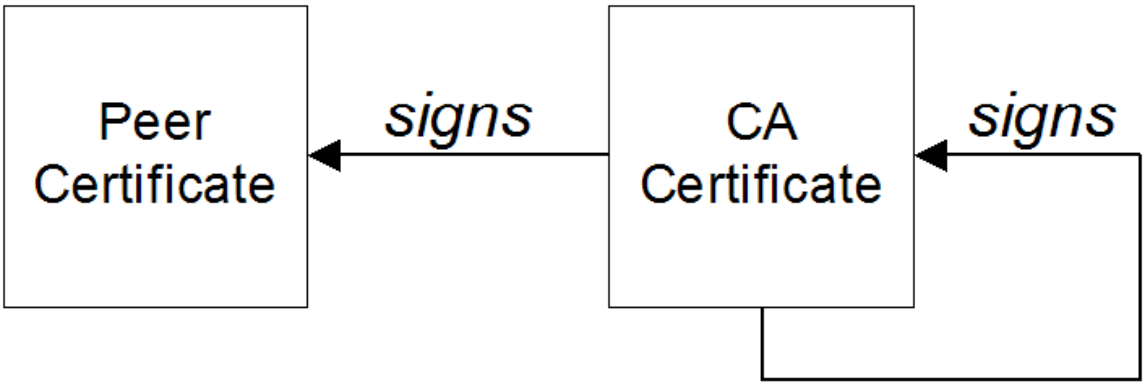

Self-signed certificate

The last certificate in the chain is normally a *self-signed certificate*—a certificate that signs itself.

Chain of trust

The purpose of a certificate chain is to establish a chain of trust from a peer certificate to a trusted CA certificate. The CA vouches for the identity in the peer certificate by signing it. If the CA is one that you trust (indicated by the presence of a copy of the CA certificate in your root certificate directory), this implies you can trust the signed peer certificate as well.

Certificates signed by multiple CAs

A CA certificate can be signed by another CA. For example, an application certificate could be signed by the CA for the finance department of Progress Software, which in turn is signed by a self-signed commercial CA.

Figure A.2, “A Certificate Chain of Depth 3” shows what this certificate chain looks like.

Figure A.2. A Certificate Chain of Depth 3

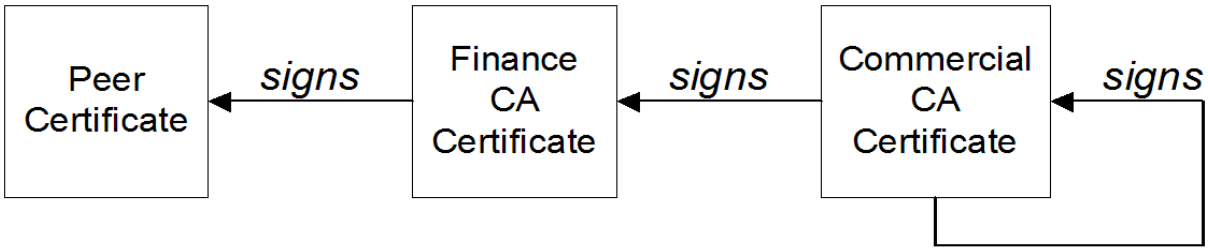

Trusted CAs

An application can accept a peer certificate, provided it trusts at least one of the CA certificates in the signing chain.

A.4. SPECIAL REQUIREMENTS ON HTTPS CERTIFICATES

Overview

The HTTPS specification mandates that HTTPS clients must be capable of verifying the identity of the

server. This can potentially affect how you generate your X.509 certificates. The mechanism for verifying the server identity depends on the type of client. Some clients might verify the server identity by accepting only those server certificates signed by a particular trusted CA. In addition, clients can inspect the contents of a server certificate and accept only the certificates that satisfy specific constraints.

In the absence of an application-specific mechanism, the HTTPS specification defines a generic mechanism, known as the *HTTPS URL integrity check*, for verifying the server identity. This is the standard mechanism used by Web browsers.

HTTPS URL integrity check

The basic idea of the URL integrity check is that the server certificate's identity must match the server host name. This integrity check has an important impact on how you generate X.509 certificates for HTTPS: **the certificate identity (usually the certificate subject DN's common name) must match the host name on which the HTTPS server is deployed**.

The URL integrity check is designed to prevent **man-in-the-middle** attacks.

Reference

The HTTPS URL integrity check is specified by RFC 2818, published by the Internet Engineering Task Force (IETF) at http://www.ietf.org/rfc/rfc2818.txt.

How to specify the certificate identity

The certificate identity used in the URL integrity check can be specified in one of the following ways:

- Using commonName
- Using subectAltName

Using commonName

The usual way to specify the certificate identity (for the purpose of the URL integrity check) is through the Common Name (CN) in the subject DN of the certificate.

For example, if a server supports secure TLS connections at the following URL:

```
https://www.redhat.com/secure
```

The corresponding server certificate would have the following subject DN:

```
C=IE,ST=Co. Dublin,L=Dublin,O=RedHat,
OU=System,CN=www.redhat.com
```

Where the CN has been set to the host name, **`www.redhat.com`**.

For details of how to set the subject DN in a new certificate, see Section A.5, "Creating Your Own Certificates".

Using subjectAltName (multi-homed hosts)

Using the subject DN's Common Name for the certificate identity has the disadvantage that only **one** host name can be specified at a time. If you deploy a certificate on a multi-homed host, however, you

might find it is practical to allow the certificate to be used with **any** of the multi-homed host names. In this case, it is necessary to define a certificate with multiple, alternative identities, and this is only possible using the **`subjectAltName`** certificate extension.

For example, if you have a multi-homed host that supports connections to either of the following host names:

```
www.redhat.com
www.jboss.org
```

Then you can define a **`subjectAltName`** that explicitly lists both of these DNS host names. If you generate your certificates using the **`openssl`** utility, edit the relevant line of your **`openssl.cnf`** configuration file to specify the value of the **`subjectAltName`** extension, as follows:

```
subjectAltName=DNS:www.redhat.com,DNS:www.jboss.org
```

Where the HTTPS protocol matches the server host name against either of the DNS host names listed in the **`subjectAltName`** (the **`subjectAltName`** takes precedence over the Common Name).

The HTTPS protocol also supports the wildcard character, `\*`, in host names. For example, you can define the **`subjectAltName`** as follows:

```
subjectAltName=DNS:*.jboss.org
```

This certificate identity matches any three-component host name in the domain **jboss.org**.

WARNING

You must **never** use the wildcard character in the domain name (and you must take care never to do this accidentally by forgetting to type the dot, `.`, delimiter in front of the domain name). For example, if you specified **`*jboss.org`**, your certificate could be used on *any* domain that ends in the letters **`jboss`**.

A.5. CREATING YOUR OWN CERTIFICATES

Abstract

This chapter describes the techniques and procedures to set up your own private Certificate Authority (CA) and to use this CA to generate and sign your own certificates.

WARNING

Creating and managing your own certificates requires an expert knowledge of security. While the procedures described in this chapter can be convenient for generating your own certificates for demonstration and testing environments, it is **not recommended** to use these certificates in a production environment.

A.5.1. Install the OpenSSL Utilities

Installing OpenSSL on RHEL and Fedora platforms

On Red Hat Enterprise Linux (RHEL) 5 and 6 and Fedora platforms, are made available as an RPM package. To install OpenSSL, enter the following command (executed with administrator privileges):

```
yum install openssl
```

Source code distribution

The source distribution of OpenSSL is available from http://www.openssl.org/docs. The OpenSSL project provides source code distributions **only**. You cannot download a binary install of the OpenSSL utilities from the OpenSSL Web site.

A.5.2. Set Up a Private Certificate Authority

Overview

If you choose to use a private CA you need to generate your own certificates for your applications to use. The OpenSSL project provides free command-line utilities for setting up a private CA, creating signed certificates, and adding the CA to your Java keystore.

WARNING

Setting up a private CA for a production environment requires a high level of expertise and extra care must be taken to protect the certificate store from external threats.

Steps to set up a private Certificate Authority

To set up your own private Certificate Authority:

1. Create the directory structure for the CA, as follows:

```
X509CA/demoCA
X509CA/demoCA/private
```

```
X509CA/demoCA/certs
X509CA/demoCA/newcerts
X509CA/demoCA/crl
```

2. Using a text editor, create the file, ***X509CA/*openssl.cfg**, and add the following contents to this file:

Example A.1. OpenSSL Configuration

```
#
# SSLeay example configuration file.
# This is mostly being used for generation of certificate
requests.
#

RANDFILE            = ./.rnd

#####################################################################
##
[ req ]
default_bits        = 2048
default_keyfile     = keySS.pem
distinguished_name  = req_distinguished_name
encrypt_rsa_key     = yes
default_md          = sha1

[ req_distinguished_name ]
countryName         = Country Name (2 letter code)

organizationName    = Organization Name (eg, company)

commonName          = Common Name (eg, YOUR name)

#####################################################################
##
[ ca ]
default_ca          = CA_default        # The default ca section

#####################################################################
##
[ CA_default ]

dir                 = ./demoCA              # Where everything is
kept
certs               = $dir/certs            # Where the issued
certs are kept
crl_dir             = $dir/crl              # Where the issued crl
are kept
database            = $dir/index.txt        # database index file.
#unique_subject     = no                    # Set to 'no' to
allow creation of
                                            # several
certificates with same subject.
new_certs_dir       = $dir/newcerts         # default place for
new certs.
```

```
certificate        = $dir/cacert.pem       # The CA certificate
serial             = $dir/serial           # The current serial
number
crl                = $dir/crl.pem          # The current CRL
private_key        = $dir/private/cakey.pem# The private key
RANDFILE           = $dir/private/.rand    # private random
number file

name_opt           = ca_default            # Subject Name
options
cert_opt           = ca_default            # Certificate field
options

default_days       = 365                   # how long to certify
for
default_crl_days   = 30                    # how long before next
CRL
default_md         = md5                   # which md to use.
preserve           = no                    # keep passed DN
ordering

policy             = policy_anything

[ policy_anything ]
countryName            = optional
stateOrProvinceName    = optional
localityName           = optional
organizationName       = optional
organizationalUnitName = optional
commonName             = supplied
emailAddress           = optional
```

IMPORTANT

The preceding **openssl.cfg** configuration file is provided **as a demonstration only**. In a production environment, this configuration file would need to be carefully elaborated by an engineer with a high level of security expertise, and actively maintained to protect against evolving security threats.

3. Initialize the **demoCA/serial** file, which must have the initial contents **01** (zero one). Enter the following command:

```
echo 01 > demoCA/serial
```

4. Initialize the **demoCA/index.txt**, which **must** initially be completely empty. Enter the following command:

```
touch demoCA/index.txt
```

5. Create a new self-signed CA certificate and private key with the command:

```
openssl req -x509 -new -config openssl.cfg -days 365 -out
demoCA/cacert.pem -keyout demoCA/private/cakey.pem
```

You are prompted for a pass phrase for the CA private key and details of the CA distinguished name as shown in Example A.2, “Creating a CA Certificate”.

Example A.2. Creating a CA Certificate

```
Generating a 2048 bit RSA private key
..............................................................
.........+++
..................+++
writing new private key to 'demoCA/private/cakey.pem'
Enter PEM pass phrase:
Verifying - Enter PEM pass phrase:
-----
You are about to be asked to enter information that will be
incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name
or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) []:DE
Organization Name (eg, company) []:Red Hat
Common Name (eg, YOUR name) []:Scooby Doo
```

NOTE

The security of the CA depends on the security of the private key file and the private key pass phrase used in this step.

You must ensure that the file names and location of the CA certificate and private key, **`cacert.pem`** and **`cakey.pem`**, are the same as the values specified in **`openssl.cfg`**.

A.5.3. Create a CA Trust Store File

Overview

A trust store file is commonly required on the client side of an SSL/TLS connection, in order to verify a server’s identity. A trust store file can also be used to check digital signatures (for example, to check that a signature was made using the private key corresponding to one of the trusted certificates in the trust store file).

Steps to create a CA trust store

To add one of more CA certificates to a trust store file:

1. Assemble the collection of trusted CA certificates that you want to deploy.
   The trusted CA certificates can be obtained from public CAs or private CAs. The trusted CA certificates can be in any format that is compatible with the Java **`keystore`** utility; for example, PEM format. All you need are the certificates themselves—the private keys and passwords are

not required.

2. Add a CA certificate to the trust store using the **keytool -import** command.
   Enter the following command to add the CA certificate, **cacert.pem**, in PEM format, to a JKS trust store.

   ```
   keytool -import -file cacert.pem -alias CAAlias -keystore
   truststore.ts -storepass StorePass
   ```

   Where **truststore.ts** is a keystore file containing CA certificates. If this file does not already exist, the **keytool** command creates it. The ***CAAlias*** is a convenient identifier for the imported CA certificate and ***StorePass*** is the password required to access the keystore file.

3. Repeat the previous step to add all of the CA certificates to the trust store.

A.5.4. Generate and Sign a New Certificate

Overview

In order for a certificate to be useful in the real world, it must be signed by a CA, which vouches for the authenticity of the certificate. This facilitates a scalable solution for certificate verification, because it means that a single CA certificate can be used to verify a large collection of certificates.

Steps to generate and sign a new certificate

To generate and sign a new certificate, using your own private CA, perform the following steps:

1. Generate a certificate and private key pair using the **keytool -genkeypair** command, as follows:

   ```
   keytool -genkeypair -keyalg RSA -dname "CN=Alice, OU=Engineering,
   O=Red Hat, ST=Dublin, C=IE" -validity 365 -alias alice -keypass
   KeyPass -keystore alice.ks -storepass StorePass
   ```

   Because the specified keystore, **alice.ks**, did not exist prior to issuing the command implicitly creates a new keystore and sets its password to ***StorePass***.

   The **-dname** and **-validity** flags define the contents of the newly created X.509 certificate.

NOTE

When specifying the certificate's Distinguished Name (through the **-dname** parameter), you must be sure to observe any policy constraints specified in the **openssl.cfg** file. If those policy constraints are not heeded, you will not be able to sign the certificate using the CA (in the next steps).

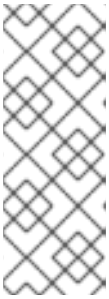

NOTE

It is essential to generate the key pair with the **-keyalg RSA** option (or a key algorithm of similar strength). The default key algorithm uses a combination of DSA encryption and SHA-1 signature. But the SHA-1 algorithm is no longer regarded as sufficiently secure and modern Web browsers will reject certificates signed using SHA-1. When you select the RSA key algorithm, the **keytool** utility uses an SHA-2 algorithm instead.

2. Create a certificate signing request using the **keystore -certreq** command.
   Create a new certificate signing request for the **alice.ks** certificate and export it to the **alice_csr.pem** file, as follows:

```
keytool -certreq -alias alice -file alice_csr.pem -keypass KeyPass -
keystore alice.ks -storepass StorePass
```

3. Sign the CSR using the **openssl ca** command.
   Sign the CSR for the Alice certificate, using your private CA, as follows:

```
openssl ca -config openssl.cfg -days 365 -in alice_csr.pem -out
alice_signed.pem
```

   You will prompted to enter the CA private key pass phrase you used when creating the CA (in the section called “Steps to set up a private Certificate Authority”).

   For more details about the **openssl ca** command see http://www.openssl.org/docs/apps/ca.html#.

4. Convert the signed certificate to PEM only format using the **openssl x509** command with the **-outform** option set to **PEM**. Enter the following command:

```
openssl x509 -in alice_signed.pem -out alice_signed.pem -outform PEM
```

5. Concatenate the CA certificate file and the converted, signed certificate file to form a certificate chain. For example, on Linux and UNIX platforms, you can concatenate the CA certificate file and the signed Alice certificate, **alice_signed.pem**, as follows:

```
cat demoCA/cacert.pem alice_signed.pem > alice.chain
```

6. Import the new certificate’s full certificate chain into the Java keystore using the **keytool -import** command. Enter the following command:

```
keytool -import -file alice.chain -keypass KeyPass -keystore
alice.ks -storepass StorePass
```

APPENDIX B. ASN.1 AND DISTINGUISHED NAMES

Abstract

The OSI Abstract Syntax Notation One (ASN.1) and X.500 Distinguished Names play an important role in the security standards that define X.509 certificates and LDAP directories.

B.1. ASN.1

Overview

The *Abstract Syntax Notation One* (ASN.1) was defined by the OSI standards body in the early 1980s to provide a way of defining data types and structures that are independent of any particular machine hardware or programming language. In many ways, ASN.1 can be considered a forerunner of modern interface definition languages, such as the OMG's IDL and WSDL, which are concerned with defining platform-independent data types.

ASN.1 is important, because it is widely used in the definition of standards (for example, SNMP, X.509, and LDAP). In particular, ASN.1 is ubiquitous in the field of security standards. The formal definitions of X.509 certificates and distinguished names are described using ASN.1 syntax. You do not require detailed knowledge of ASN.1 syntax to use these security standards, but you need to be aware that ASN.1 is used for the basic definitions of most security-related data types.

BER

The OSI's Basic Encoding Rules (BER) define how to translate an ASN.1 data type into a sequence of octets (binary representation). The role played by BER with respect to ASN.1 is, therefore, similar to the role played by GIOP with respect to the OMG IDL.

DER

The OSI's Distinguished Encoding Rules (DER) are a specialization of the BER. The DER consists of the BER plus some additional rules to ensure that the encoding is unique (BER encodings are not).

References

You can read more about ASN.1 in the following standards documents:

- ASN.1 is defined in X.208.
- BER is defined in X.209.

B.2. DISTINGUISHED NAMES

Overview

Historically, distinguished names (DN) are defined as the primary keys in an X.500 directory structure. However, DNs have come to be used in many other contexts as general purpose identifiers. In Apache CXF, DNs occur in the following contexts:

- X.509 certificates—for example, one of the DNs in a certificate identifies the owner of the certificate (the security principal).

- LDAP—DNs are used to locate objects in an LDAP directory tree.

String representation of DN

Although a DN is formally defined in ASN.1, there is also an LDAP standard that defines a UTF-8 string representation of a DN (see **RFC 2253**). The string representation provides a convenient basis for describing the structure of a DN.

NOTE

The string representation of a DN does **not** provide a unique representation of DER-encoded DN. Hence, a DN that is converted from string format back to DER format does not always recover the original DER encoding.

DN string example

The following string is a typical example of a DN:

```
C=US,O=IONA Technologies,OU=Engineering,CN=A. N. Other
```

Structure of a DN string

A DN string is built up from the following basic elements:

- OID .
- Attribute Types .
- AVA .
- RDN .

OID

An OBJECT IDENTIFIER (OID) is a sequence of bytes that uniquely identifies a grammatical construct in ASN.1.

Attribute types

The variety of attribute types that can appear in a DN is theoretically open-ended, but in practice only a small subset of attribute types are used. Table B.1, “Commonly Used Attribute Types” shows a selection of the attribute types that you are most likely to encounter:

Table B.1. Commonly Used Attribute Types

String Representation	X.500 Attribute Type	Size of Data	Equivalent OID
C	countryName	2	2.5.4.6
O	organizationName	1…64	2.5.4.10
OU	organizationalUnitName	1…64	2.5.4.11

String Representation	X.500 Attribute Type	Size of Data	Equivalent OID
CN	commonName	1…64	2.5.4.3
ST	stateOrProvinceName	1…64	2.5.4.8
L	localityName	1…64	2.5.4.7
STREET	streetAddress		
DC	domainComponent		
UID	userid		

AVA

An *attribute value assertion* (AVA) assigns an attribute value to an attribute type. In the string representation, it has the following syntax:

```
<attr-type>=<attr-value>
```

For example:

```
CN=A. N. Other
```

Alternatively, you can use the equivalent OID to identify the attribute type in the string representation (see Table B.1, “Commonly Used Attribute Types”). For example:

```
2.5.4.3=A. N. Other
```

RDN

A *relative distinguished name* (RDN) represents a single node of a DN (the bit that appears between the commas in the string representation). Technically, an RDN might contain more than one AVA (it is formally defined as a set of AVAs). However, this almost never occurs in practice. In the string representation, an RDN has the following syntax:

```
<attr-type>=<attr-value>[+<attr-type>=<attr-value> ...]
```

Here is an example of a (very unlikely) multiple-value RDN:

```
OU=Eng1+OU=Eng2+OU=Eng3
```

Here is an example of a single-value RDN:

```
OU=Engineering
```