



Fuse ESB

Web Services Security Guide

Version 4.4.1
Sept. 2011

Web Services Security Guide

Version 4.4.1

Updated: 06 Jun 2013

Copyright © 2011-2013 Red Hat, Inc. and/or its affiliates.

Trademark Disclaimer

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Apache, ServiceMix, Camel, CXF, and ActiveMQ are trademarks of Apache Software Foundation. Any other names contained herein may be trademarks of their respective owners.

Third Party Acknowledgements

One or more products in the Red Hat JBoss Fuse release includes third party components covered by licenses that require that the following documentation notices be provided:

- JLine (<http://jline.sourceforge.net>) jline:jline:jar:1.0

License: BSD (LICENSE.txt) - Copyright (c) 2002-2006, Marc Prud'hommeaux <mwp1@cornell.edu>

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of JLine nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR

SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- Stax2 API (<http://woodstox.codehaus.org/StAX2>) org.codehaus.woodstox:stax2-api:jar:3.1.1

License: The BSD License (<http://www.opensource.org/licenses/bsd-license.php>)

Copyright (c) <YEAR>, <OWNER> All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- jibx-run - JiBX runtime (<http://www.jibx.org/main-reactor/jibx-run>) org.jibx:jibx-run:bundle:1.2.3

License: BSD (<http://jibx.sourceforge.net/jibx-license.html>) Copyright (c) 2003-2010, Dennis M. Sosnoski.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of JiBX nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON

ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- JavaAssist (<http://www.jboss.org/javassist>) org.jboss.javassist:com.springsource.javassist:jar:3.9.0.GA:compile
License: MPL (<http://www.mozilla.org/MPL/MPL-1.1.html>)
- HAPI-OSGI-Base Module (<http://hl7api.sourceforge.net/hapi-osgi-base/>) ca.uhn.hapi:hapi-osgi-base:bundle:1.2
License: Mozilla Public License 1.1 (<http://www.mozilla.org/MPL/MPL-1.1.txt>)

Table of Contents

1. Security for HTTP-Compatible Bindings	13
2. Managing Certificates	21
What is an X.509 Certificate?	22
Certification Authorities	24
Choice of CAs	25
Commercial Certification Authorities	26
Private Certification Authorities	27
Certificate Chaining	28
PKCS#12 Files	30
Special Requirements on HTTPS Certificates	32
Creating Your Own Certificates	35
Prerequisites	36
Set Up Your Own CA	37
Use the CA to Create Signed Certificates in a Java Keystore	41
Use the CA to Create Signed PKCS#12 Certificates	44
3. Configuring HTTPS	49
Authentication Alternatives	50
Target-Only Authentication	51
Mutual Authentication	53
Specifying Trusted CA Certificates	55
When to Deploy Trusted CA Certificates	56
Specifying Trusted CA Certificates for HTTPS	57
Specifying an Application's Own Certificate	59
Deploying Own Certificate for HTTPS	60
4. Configuring HTTPS Cipher Suites	63
Supported Cipher Suites	64
Cipher Suite Filters	66
SSL/TLS Protocol Version	69
5. The WS-Policy Framework	71
Introduction to WS-Policy	72
Policy Expressions	76
6. Message Protection	81
Transport Layer Message Protection	82
SOAP Message Protection	86
Introduction to SOAP Message Protection	87
Basic Signing and Encryption Scenario	89
Specifying an AsymmetricBinding Policy	91
Specifying a SymmetricBinding Policy	97
Specifying Parts of Message to Encrypt and Sign	101
Providing Encryption Keys and Signing Keys	104
Specifying the Algorithm Suite	111

7. Authentication	117
Introduction to Authentication	118
Specifying an Authentication Policy	119
Providing Client Credentials	127
Authenticating Received Credentials	132
8. WS-Trust	135
Introduction to WS-Trust	136
Basic Scenarios	139
WS-Trust Single Sign-On Demonstration	143
Overview of the Demonstration	144
Configure the Security Token Service	146
Define the Security Policy	152
Configure the Client-STS Connection	155
Configure the Server-Side Interceptor	160
Build and Run the Demonstration	165
Sample Message Exchanges	166
Defining an IssuedToken Policy	173
Creating an STSClient Instance	178
A. ASN.1 and Distinguished Names	181
ASN.1	182
Distinguished Names	183
B. OpenSSL Utilities	187
Using OpenSSL Utilities	188
Utilities Overview	189
The x509 Utility	191
The req Utility	193
The rsa Utility	195
The ca Utility	197
The s_client Utility	199
The s_server Utility	202
The OpenSSL Configuration File	206
Configuration Overview	207
[req] Variables	208
[ca] Variables	209
[policy] Variables	210
Example openssl.cnf File	212
C. Licenses	215
OpenSSL License	216
Index	219

List of Figures

2.1. A Certificate Chain of Depth 2	28
2.2. A Certificate Chain of Depth 3	28
2.3. Elements in a PKCS#12 File	30
3.1. Target Authentication Only	51
3.2. Mutual Authentication	53
6.1. Basic Signing and Encryption Scenario	89
8.1. WS-Trust Architecture	137
8.2. Server-Vouches Scenario	140
8.3. Bearer Scenario	141
8.4. Holder-of-Key Scenario	142
8.5. WS-Trust Single Sign-On Scenario	144
8.6. Installing Requisite Certificates and WSDL File	146
8.7. Injecting Parameters into the Outgoing RequestSecurityToken Message	173

List of Tables

4.1. Namespaces Used for Configuring Cipher Suite Filters	66
4.2. SSL/TLS Protocols Supported by SUN's JSSE Provider	69
6.1. Encryption and Signing Properties	104
6.2. WSS4J Keystore Properties	107
6.3. Properties for Specifying Crypto Objects	108
6.4. Algorithm Suites	112
6.5. Key Length Properties	116
7.1. Client Credentials Properties	127
8.1. XML Namespaces used with IssuedToken	173
8.2. Token Type URIs	176
A.1. Commonly Used Attribute Types	184

List of Examples

1.1. Specifying HTTPS in the WSDL	14
1.2. Specifying HTTPS in the Server Code	15
1.3. Sample HTTPS Client with No Certificate	15
1.4. Sample HTTPS Client with Certificate	16
1.5. Sample HTTPS Server Configuration	18
4.1. Structure of a sec:cipherSuitesFilter Element	66
5.1. The Empty Policy	78
5.2. The Null Policy	78
5.3. Normal Form Syntax	79
6.1. Client HTTPS Configuration in Spring	82
6.2. Server HTTPS Configuration in Spring	83
6.3. Example of a Transport Binding	84
6.4. Example of an Asymmetric Binding	92
6.5. Example of a Symmetric Binding	98
6.6. Integrity and Encryption Policy Assertions	102
6.7. WSS4J Crypto Interface	109
7.1. Example of a Supporting Tokens Policy	120
7.2. Callback Handler for UsernameToken Passwords	128
8.1. Sample Security Policy for Single Sign-On	152

Chapter 1. Security for HTTP-Compatible Bindings

This chapter describes the security features supported by the Fuse Services Framework HTTP transport. These security features are available to any Fuse Services Framework binding that can be layered on top of the HTTP transport.

Overview

This section describes how to configure the HTTP transport to use SSL/TLS security, a combination usually referred to as HTTPS. In Fuse Services Framework, HTTPS security is configured by specifying settings in XML configuration files.

The following topics are discussed in this chapter:

- ["Generating X.509 certificates"](#)
- ["Enabling HTTPS"](#)
- ["HTTPS client with no certificate"](#)
- ["HTTPS client with certificate"](#)
- ["HTTPS server configuration"](#)

Generating X.509 certificates

A basic prerequisite for using SSL/TLS security is to have a collection of X.509 certificates available to identify your server applications and, optionally, to identify your client applications. You can generate X.509 certificates in one of the following ways:

- Use a commercial third-party tool to generate and manage your X.509 certificates.
- Use the free **openssl** utility (which can be downloaded from <http://www.openssl.org>) and the Java **keystore** utility to generate certificates (see ["Use the CA to Create Signed Certificates in a Java Keystore"](#) on page 41).



Note

The HTTPS protocol mandates a *URL integrity check*, which requires a certificate's identity to match the hostname on which the server is deployed. See ["Special Requirements on HTTPS Certificates" on page 32](#) for details.

Certificate format

In the Java runtime, you must deploy X.509 certificate chains and trusted CA certificates in the form of Java keystores. See ["Configuring HTTPS" on page 49](#) for details.

Enabling HTTPS

A prerequisite for enabling HTTPS on a WSDL endpoint is that the endpoint address must be specified as a HTTPS URL. There are two different locations where the endpoint address is set and both must be modified to use a HTTPS URL:

- HTTPS specified in the WSDL contract—you must specify the endpoint address in the WSDL contract to be a URL with the `https:` prefix, as shown in [Example 1.1 on page 14](#).

Example 1.1. Specifying HTTPS in the WSDL

```
<wsdl:definitions name="HelloWorld"
    targetNamespace="http://apache.org/hello_world_soap_http"
    xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
    xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/" ... >
  ...
  <wsdl:service name="SOAPService">
    <wsdl:port binding="tns:Greeter_SOAPBinding"
      name="SoapPort">
      <soap:address location="https://localhost:9001/SoapContext/SoapPort"/>
    </wsdl:port>
  </wsdl:service>
</wsdl:definitions>
```

Where the `location` attribute of the `soap:address` element is configured to use a HTTPS URL. For bindings other than SOAP, you edit the URL appearing in the `location` attribute of the `http:address` element.

- HTTPS specified in the server code—you must ensure that the URL published in the server code by calling `Endpoint.publish()` is defined with a `https:` prefix, as shown in [Example 1.2 on page 15](#).

Example 1.2. Specifying HTTPS in the Server Code

```
// Java
package demo.hw_https.server;
import javax.xml.ws.Endpoint;

public class Server {
    protected Server() throws Exception {
        Object implementor = new GreeterImpl();
        String address = "https://localhost:9001/SoapContext/SoapPort";
        Endpoint.publish(address, implementor);
    }
    ...
}
```

HTTPS client with no certificate

For example, consider the configuration for a secure HTTPS client with no certificate, as shown in [Example 1.3 on page 15](#).

Example 1.3. Sample HTTPS Client with No Certificate

```
<?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:sec="http://cxf.apache.org/configuration/security"
    xmlns:http="http://cxf.apache.org/transports/http/configuration"
    xmlns:jaxws="http://java.sun.com/xml/ns/jaxws"
    xsi:schemaLocation="...">

❶ <http:conduit name="{http://apache.org/hello_world_soap_http}SoapPort.http-conduit">
❷   <http:tlsClientParameters>
❸     <sec:trustManagers>
        <sec:keyStore type="JKS" password="password"
            file="certs/truststore.jks"/>
      </sec:trustManagers>
❹     <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>
```

The preceding client configuration is described as follows:

- ❶ The TLS security settings are defined on a specific WSDL port. In this example, the WSDL port being configured has the QName, {http://apache.org/hello_world_soap_http}SoapPort.
- ❷ The http:tlsClientParameters element contains all of the client's TLS configuration details.
- ❸ The sec:trustManagers element is used to specify a list of trusted CA certificates (the client uses this list to decide whether or not to trust certificates received from the server side).

The file attribute of the sec:keyStore element specifies a Java keystore file, truststore.jks, containing one or more trusted CA certificates. The password attribute specifies the password required to access the keystore, truststore.jks. See ["Specifying Trusted CA Certificates for HTTPS" on page 57](#).



Note

Instead of the file attribute, you can specify the location of the keystore using either the resource attribute or the url attribute. You must be extremely careful not to load the truststore from an untrustworthy source.

- ❹ The sec:cipherSuitesFilter element can be used to narrow the choice of cipher suites that the client is willing to use for a TLS connection. See ["Configuring HTTPS Cipher Suites" on page 63](#) for details.

HTTPS client with certificate

Consider a secure HTTPS client that is configured to have its own certificate. [Example 1.4 on page 16](#) shows how to configure such a sample client.

Example 1.4. Sample HTTPS Client with Certificate

```
<?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:sec="http://cxf.apache.org/configuration/security"
  xmlns:http="http://cxf.apache.org/transports/http/configuration"
  xmlns:jaxws="http://java.sun.com/xml/ns/jaxws"
  xsi:schemaLocation="...">

  <http:conduit name="{http://apache.org/hello_world_soap_http}SoapPort.http-conduit">
    <http:tlsClientParameters>
      <sec:trustManagers>
        <sec:keyStore type="JKS" password="password"
          file="certs/truststore.jks"/>
      </sec:trustManagers>
      ❶ <sec:keyManagers keyPassword="password">
      ❷ <sec:keyStore type="JKS" password="password"
```



```

        file="certs/wibble.jks"/>
    </sec:keyManagers>
    <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>

    <bean id="cxf" class="org.apache.cxf.bus.CXFBusImpl"/>
</beans>

```

The preceding client configuration is described as follows:

- ❶ The `sec:keyManagers` element is used to attach an X.509 certificate and a private key to the client. The password specified by the `keyPasswod` attribute is used to decrypt the certificate's private key.
- ❷ The `sec:keyStore` element is used to specify an X.509 certificate and a private key that are stored in a Java keystore. This sample declares that the keystore is in Java Keystore format (JKS).

The `file` attribute specifies the location of the keystore file, `wibble.jks`, that contains the client's X.509 certificate chain and private key in a *key entry*. The `password` attribute specifies the keystore password which is required to access the contents of the keystore. It is expected that the keystore file contains just one key entry, so it is not necessary to specify a key alias to identify the entry.

For details of how to create such a keystore file, see ["Use the CA to Create Signed Certificates in a Java Keystore" on page 41](#).

Note

Instead of the `file` attribute, you can specify the location of the keystore using either the `resource` attribute or the `url` attribute. You must be extremely careful not to load the truststore from an untrustworthy source.

HTTPS server configuration

Consider a secure HTTPS server that requires clients to present an X.509 certificate. [Example 1.5 on page 18](#) shows how to configure such a server.

Example 1.5. Sample HTTPS Server Configuration

```

<?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:sec="http://cxf.apache.org/configuration/security"
xmlns:http="http://cxf.apache.org/transport/http/configuration"
xmlns:httpj="http://cxf.apache.org/transport/http-jetty/configuration"
xmlns:jaxws="http://java.sun.com/xml/ns/jaxws"
xsi:schemaLocation="...">

  <httpj:engine-factory bus="cxf">
❶   <httpj:engine port="9001">
❷     <httpj:tlsServerParameters>
❸       <sec:keyManagers keyPassword="password">
❹         <sec:keyStore type="JKS" password="password"
           file="certs/cherry.jks"/>
        </sec:keyManagers>
❺       <sec:trustManagers>
        <sec:keyStore type="JKS" password="password"
           file="certs/truststore.jks"/>
        </sec:trustManagers>
❻       <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>
❼       <sec:clientAuthentication want="true" required="true"/>
      </httpj:tlsServerParameters>
    </httpj:engine>
  </httpj:engine-factory>

  <!-- We need a bean named "cxf" -->
  <bean id="cxf" class="org.apache.cxf.bus.CXFBusImpl"/>
</beans>

```

The preceding server configuration is described as follows:

- ❶ On the server side, TLS is *not* configured for each WSDL port. Instead of configuring each WSDL port, the TLS security settings are applied to a specific *IP port*, which is 9001 in this example. All of the WSDL ports that share this IP port are therefore configured with the same TLS security settings.
- ❷ The `http:tlsServerParameters` element contains all of the server's TLS configuration details.
- ❸ The `sec:keyManagers` element is used to attach an X.509 certificate and a private key to the server. The password specified by the `keyPassword` attribute is used to decrypt the certificate's private key.
- ❹ The `sec:keyStore` element is used to specify an X.509 certificate and a private key that are stored in a Java keystore. This sample declares that the keystore is in Java Keystore format (JKS).

The `file` attribute specifies the location of the keystore file, `cherry.jks`, that contains the client's X.509 certificate chain and private key in a *key entry*. The `password` attribute specifies the keystore password, which is needed to access the contents of the keystore. It is expected that the keystore file contains just one key entry, so there is no need to specify a key alias.



Note

Instead of the `file` attribute, you can specify the location of the keystore using either the `resource` attribute or the `url` attribute. You must be extremely careful not to load the truststore from an untrustworthy source.

For details of how to create such a keystore file, see ["Use the CA to Create Signed Certificates in a Java Keystore" on page 41](#).

- ⑥ The `sec:trustManagers` element is used to specify a list of trusted CA certificates (the server uses this list to decide whether or not to trust certificates presented by clients).

The `file` attribute of the `sec:keyStore` element specifies a Java keystore file, `truststore.jks`, containing one or more trusted CA certificates. The `password` attribute specifies the password required to access the keystore, `truststore.jks`. See ["Specifying Trusted CA Certificates for HTTPS" on page 57](#).



Note

Instead of the `file` attribute, you can specify the location of the keystore using either the `resource` attribute or the `url` attribute.

- ⑥ The `sec:cipherSuitesFilter` element can be used to narrow the choice of cipher suites that the server is willing to use for a TLS connection. See ["Configuring HTTPS Cipher Suites" on page 63](#) for details.
- ⑦ The `sec:clientAuthentication` element determines the server's disposition towards the presentation of client certificates. The element has the following attributes:
 - `want` attribute—If `true` (the default), the server requests the client to present an X.509 certificate during the TLS handshake; if `false`, the server does *not* request the client to present an X.509 certificate.
 - `required` attribute—If `true`, the server raises an exception if a client fails to present an X.509 certificate during the TLS handshake; if `false` (the default), the server does *not* raise an exception if the client fails to present an X.509 certificate.

Chapter 2. Managing Certificates

TLS authentication uses X.509 certificates—a common, secure and reliable method of authenticating your application objects. This chapter explains how to create X.509 certificates that identify your Fuse Services Framework applications.

What is an X.509 Certificate?	22
Certification Authorities	24
Choice of CAs	25
Commercial Certification Authorities	26
Private Certification Authorities	27
Certificate Chaining	28
PKCS#12 Files	30
Special Requirements on HTTPS Certificates	32
Creating Your Own Certificates	35
Prerequisites	36
Set Up Your Own CA	37
Use the CA to Create Signed Certificates in a Java Keystore	41
Use the CA to Create Signed PKCS#12 Certificates	44

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 demonstration certificates supplied with Fuse Services Framework are signed by the demonstration CA. This CA is completely insecure because anyone can access its private key. To secure your system, you must create new certificates signed by a trusted CA. This chapter describes the set of certificates required by a Fuse Services Framework application and describes how to replace the default certificates.

The 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:

- X.509 version information.

- A *serial number* that uniquely identifies the certificate.
- A *subject distinguished name (DN)* that identifies the certificate owner.
- The *public key* associated with the subject.
- 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 A on page 181](#) for more details about DNs.

Certification Authorities

Choice of CAs	25
Commercial Certification Authorities	26
Private Certification Authorities	27

Choice of CAs

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 Fuse Services Framework 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:

- A *commercial CA* is a company that signs certificates for many systems.
- A *private CA* is a trusted node that you set up and use to sign certificates for your system only.

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 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?

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>. OpenSSL is derived from SSLeay, an implementation of SSL developed by Eric Young (<ey@cryptsoft.com>). Complete license information can be found in [Appendix C on page 215](#). 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 ["Creating Your Own Certificates" on page 35](#).

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 Fuse Services Framework 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.

Certificate Chaining

Certificate chain

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

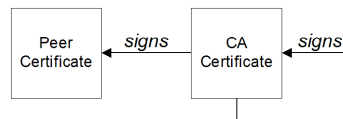
Self-signed certificate

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

Example

Figure 2.1 on page 28 shows an example of a simple certificate chain.

Figure 2.1. A Certificate Chain of Depth 2



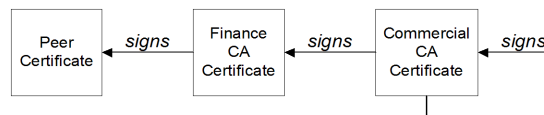
Chain of trust

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 2.2 on page 28 shows what this certificate chain looks like.

Figure 2.2. 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.

See ["Specifying Trusted CA Certificates"](#) on page 55.

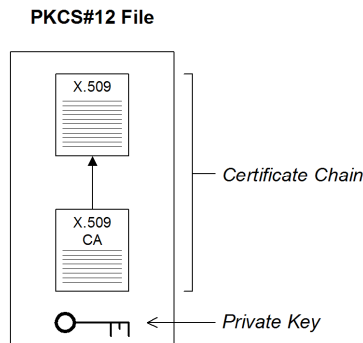
PKCS#12 Files

Overview

PKCS#12 is an industry-standard format for deploying certificates and private keys as a file.

Figure 2.3 on page 30 shows the typical elements in a PKCS#12 file.

Figure 2.3. Elements in a PKCS#12 File



Contents of a PKCS#12 file

A PKCS#12 file contains the following:

- An X.509 peer certificate (first in a chain).
- All the CA certificates in the certificate chain.
- A private key.

The file is encrypted with a pass phrase.



Note

The same pass phrase is used both for the encryption of the private key within the PKCS#12 file, and for the encryption of the PKCS#12 file overall. This condition (same pass phrase) is not officially part of the PKCS#12 standard, but it is enforced by most Web browsers and by Fuse Services Framework.

Creating a PKCS#12 file

To create a PKCS#12 file, see ["Use the CA to Create Signed Certificates in a Java Keystore" on page 41](#) .

Viewing a PKCS#12 file

To view a PKCS#12 file, *CertName.p12*, enter the following command:

```
openssl pkcs12 -in CertName.p12
```

Importing and exporting PKCS#12 files

The generated PKCS#12 files generated by OpenSSL can be imported into browsers such as Internet Explorer or Firefox. Exported PKCS#12 files from these browsers can be used in Fuse Services Framework.



Note

Use OpenSSL v0.9.2 or later.

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.progress.com/secure
```

The corresponding server certificate would have the following subject DN:

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

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

For details of how to set the subject DN in a new certificate, see ["Use the CA to Create Signed Certificates in a Java Keystore" on page 41](#) and ["Use the CA to Create Signed Certificates in a Java Keystore" on page 41](#).

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.progress.com
fusesource.com
```

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.progress.com,DNS:fusesource.com
```

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:*.progress.com
```

This certificate identity matches any three-component host name in the domain `progress.com`. For example, the wildcarded host name matches either `www.progress.com` or `fusesource.com`, but does not match `www.fusesource.com`.



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 `*progress.com`, your certificate could be used on *any* domain that ends in the letters `progress`.

For details of how to set up the `openssl.cnf` configuration file to generate certificates with the `subjectAltName` certificate extension, see ["Use the CA to Create Signed PKCS#12 Certificates" on page 44](#).

Creating Your Own Certificates

Prerequisites	36
Set Up Your Own CA	37
Use the CA to Create Signed Certificates in a Java Keystore	41
Use the CA to Create Signed PKCS#12 Certificates	44

Prerequisites

OpenSSL utilities

The steps described in this section are based on the OpenSSL command-line utilities from the OpenSSL project, <http://www.openssl.org> (see [Appendix B on page 187](#)). Further documentation of the OpenSSL command-line utilities can be obtained at <http://www.openssl.org/docs>.

Sample CA directory structure

For the purposes of illustration, the CA database is assumed to have the following directory structure:

```
x509CA/ca  
x509CA/certs  
x509CA/newcerts  
x509CA/cr1
```

Where *x509CA* is the parent directory of the CA database.

Set Up Your Own CA

Substeps to perform

This section describes how to set up your own private CA. Before setting up a CA for a real deployment, read the additional notes in ["Choosing a host for a private certification authority" on page 27](#) .

To set up your own CA, perform the following steps:

1. ["Add the bin directory to your PATH"](#)
2. ["Create the CA directory hierarchy"](#)
3. ["Copy and edit the openssl.cnf file"](#)
4. ["Initialize the CA database"](#)
5. ["Create a self-signed CA certificate and private key"](#)

Add the bin directory to your PATH

On the secure CA host, add the OpenSSL bin directory to your path:

Windows

```
> set PATH=OpenSSLDir\bin;%PATH%
```

UNIX

```
% PATH=OpenSSLDir/bin:$PATH; export PATH
```

This step makes the **openssl** utility available from the command line.

Create the CA directory hierarchy

Create a new directory, *X509CA*, to hold the new CA. This directory is used to hold all of the files associated with the CA. Under the *X509CA* directory, create the following hierarchy of directories:

```
X509CA/ca
X509CA/certs
X509CA/newcerts
X509CA/crl
```

Copy and edit the openssl.cnf file

Copy the sample openssl.cnf from your OpenSSL installation to the *x509CA* directory.

Edit the openssl.cnf to reflect the directory structure of the *x509CA* directory, and to identify the files used by the new CA.

Edit the [CA_default] section of the openssl.cnf file to look like the following:

```
#####
[ CA_default ]

dir            = x509CA           # Where CA files are kept
certs         = $dir/certs       # Where issued certs are kept
crl_dir       = $dir/crl         # Where the issued crl are kept
database      = $dir/index.txt   # Database index file
new_certs_dir = $dir/newcerts    # Default place for new certs

certificate    = $dir/ca/new_ca.pem # The CA certificate
serial        = $dir/serial       # The current serial number
crl            = $dir/crl.pem     # The current CRL
private_key    = $dir/ca/new_ca_pk.pem # The private key
RANDFILE      = $dir/ca/.rand    # Private random number file

x509_extensions = usr_cert      # The extensions to add to the cert
...
```

You might decide to edit other details of the OpenSSL configuration at this point—for more details, see ["The OpenSSL Configuration File" on page 206](#).

Initialize the CA database

In the *x509CA* directory, initialize two files, serial and index.txt.

Windows

To initialize the serial file in Windows, enter the following command:

```
> echo 01 > serial
```

To create an empty file, index.txt, in Windows start Windows Notepad at the command line in the *x509CA* directory, as follows:

```
> notepad index.txt
```

In response to the dialog box with the text, Cannot find the text.txt file. Do you want to create a new file?, click **Yes**, and close Notepad.

UNIX

To initialize the serial file and the index.txt file in UNIX, enter the following command:

```
% echo "01" > serial
% touch index.txt
```

These files are used by the CA to maintain its database of certificate files.



Note

The index.txt file must initially be completely empty, not even containing white space.

Create a self-signed CA certificate and private key

Create a new self-signed CA certificate and private key with the following command:

```
openssl req -x509 -new -config x509CA/openssl.cnf -days 365 -out x509CA/ca/new_ca.pem -keyout
x509CA/ca/new_ca_pk.pem
```

The command prompts you for a pass phrase for the CA private key and details of the CA distinguished name. For example:

```
Using configuration from x509CA/openssl.cnf
Generating a 512 bit RSA private key
...+++++
.+++++
writing new private key to 'new_ca_pk.pem'
Enter PEM pass phrase:
Verifying password - 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) []:IE
State or Province Name (full name) []:Co. Dublin
Locality Name (eg, city) []:Dublin
Organization Name (eg, company) []:Progress
```

```
Organizational Unit Name (eg, section) []:Finance  
Common Name (eg, YOUR name) []:Gordon Brown  
Email Address []:gbrown@progress.com
```



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, `new_ca.pem` and `new_ca_pk.pem`, are the same as the values specified in `openssl.cnf` (see the preceding step).

You are now ready to sign certificates with your CA.

Use the CA to Create Signed Certificates in a Java Keystore

Substeps to perform

To create and sign a certificate in a Java keystore (JKS), *CertName.jks*, perform the following substeps:

1. "Add the Java bin directory to your PATH"
2. "Generate a certificate and private key pair"
3. "Create a certificate signing request"
4. "Sign the CSR"
5. "Convert to PEM format"
6. "Concatenate the files"
7. "Update keystore with the full certificate chain"
8. "Repeat steps as required"

Add the Java bin directory to your PATH

If you have not already done so, add the Java bin directory to your path:

Windows

```
> set PATH=JAVA_HOME\bin;%PATH%
```

UNIX

```
% PATH=JAVA_HOME/bin:$PATH; export PATH
```

This step makes the **keytool** utility available from the command line.

Generate a certificate and private key pair

Open a command prompt and change directory to the directory where you store your keystore files, *keystoreDir*. Enter the following command:

```
keytool -genkey -dname "CN=Alice, OU=Engineering, O=Progress, ST=Co. Dublin, C=IE" -validity  
365 -alias CertAlias -keypass CertPassword -keystore CertName.jks -storepass CertPassword
```

This `keytool` command, invoked with the `-genkey` option, generates an X.509 certificate and a matching private key. The certificate and the key are both placed in a *key entry* in a newly created keystore, `CertName.jks`. Because the specified keystore, `CertName.jks`, did not exist prior to issuing the command, **keytool** implicitly creates a new keystore.

The `-dname` and `-validity` flags define the contents of the newly created X.509 certificate, specifying the subject DN and the days before expiration respectively. For more details about DN format, see [Appendix A on page 181](#).

Some parts of the subject DN must match the values in the CA certificate (specified in the CA Policy section of the `openssl.cnf` file). The default `openssl.cnf` file requires the following entries to match:

- Country Name (C)
- State or Province Name (ST)
- Organization Name (O)



Note

If you do not observe the constraints, the OpenSSL CA will refuse to sign the certificate (see ["Sign the CSR" on page 42](#)).

Create a certificate signing request

Create a new certificate signing request (CSR) for the `CertName.jks` certificate, as follows:

```
keytool -certreq -alias CertAlias -file CertName_csr.pem -keypass CertPassword -keystore CertName.jks -storepass CertPassword
```

This command exports a CSR to the file, `CertName_csr.pem`.

Sign the CSR

Sign the CSR using your CA, as follows:

```
openssl ca -config x509CA/openssl.cnf -days 365 -in CertName_csr.pem -out CertName.pem
```

To sign the certificate successfully, you must enter the CA private key pass phrase (see ["Set Up Your Own CA" on page 37](#)).



Note

If you want to sign the CSR using a CA certificate *other* than the default CA, use the `-cert` and `-keyfile` options to specify the CA certificate and its private key file, respectively.

Convert to PEM format

Convert the signed certificate, `CertName.pem`, to PEM only format, as follows:

```
openssl x509 -in CertName.pem -out CertName.pem -outform PEM
```

Concatenate the files

Concatenate the CA certificate file and `CertName.pem` certificate file, as follows:

Windows

```
copy CertName.pem + X509CA\ca\new_ca.pem CertName.chain
```

UNIX

```
cat CertName.pem X509CA/ca/new_ca.pem > CertName.chain
```

Update keystore with the full certificate chain

Update the keystore, `CertName.jks`, by importing the full certificate chain for the certificate, as follows:

```
keytool -import -file CertName.chain -keypass CertPassword -keystore CertName.jks -storepass CertPassword
```

Repeat steps as required

Repeat steps 2 through 7, to create a complete set of certificates for your system.

Use the CA to Create Signed PKCS#12 Certificates

Substeps to perform

If you have set up a private CA, as described in ["Set Up Your Own CA" on page 37](#) , you are now ready to create and sign your own certificates.

To create and sign a certificate in PKCS#12 format, *CertName.p12*, perform the following substeps:

1. ["Add the bin directory to your PATH"](#) .
2. ["Configure the subjectAltName extension \(Optional\)"](#) .
3. ["Create a certificate signing request"](#) .
4. ["Sign the CSR"](#) .
5. ["Concatenate the files"](#) .
6. ["Create a PKCS#12 file"](#) .
7. ["Repeat steps as required"](#) .
8. ["\(Optional\) Clear the subjectAltName extension"](#) .

Add the bin directory to your PATH

If you have not already done so, add the OpenSSL bin directory to your path, as follows:

Windows

```
> set PATH=OpenSSLDir\bin;%PATH%
```

UNIX

```
% PATH=OpenSSLDir/bin:$PATH; export PATH
```

This step makes the **openssl** utility available from the command line.

Configure the subjectAltName extension (Optional)

Perform this step, if the certificate is intended for a HTTPS server whose clients enforce URL integrity check, and if you plan to deploy the server on a multi-homed host or a host with several DNS name aliases (for example, if you are deploying the certificate on a multi-homed Web server). In this case, the certificate identity

must match multiple host names and this can be done only by adding a `subjectAltName` certificate extension (see ["Special Requirements on HTTPS Certificates" on page 32](#)).

To configure the `subjectAltName` extension, edit your CA's `openssl.cnf` file as follows:

1. Add the following `req_extensions` setting to the `[req]` section (if not already present in your `openssl.cnf` file):

```
# openssl Configuration File
...
[req]
req_extensions=v3_req
```

2. Add the `[v3_req]` section header (if not already present in your `openssl.cnf` file). Under the `[v3_req]` section, add or modify the `subjectAltName` setting, setting it to the list of your DNS host names. For example, if the server host supports the alternative DNS names, `www.progress.com` and `fusesource.com`, set the `subjectAltName` as follows:

```
# openssl Configuration File
...
[v3_req]
subjectAltName=DNS:www.progress.com,DNS:fusesource.com
```

3. Add a `copy_extensions` setting to the appropriate CA configuration section. The CA configuration section used for signing certificates is one of the following:

- The section specified by the `-name` option of the **openssl ca** command,
- The section specified by the `default_ca` setting under the `[ca]` section (usually `[CA_default]`).

For example, if the appropriate CA configuration section is `[CA_default]`, set the `copy_extensions` property as follows:

```
# openssl Configuration File
...
[CA_default]
copy_extensions=copy
```

This setting ensures that certificate extensions present in the certificate signing request are copied into the signed certificate.

Create a certificate signing request

Create a new certificate signing request (CSR) for the `CertName.p12` certificate, as shown:

```
openssl req -new -config x509CA/openssl.cnf -days 365 -out x509CA/certs/CertName_csr.pem -  
keyout x509CA/certs/CertName_pk.pem
```

This command prompts you for a pass phrase for the certificate's private key, and for information about the certificate's distinguished name.

Some of the entries in the CSR distinguished name must match the values in the CA certificate (specified in the CA Policy section of the `openssl.cnf` file). The default `openssl.cnf` file requires that the following entries match:

- Country Name
- State or Province Name
- Organization Name

The certificate subject DN's Common Name is the field that is usually used to represent the certificate owner's identity. The Common Name must comply with the following conditions:

- The Common Name must be *distinct* for every certificate generated by the OpenSSL certificate authority.
- If your HTTPS clients implement the URL integrity check, you must ensure that the Common Name is identical to the DNS name of the host where the certificate is to be deployed (see ["Special Requirements on HTTPS Certificates" on page 32](#)).



Note

For the purpose of the HTTPS URL integrity check, the `subjectAltName` extension takes precedence over the Common Name.

```
Using configuration from x509CA/openssl.cnf  
Generating a 512 bit RSA private key  
.+++++  
.+++++  
writing new private key to  
    'x509CA/certs/CertName_pk.pem'  
Enter PEM pass phrase:  
Verifying password - 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) []:IE
State or Province Name (full name) []:Co. Dublin
Locality Name (eg, city) []:Dublin
Organization Name (eg, company) []:Progress
Organizational Unit Name (eg, section) []:Systems
Common Name (eg, YOUR name) []:Artix
Email Address []:info@progress.com

```

Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:password
An optional company name []:Progress

Sign the CSR

Sign the CSR using your CA, as follows:

```
openssl ca -config X509CA/openssl.cnf -days 365 -in X509CA/certs/CertName_csr.pem -out
X509CA/certs/CertName.pem
```

This command requires the pass phrase for the private key associated with the new_ca.pem CA certificate.
For example:

```

Using configuration from X509CA/openssl.cnf
Enter PEM pass phrase:
Check that the request matches the signature
Signature ok
The Subjects Distinguished Name is as follows
countryName :PRINTABLE:'IE'
stateOrProvinceName :PRINTABLE:'Co. Dublin'
localityName :PRINTABLE:'Dublin'
organizationName :PRINTABLE:'Progress'
organizationalUnitName:PRINTABLE:'Systems'
commonName :PRINTABLE:'Bank Server Certificate'
emailAddress :IA5STRING:'info@progress.com'
Certificate is to be certified until May 24 13:06:57 2000 GMT (365 days)
Sign the certificate? [y/n]:y
1 out of 1 certificate requests certified, commit? [y/n]y
Write out database with 1 new entries
Data Base Updated

```

To sign the certificate successfully, you must enter the CA private key pass phrase (see ["Set Up Your Own CA" on page 37](#)).



Note

If you did not set `copy_extensions=copy` under the `[CA_default]` section in the `openssl.cnf` file, the signed certificate will not include any of the certificate extensions that were in the original CSR.

Concatenate the files

Concatenate the CA certificate file, `CertName.pem` certificate file, and `CertName_pk.pem` private key file as follows:

Windows

```
copy X509CA\ca\new_ca.pem + X509CA\certs\CertName.pem + X509CA\certs\CertName_pk.pem  
X509CA\certs\CertName_list.pem
```

UNIX

```
cat X509CA/ca/new_ca.pem X509CA/certs/CertName.pem X509CA/certs/CertName_pk.pem >  
X509CA/certs/CertName_list.pem
```

Create a PKCS#12 file

Create a PKCS#12 file from the `CertName_list.pem` file as follows:

```
openssl pkcs12 -export -in X509CA/certs/CertName_list.pem -out X509CA/certs/CertName.p12 -name  
"New cert"
```

You are prompted to enter a password to encrypt the PKCS#12 certificate. Usually this password is the same as the CSR password (this is required by many certificate repositories).

Repeat steps as required

Repeat steps 3 through 6, to create a complete set of certificates for your system.

(Optional) Clear the subjectAltName extension

After generating certificates for a particular host machine, it is advisable to clear the `subjectAltName` setting in the `openssl.cnf` file to avoid accidentally assigning the wrong DNS names to another set of certificates.

In the `openssl.cnf` file, comment out the `subjectAltName` setting (by adding a `#` character at the start of the line), and also comment out the `copy_extensions` setting.

Chapter 3. Configuring HTTPS

This chapter describes how to configure HTTPS endpoints.

Authentication Alternatives	50
Target-Only Authentication	51
Mutual Authentication	53
Specifying Trusted CA Certificates	55
When to Deploy Trusted CA Certificates	56
Specifying Trusted CA Certificates for HTTPS	57
Specifying an Application's Own Certificate	59
Deploying Own Certificate for HTTPS	60

Authentication Alternatives

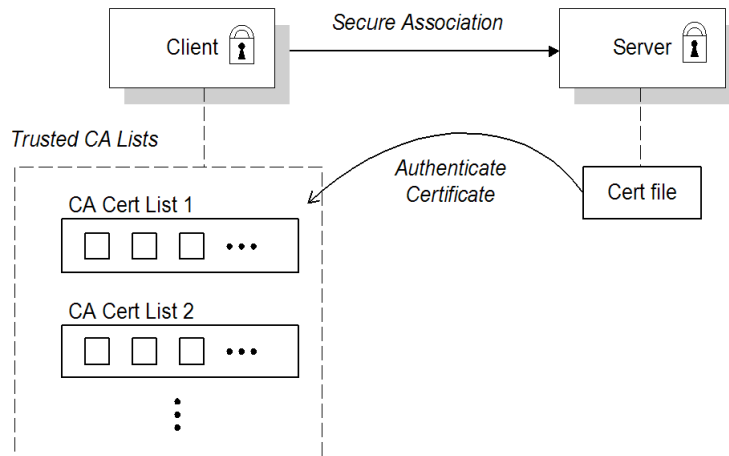
Target-Only Authentication	51
Mutual Authentication	53

Target-Only Authentication

Overview

When an application is configured for target-only authentication, the target authenticates itself to the client but the client is not authentic to the target object, as shown in [Figure 3.1 on page 51](#).

Figure 3.1. Target Authentication Only



Security handshake

Prior to running the application, the client and server should be set up as follows:

- A certificate chain is associated with the server. The certificate chain is provided in the form of a Java keystore (see ["Specifying an Application's Own Certificate" on page 59](#)).
- One or more lists of trusted certification authorities (CA) are made available to the client. (see ["Specifying Trusted CA Certificates" on page 55](#)).

During the security handshake, the server sends its certificate chain to the client (see [Figure 3.1 on page 51](#)). The client then searches its trusted CA lists to find a CA certificate that matches one of the CA certificates in the server's certificate chain.

HTTPS example

On the client side, there are no policy settings required for target-only authentication. Simply configure your client *without* associating an X.509 certificate with the HTTPS port. You must provide the client with a list of trusted CA certificates, however (see ["Specifying Trusted CA Certificates" on page 55](#)).

On the server side, in the server's XML configuration file, make sure that the `sec:clientAuthentication` element does not require client authentication. This element can be omitted, in which case the default policy is to *not* require client authentication. However, if the `sec:clientAuthentication` element is present, it should be configured as follows:

```
<http:destination id="{Namespace}PortName.http-destination">
  <http:tlsServerParameters>
    ...

    <sec:clientAuthentication want="false" required="false"/>
  </http:tlsServerParameters>
</http:destination>
```

Where the `want` attribute is set to `false` (the default), specifying that the server does not request an X.509 certificate from the client during a TLS handshake. The `required` attribute is also set to `false` (the default), specifying that the absence of a client certificate does not trigger an exception during the TLS handshake.



Note

The `want` attribute can be set either to `true` or to `false`. If set to `true`, the `want` setting causes the server to request a client certificate during the TLS handshake, but no exception is raised for clients lacking a certificate, so long as the `required` attribute is set to `false`.

It is also necessary to associate an X.509 certificate with the server's HTTPS port (see ["Specifying an Application's Own Certificate" on page 59](#)) and to provide the server with a list of trusted CA certificates (see ["Specifying Trusted CA Certificates" on page 55](#)).



Note

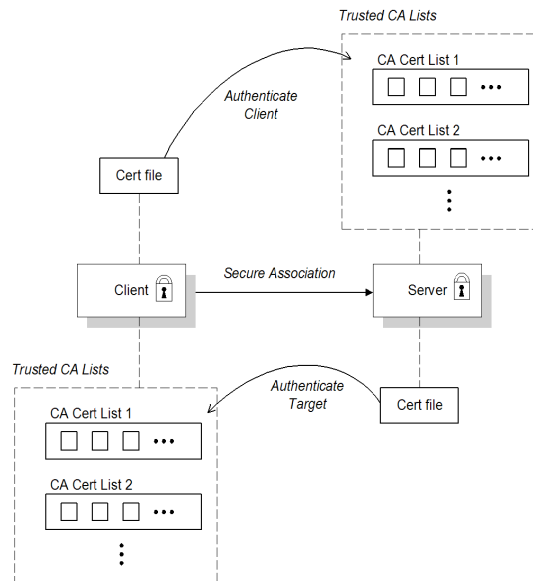
The choice of cipher suite can potentially affect whether or not target-only authentication is supported (see ["Configuring HTTPS Cipher Suites" on page 63](#)).

Mutual Authentication

Overview

When an application is configured for mutual authentication, the target authenticates itself to the client and the client authenticates itself to the target. This scenario is illustrated in [Figure 3.2 on page 53](#) . In this case, the server and the client each require an X.509 certificate for the security handshake.

Figure 3.2. Mutual Authentication



Security handshake

Prior to running the application, the client and server must be set up as follows:

- Both client and server have an associated certificate chain (see ["Specifying an Application's Own Certificate" on page 59](#)).
- Both client and server are configured with lists of trusted certification authorities (CA) (see ["Specifying Trusted CA Certificates" on page 55](#)).

During the TLS handshake, the server sends its certificate chain to the client, and the client sends its certificate chain to the server—see [Figure 3.1 on page 51](#) .

HTTPS example

On the client side, there are no policy settings required for mutual authentication. Simply associate an X.509 certificate with the client's HTTPS port (see ["Specifying an Application's Own Certificate" on page 59](#)). You also need to provide the client with a list of trusted CA certificates (see ["Specifying Trusted CA Certificates" on page 55](#)).

On the server side, in the server's XML configuration file, make sure that the `sec:clientAuthentication` element is configured to *require* client authentication. For example:

```
<http:destination id="{Namespace}PortName.http-destination">
  <http:tlsServerParameters>
    ...
    <sec:clientAuthentication want="true" required="true"/>
  </http:tlsServerParameters>
</http:destination>
```

Where the `want` attribute is set to `true`, specifying that the server requests an X.509 certificate from the client during a TLS handshake. The `required` attribute is also set to `true`, specifying that the absence of a client certificate triggers an exception during the TLS handshake.

It is also necessary to associate an X.509 certificate with the server's HTTPS port (see ["Specifying an Application's Own Certificate" on page 59](#)) and to provide the server with a list of trusted CA certificates (see ["Specifying Trusted CA Certificates" on page 55](#)).



Note

The choice of cipher suite can potentially affect whether or not mutual authentication is supported (see ["Configuring HTTPS Cipher Suites" on page 63](#)).

Specifying Trusted CA Certificates

When to Deploy Trusted CA Certificates	56
Specifying Trusted CA Certificates for HTTPS	57

When to Deploy Trusted CA Certificates

Overview

When an application receives an X.509 certificate during an SSL/TLS handshake, the application decides whether or not to trust the received certificate by checking whether the issuer CA is one of a pre-defined set of trusted CA certificates. If the received X.509 certificate is validly signed by one of the application's trusted CA certificates, the certificate is deemed trustworthy; otherwise, it is rejected.

Which applications need to specify trusted CA certificates?

Any application that is likely to receive an X.509 certificate as part of an HTTPS or IIOP/TLS handshake must specify a list of trusted CA certificates. For example, this includes the following types of application:

- All HTTPS clients.
- Any HTTPS servers that support *mutual authentication*.

Specifying Trusted CA Certificates for HTTPS

CA certificate format

CA certificates must be provided in Java keystore format.

CA certificate deployment in the Fuse Services Framework configuration file

To deploy one or more trusted root CAs for the HTTPS transport, perform the following steps:

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 (for details of how to generate your own CA certificates, see ["Set Up Your Own CA" on page 37](#)). 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. Given a CA certificate, `cacert.pem`, in PEM format, you can add the certificate to a JKS truststore (or create a new truststore) by entering the following command:

```
keytool -import -file cacert.pem -alias CAAlias -keystore truststore.jks -storepass StorePass
```

Where *CAAlias* is a convenient tag that enables you to access this particular CA certificate using the `keytool` utility. The file, `truststore.jks`, is a keystore file containing CA certificates—if this file does not already exist, the `keytool` utility creates one. The *StorePass* password provides access to the keystore file, `truststore.jks`.

3. Repeat step 2 as necessary, to add all of the CA certificates to the truststore file, `truststore.jks`.
4. Edit the relevant XML configuration files to specify the location of the truststore file. You must include the `sec:trustManagers` element in the configuration of the relevant HTTPS ports.

For example, you can configure a client port as follows:

```
<!-- Client port configuration -->
<http:conduit id="{Namespace}PortName.http-conduit">
  <http:tlsClientParameters>
    ...
    <sec:trustManagers>
      <sec:keyStore type="JKS"
                    password="StorePass"
                    file="certs/truststore.jks"/>
    </sec:trustManagers>
    ...
  </http:tlsClientParameters>
```

```
</http:conduit>
```

Where the `type` attribute specifies that the truststore uses the JKS keystore implementation and `StorePass` is the password needed to access the `truststore.jks` keystore.

Configure a server port as follows:

```
<!-- Server port configuration -->
<http:destination id="{Namespace}PortName.http-destination">
  <http:tlsServerParameters>
    ...
    <sec:trustManagers>
      <sec:keyStore type="JKS"
                    password="{StorePass}"
                    file="certs/truststore.jks"/>
    </sec:trustManagers>
    ...
  </http:tlsServerParameters>
</http:destination>
```



Warning

The directory containing the truststores (for example, `X509Deploy/truststores/`) should be a secure directory (that is, writable only by the administrator).

Specifying an Application's Own Certificate

Deploying Own Certificate for HTTPS 60

Deploying Own Certificate for HTTPS

Overview

When working with the HTTPS transport the application's certificate is deployed using the XML configuration file.

Procedure

To deploy an application's own certificate for the HTTPS transport, perform the following steps:

1. Obtain an application certificate in Java keystore format, *CertName.jks*. For instructions on how to create a certificate in Java keystore format, see ["Use the CA to Create Signed Certificates in a Java Keystore" on page 41](#).



Note

Some HTTPS clients (for example, Web browsers) perform a *URL integrity check*, which requires a certificate's identity to match the hostname on which the server is deployed. See ["Special Requirements on HTTPS Certificates" on page 32](#) for details.

2. Copy the certificate's keystore, *CertName.jks*, to the certificates directory on the deployment host; for example, *X509Deploy/certs*.

The certificates directory should be a secure directory that is writable only by administrators and other privileged users.

3. Edit the relevant XML configuration file to specify the location of the certificate keystore, *CertName.jks*. You must include the `sec:keyManagers` element in the configuration of the relevant HTTPS ports.

For example, you can configure a client port as follows:

```
<http:conduit id="{Namespace}PortName.http-conduit">
  <http:tlsClientParameters>
    ...
    <sec:keyManagers keyPassword="CertPassword">
      <sec:keyStore type="JKS"
        password="KeystorePassword"
        file="certs/CertName.jks"/>
    </sec:keyManagers>
    ...
  </http:tlsClientParameters>
</http:conduit>
```

```
</http:tlsClientParameters>
</http:conduit>
```

Where the `keyPassword` attribute specifies the password needed to decrypt the certificate's private key (that is, *CertPassword*), the `type` attribute specifies that the truststore uses the JKS keystore implementation, and the `password` attribute specifies the password required to access the *CertName.jks* keystore (that is, *KeystorePassword*).

Configure a server port as follows:

```
<http:destination id="{Namespace}PortName.http-destination">
  <http:tlsServerParameters>
    ...
    <sec:keyManagers keyPassword="CertPassword">
      <sec:keyStore type="JKS"
        password="KeystorePassword"
        file="certs/CertName.jks"/>
    </sec:keyManagers>
    ...
  </http:tlsServerParameters>
</http:destination>
```



Warning

The directory containing the application certificates (for example, *x509Deploy/certs/*) should be a secure directory (that is, readable and writable only by the administrator).



Warning

The directory containing the XML configuration file should be a secure directory (that is, readable and writable only by the administrator), because the configuration file contains passwords in plain text.

Chapter 4. Configuring HTTPS Cipher Suites

This chapter explains how to specify the list of cipher suites that are made available to clients and servers for the purpose of establishing HTTPS connections. During a security handshake, the client chooses a cipher suite that matches one of the cipher suites available to the server.

Supported Cipher Suites	64
Cipher Suite Filters	66
SSL/TLS Protocol Version	69

Supported Cipher Suites

Overview

A *cipher suite* is a collection of security algorithms that determine precisely how an SSL/TLS connection is implemented.

For example, the SSL/TLS protocol mandates that messages be signed using a message digest algorithm. The choice of digest algorithm, however, is determined by the particular cipher suite being used for the connection. Typically, an application can choose either the MD5 or the SHA digest algorithm.

The cipher suites available for SSL/TLS security in Fuse Services Framework depend on the particular *JSSE provider* that is specified on the endpoint.

JCE/JSSE and security providers

The Java Cryptography Extension (JCE) and the Java Secure Socket Extension (JSSE) constitute a pluggable framework that allows you to replace the Java security implementation with arbitrary third-party toolkits, known as *security providers*.

SunJSSE provider

In practice, the security features of Fuse Services Framework have been tested only with SUN's JSSE provider, which is named SunJSSE.

Hence, the SSL/TLS implementation and the list of available cipher suites in Fuse Services Framework are effectively determined by what is available from SUN's JSSE provider.

Cipher suites supported by SunJSSE

The following cipher suites are supported by SUN's JSSE provider in the J2SE 1.5.0 Java development kit (see also [Appendix A](#)¹ of SUN's *JSSE Reference Guide*):

- Standard ciphers:

```
SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA
SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA
SSL_DHE_DSS_WITH_DES_CBC_SHA
SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA
SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
SSL_DHE_RSA_WITH_DES_CBC_SHA
SSL_RSA_EXPORT_WITH_DES40_CBC_SHA
SSL_RSA_EXPORT_WITH_RC4_40_MD5
```

¹ <http://java.sun.com/j2se/1.5.0/docs/guide/security/jsse/JSSERefGuide.html#AppA>


```

SSL_RSA_WITH_3DES_EDE_CBC_SHA
SSL_RSA_WITH_DES_CBC_SHA
SSL_RSA_WITH_RC4_128_MD5
SSL_RSA_WITH_RC4_128_SHA
TLS_DHE_DSS_WITH_AES_128_CBC_SHA
TLS_DHE_DSS_WITH_AES_256_CBC_SHA
TLS_DHE_RSA_WITH_AES_128_CBC_SHA
TLS_DHE_RSA_WITH_AES_256_CBC_SHA
TLS_KRB5_EXPORT_WITH_DES_CBC_40_MD5
TLS_KRB5_EXPORT_WITH_DES_CBC_40_SHA
TLS_KRB5_EXPORT_WITH_RC4_40_MD5
TLS_KRB5_EXPORT_WITH_RC4_40_SHA
TLS_KRB5_WITH_3DES_EDE_CBC_MD5
TLS_KRB5_WITH_3DES_EDE_CBC_SHA
TLS_KRB5_WITH_DES_CBC_MD5
TLS_KRB5_WITH_DES_CBC_SHA
TLS_KRB5_WITH_RC4_128_MD5
TLS_KRB5_WITH_RC4_128_SHA
TLS_RSA_WITH_AES_128_CBC_SHA
TLS_RSA_WITH_AES_256_CBC_SHA

```

- Null encryption, integrity-only ciphers:

```

SSL_RSA_WITH_NULL_MD5
SSL_RSA_WITH_NULL_SHA

```

- Anonymous Diffie-Hellman ciphers (no authentication):

```

SSL_DH_anon_EXPORT_WITH_DES40_CBC_SHA
SSL_DH_anon_EXPORT_WITH_RC4_40_MD5
SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
SSL_DH_anon_WITH_DES_CBC_SHA
SSL_DH_anon_WITH_RC4_128_MD5
TLS_DH_anon_WITH_AES_128_CBC_SHA
TLS_DH_anon_WITH_AES_256_CBC_SHA

```

JSSE reference guide

For more information about SUN's JSSE framework, please consult the *JSSE Reference Guide* at the following location:

<http://java.sun.com/j2se/1.5.0/docs/guide/security/jsse/JSSERefGuide.html>

Cipher Suite Filters

Overview

In a typical application, you usually want to restrict the list of available cipher suites to a subset of the ciphers supported by the JSSE provider.

Namespaces

[Table 4.1 on page 66](#) shows the XML namespaces that are referenced in this section:

Table 4.1. Namespaces Used for Configuring Cipher Suite Filters

Prefix	Namespace URI
http	http://cxf.apache.org/transport/http/configuration
httpj	http://cxf.apache.org/transport/http-jetty/configuration
sec	http://cxf.apache.org/configuration/security

sec:cipherSuitesFilter element

You define a cipher suite filter using the `sec:cipherSuitesFilter` element, which can be a child of either a `http:tlsClientParameters` element or a `httpj:tlsServerParameters` element. A typical `sec:cipherSuitesFilter` element has the outline structure shown in [Example 4.1 on page 66](#).

Example 4.1. Structure of a sec:cipherSuitesFilter Element

```
<sec:cipherSuitesFilter>
  <sec:include>RegularExpression</sec:include>
  <sec:include>RegularExpression</sec:include>
  ...
  <sec:exclude>RegularExpression</sec:exclude>
  <sec:exclude>RegularExpression</sec:exclude>
  ...
</sec:cipherSuitesFilter>
```

Semantics

The following semantic rules apply to the `sec:cipherSuitesFilter` element:

1. If a `sec:cipherSuitesFilter` element does *not* appear in an endpoint's configuration (that is, it is absent from the relevant `http:conduit` or `httpj:engine-factory` element), the following default filter is used:

```
<sec:cipherSuitesFilter>
  <sec:include>.*_EXPORT_.*</sec:include>
  <sec:include>.*_EXPORT1024.*</sec:include>
  <sec:include>.*_DES_.*</sec:include>
  <sec:include>.*_WITH_NULL_.*</sec:include>
</sec:cipherSuitesFilter>
```

2. If the `sec:cipherSuitesFilter` element *does* appear in an endpoint's configuration, all cipher suites are *excluded* by default.
3. To include cipher suites, add a `sec:include` child element to the `sec:cipherSuitesFilter` element. The content of the `sec:include` element is a regular expression that matches one or more cipher suite names (for example, see the cipher suite names in ["Cipher suites supported by SunJSSE" on page ?](#)).
4. To refine the selected set of cipher suites further, you can add a `sec:exclude` element to the `sec:cipherSuitesFilter` element. The content of the `sec:exclude` element is a regular expression that matches zero or more cipher suite names from the currently included set.



Note

Sometimes it makes sense to explicitly exclude cipher suites that are currently not included, in order to future-proof against accidental inclusion of undesired cipher suites.

Regular expression matching

The grammar for the regular expressions that appear in the `sec:include` and `sec:exclude` elements is defined by the Java regular expression utility, `java.util.regex.Pattern`. For a detailed description of the grammar, please consult the Java reference guide, <http://java.sun.com/j2se/1.5.0/docs/api/java/util/regex/Pattern.html>.

Client conduit example

The following XML configuration shows an example of a client that applies a cipher suite filter to the remote endpoint, `{WSDLPortNamespace}PortName`. Whenever the client attempts to open an SSL/TLS connection to this endpoint, it restricts the available cipher suites to the set selected by the `sec:cipherSuitesFilter` element.

```
<beans ... >
  <http:conduit name="{WSDLPortNamespace}PortName.http-conduit">
    <http:tlsClientParameters>
      ...
    <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>

<bean id="cxf" class="org.apache.cxf.bus.CXFBusImpl"/>
</beans>
```

SSL/TLS Protocol Version

Overview

The versions of the SSL/TLS protocol that are supported by Fuse Services Framework depend on the particular *JSSE provider* configured. By default, the JSSE provider is configured to be SUN's JSSE provider implementation.

SSL/TLS protocol versions supported by SunJSSE

[Table 4.2 on page 69](#) shows the SSL/TLS protocol versions supported by SUN's JSSE provider.

Table 4.2. SSL/TLS Protocols Supported by SUN's JSSE Provider

Protocol	Description
SSL	Supports some version of SSL; may support other versions
SSLv2	Supports SSL version 2 or higher
SSLv3	Supports SSL version 3; may support other versions
TLS	Supports some version of TLS; may support other versions
TLSv1	Supports TLS version 1; may support other versions

Specifying the SSL/TLS protocol version

You can specify the preferred SSL/TLS protocol version as an attribute on the `http:tlsClientParameters` element (client side) or on the `httpj:tlsServerParameters` element (server side).

Client side SSL/TLS protocol version

You can specify the protocol to be TLS on the client side by setting the `secureSocketProtocol` attribute as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans ... >
  ...
  <http:conduit name="{Namespace}PortName.http-conduit">
    ...
    <http:tlsClientParameters secureSocketProtocol="TLS">
      ...
    </http:tlsClientParameters>
  </http:conduit>
</beans>
```

```
</http:conduit>
...
</beans>
```

Server side SSL/TLS protocol version

You can specify the protocol to be TLS on the server side by setting the `secureSocketProtocol` attribute as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans ... >
  ...
  <httpj:engine-factory bus="cxf">
    <httpj:engine port="9001">
      ...
      <httpj:tlsServerParameters secureSocketProtocol="TLS">
        ...
      </httpj:tlsClientParameters>
    </httpj:engine>
  </httpj:engine-factory>
  ...
</beans>
```

Chapter 5. The WS-Policy Framework

This chapter provides an introduction to the basic concepts of the WS-Policy framework, defining policy subjects and policy assertions, and explaining how policy assertions can be combined to make policy expressions.

Introduction to WS-Policy	72
Policy Expressions	76

Introduction to WS-Policy

Overview

The WS-Policy [specification](http://www.w3.org/TR/ws-policy/)¹ provides a general framework for applying policies that modify the semantics of connections and communications at runtime in a Web services application. Fuse Services Framework security uses the WS-Policy framework to configure message protection and authentication requirements.

Policies and policy references

The simplest way to specify a policy is to embed it directly where you want to apply it. For example, to associate a policy with a specific port in the WSDL contract, you can specify it as follows:

```
<wsdl:definitions targetNamespace="http://tempuri.org/"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-utility-1.0.xsd"
  xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy" ... >
  ...
  <wsdl:service name="PingService10">
    <wsdl:port name="UserNameOverTransport_IPingService" binding="BindingName">
      <wsp:Policy>
        <!-- Policy expression comes here! -->
      </wsp:Policy>
      <soap:address location="SOAPAddress"/>
    </wsdl:port>
  </wsdl:service>
</wsdl:definitions>
```

An alternative way to specify a policy is to insert a policy reference element, `wsp:PolicyReference`, at the point where you want to apply the policy and then insert the policy element, `wsp:Policy`, at some other point in the XML file. For example, to associate a policy with a specific port using a policy reference, you could use a configuration like the following:

```
<wsdl:definitions targetNamespace="http://tempuri.org/"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-utility-1.0.xsd"
  xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy" ... >
  ...
  <wsdl:service name="PingService10">
```

¹ <http://www.w3.org/TR/ws-policy/>


```

    <wsdl:port name="UserNameOverTransport_IPingService" binding="BindingName">
      <wsp:PolicyReference URI="#PolicyID"/>
      <soap:address location="SOAPAddress"/>
    </wsdl:port>
  </wsdl:service>
  ...
  <wsp:Policy wsu:Id="PolicyID">
    <!-- Policy expression comes here ... -->
  </wsp:Policy>
</wsdl:definitions>

```

Where the policy reference, `wsp:PolicyReference`, locates the referenced policy using the ID, *PolicyID* (note the addition of the # prefix character in the URI attribute). The policy itself, `wsp:Policy`, must be identified by adding the attribute, `wsu:Id="PolicyID"`.

Policy subjects

The entities with which policies are associated are called *policy subjects*. For example, you can associate a policy with an endpoint, in which case the *endpoint* is the policy subject. It is possible to associate multiple policies with any given policy subject. The WS-Policy framework supports the following kinds of policy subject:

- "Service policy subject" on page 73.
- "Endpoint policy subject" on page 73.
- "Operation policy subject" on page 74.
- "Message policy subject" on page 75.

Service policy subject

To associate a policy with a service, insert either a `<wsp:Policy>` element or a `<wsp:PolicyReference>` element as a sub-element of the following WSDL 1.1 element:

- `wsdl:service`—apply the policy to all of the ports (endpoints) offered by this service.

Endpoint policy subject

To associate a policy with an endpoint, insert either a `<wsp:Policy>` element or a `<wsp:PolicyReference>` element as a sub-element of any of the following WSDL 1.1 elements:

- `wsdl:portType`—apply the policy to all of the ports (endpoints) that use this port type.
- `wsdl:binding`—apply the policy to all of the ports that use this binding.

- `wsdl:port`—apply the policy to this endpoint only.

For example, you can associate a policy with an endpoint binding as follows (using a policy reference):

```
<wsdl:definitions targetNamespace="http://tempuri.org/"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
  xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-utility-1.0.xsd"
  xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy" ... >
  ...
  <wsdl:binding name="EndpointBinding" type="i0:IPingService">
    <wsp:PolicyReference URI="#PolicyID"/>
    ...
  </wsdl:binding>
  ...
  <wsp:Policy wsu:Id="PolicyID"> ... </wsp:Policy>
  ...
</wsdl:definitions>
```

Operation policy subject

To associate a policy with an operation, insert either a `<wsp:Policy>` element or a `<wsp:PolicyReference>` element as a sub-element of any of the following WSDL 1.1 elements:

- `wsdl:portType/wsdl:operation`
- `wsdl:binding/wsdl:operation`

For example, you can associate a policy with an operation in a binding as follows (using a policy reference):

```
<wsdl:definitions targetNamespace="http://tempuri.org/"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-utility-1.0.xsd"
  xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy" ... >
  ...
  <wsdl:binding name="EndpointBinding" type="i0:IPingService">
    <wsdl:operation name="Ping">
      <wsp:PolicyReference URI="#PolicyID"/>
      <soap:operation soapAction="http://xmlsoap.org/Ping" style="document"/>
      <wsdl:input name="PingRequest"> ... </wsdl:input>
      <wsdl:output name="PingResponse"> ... </wsdl:output>
    </wsdl:operation>
    ...
  </wsdl:binding>
  ...
```

```

    <wsp:Policy wsu:Id="PolicyID"> ... </wsp:Policy>
    ...
</wsdl:definitions>

```

Message policy subject

To associate a policy with a message, insert either a `<wsp:Policy>` element or a `<wsp:PolicyReference>` element as a sub-element of any of the following WSDL 1.1 elements:

- `wsdl:message`
- `wsdl:portType/wsdl:operation/wsdl:input`
- `wsdl:portType/wsdl:operation/wsdl:output`
- `wsdl:portType/wsdl:operation/wsdl:fault`
- `wsdl:binding/wsdl:operation/wsdl:input`
- `wsdl:binding/wsdl:operation/wsdl:output`
- `wsdl:binding/wsdl:operation/wsdl:fault`

For example, you can associate a policy with a message in a binding as follows (using a policy reference):

```

<wsdl:definitions targetNamespace="http://tempuri.org/"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-utility-
1.0.xsd"
  xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy" ... >
  ...
  <wsdl:binding name="EndpointBinding" type="i0:IPingService">
    <wsdl:operation name="Ping">
      <soap:operation soapAction="http://xmlsoap.org/Ping" style="document"/>
      <wsdl:input name="PingRequest">
        <wsp:PolicyReference URI="#PolicyID"/>
        <soap:body use="literal"/>
      </wsdl:input>
      <wsdl:output name="PingResponse"> ... </wsdl:output>
    </wsdl:operation>
    ...
  </wsdl:binding>
  ...
  <wsp:Policy wsu:Id="PolicyID"> ... </wsp:Policy>
  ...
</wsdl:definitions>

```

Policy Expressions

Overview

In general, a `wsp:Policy` element is composed of multiple different policy settings (where individual policy settings are specified as *policy assertions*). Hence, the policy defined by a `wsp:Policy` element is really a composite object. The content of the `wsp:Policy` element is called a *policy expression*, where the policy expression consists of various logical combinations of the basic policy assertions. By tailoring the syntax of the policy expression, you can determine what combinations of policy assertions must be satisfied at runtime in order to satisfy the policy overall.

This section describes the syntax and semantics of policy expressions in detail.

Policy assertions

Policy assertions are the basic building blocks that can be combined in various ways to produce a policy. A policy assertion has two key characteristics: it adds a basic unit of functionality to the policy subject *and* it represents a boolean assertion to be evaluated at runtime. For example, consider the following policy assertion that requires a WS-Security username token to be propagated with request messages:

```
<sp:SupportingTokens xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">
  <wsp:Policy>
    <sp:UsernameToken/>
  </wsp:Policy>
</sp:SupportingTokens>
```

When associated with an endpoint policy subject, this policy assertion has the following effects:

- The Web service endpoint marshals/unmarshals the UsernameToken credentials.
- At runtime, the policy assertion returns `true`, if UsernameToken credentials are provided (on the client side) or received in the incoming message (on the server side); otherwise the policy assertion returns `false`.

Note that if a policy assertion returns `false`, this does not necessarily result in an error. The net effect of a particular policy assertion depends on how it is inserted into a policy and on how it is combined with other policy assertions.

Policy alternatives

A policy is built up using policy assertions, which can additionally be qualified using the `wsp:Optional` attribute, and various nested combinations of the `wsp:All` and `wsp:ExactlyOne` elements. The net effect of composing these elements is to produce a range of acceptable *policy alternatives*. As long as one of these acceptable policy alternatives is satisfied, the overall policy is also satisfied (evaluates to `true`).

wsp:All element

When a list of policy assertions is wrapped by the `wsp:All` element, *all* of the policy assertions in the list must evaluate to true. For example, consider the following combination of authentication and authorization policy assertions:

```
<wsp:Policy wsu:Id="AuthenticateAndAuthorizeWSSUsernameTokenPolicy">
  <wsp:All>
    <sp:SupportingTokens>
      <wsp:Policy>
        <sp:UsernameToken/>
      </wsp:Policy>
    </sp:SupportingTokens>
    <sp:SupportingTokens>
      <wsp:Policy>
        <sp:SamlToken/>
      </wsp:Policy>
    </sp:SupportingTokens>
  </wsp:All>
</wsp:Policy>
```

The preceding policy will be satisfied for a particular incoming request, if the following conditions *both* hold:

- WS-Security UsernameToken credentials must be present; *and*
- A SAML token must be present.



Note

The `wsp:Policy` element is semantically equivalent to `wsp:All`. Hence, if you removed the `wsp:All` element from the preceding example, you would obtain a semantically equivalent example

wsp:ExactlyOne element

When a list of policy assertions is wrapped by the `wsp:ExactlyOne` element, *at least one* of the policy assertions in the list must evaluate to true. The runtime goes through the list, evaluating policy assertions until it finds a policy assertion that returns true. At that point, the `wsp:ExactlyOne` expression is satisfied (returns true) and any remaining policy assertions from the list will not be evaluated. For example, consider the following combination of authentication policy assertions:

```
<wsp:Policy wsu:Id="AuthenticateUsernamePasswordPolicy">
  <wsp:ExactlyOne>
    <sp:SupportingTokens>
```

```

    <wsp:Policy>
      <sp:UsernameToken/>
    </wsp:Policy>
  </sp:SupportingTokens>
<sp:SupportingTokens>
  <wsp:Policy>
    <sp:SamlToken/>
  </wsp:Policy>
</sp:SupportingTokens>
</wsp:ExactlyOne>
</wsp:Policy>

```

The preceding policy will be satisfied for a particular incoming request, if *either* of the following conditions hold:

- WS-Security UsernameToken credentials are present; *or*
- A SAML token is present.

Note, in particular, that if *both* credential types are present, the policy would be satisfied after evaluating one of the assertions, but no guarantees can be given as to which of the policy assertions actually gets evaluated.

The empty policy

A special case is the *empty policy*, an example of which is shown in [Example 5.1 on page 78](#).

Example 5.1. The Empty Policy

```

<wsp:Policy ... >
  <wsp:ExactlyOne>
    <wsp:All/>
  </wsp:ExactlyOne>
</wsp:Policy>

```

Where the empty policy alternative, `<wsp:All/>`, represents an alternative for which no policy assertions need be satisfied. In other words, it always returns true. When `<wsp:All/>` is available as an alternative, the overall policy can be satisfied even when no policy assertions are true.

The null policy

A special case is the *null policy*, an example of which is shown in [Example 5.2 on page 78](#).

Example 5.2. The Null Policy

```

<wsp:Policy ... >
  <wsp:ExactlyOne/>
</wsp:Policy>

```

Where the null policy alternative, `<wsp:ExactlyOne/>`, represents an alternative that is never satisfied. In other words, it always returns false.

Normal form

In practice, by nesting the `<wsp:All>` and `<wsp:ExactlyOne>` elements, you can produce fairly complex policy expressions, whose policy alternatives might be difficult to work out. To facilitate the comparison of policy expressions, the WS-Policy specification defines a canonical or *normal form* for policy expressions, such that you can read off the list of policy alternatives unambiguously. Every valid policy expression can be reduced to the normal form.

In general, a normal form policy expression conforms to the syntax shown in [Example 5.3 on page 79](#).

Example 5.3. Normal Form Syntax

```
<wsp:Policy ... >
  <wsp:ExactlyOne>
    <wsp:All> <Assertion .../> ... <Assertion .../> </wsp:All>
    <wsp:All> <Assertion .../> ... <Assertion .../> </wsp:All>
    ...
  </wsp:ExactlyOne>
</wsp:Policy>
```

Where each line of the form, `<wsp:All>...</wsp:All>`, represents a valid policy alternative. If one of these policy alternatives is satisfied, the policy is satisfied overall.

Chapter 6. Message Protection

The following message protection mechanisms are described in this chapter: protection against eavesdropping (by employing encryption algorithms) and protection against message tampering (by employing message digest algorithms). The protection can be applied at various levels of granularity and to different protocol layers. At the transport layer, you have the option of applying protection to the entire contents of the message; while at the SOAP layer, you have the option of applying protection to various parts of the message (bodies, headers, or attachments).

Transport Layer Message Protection	82
SOAP Message Protection	86
Introduction to SOAP Message Protection	87
Basic Signing and Encryption Scenario	89
Specifying an AsymmetricBinding Policy	91
Specifying a SymmetricBinding Policy	97
Specifying Parts of Message to Encrypt and Sign	101
Providing Encryption Keys and Signing Keys	104
Specifying the Algorithm Suite	111

Transport Layer Message Protection

Overview

Transport layer message protection refers to the message protection (encryption and signing) that is provided by the transport layer. For example, HTTPS provides encryption and message signing features using SSL/TLS. In fact, WS-SecurityPolicy does not add much to the HTTPS feature set, because HTTPS is already fully configurable using Spring XML configuration (see ["Configuring HTTPS" on page 49](#)). An advantage of specifying a transport binding policy for HTTPS, however, is that it enables you to embed security requirements in the WSDL contract. Hence, any client that obtains a copy of the WSDL contract can discover what the transport layer security requirements are for the endpoints in the WSDL contract.

Prerequisites

If you use WS-SecurityPolicy to configure the HTTPS transport, you must also configure HTTPS security appropriately in the Spring configuration.

[Example 6.1 on page 82](#) shows how to configure a client to use the HTTPS transport protocol. The `sec:keyManagers` element specifies the client's own certificate, `alice.pfx`, and the `sec:trustManagers` element specifies the trusted CA list. Note how the `http:conduit` element's name attribute uses wildcards to match the endpoint address. For details of how to configure HTTPS on the client side, see ["Configuring HTTPS" on page 49](#).

Example 6.1. Client HTTPS Configuration in Spring

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:http="http://cxf.apache.org/transports/http/configuration"
  xmlns:sec="http://cxf.apache.org/configuration/security" ... >

  <http:conduit name="https://.*/UserNameOverTransport.*">
    <http:tlsClientParameters disableCNCheck="true">
      <sec:keyManagers keyPassword="password">
        <sec:keyStore type="pkcs12" password="password" resource="certs/alice.pfx"/>
      </sec:keyManagers>
      <sec:trustManagers>
        <sec:keyStore type="pkcs12" password="password" resource="certs/bob.pfx"/>
      </sec:trustManagers>
    </http:tlsClientParameters>
  </http:conduit>
  ...
</beans>
```

[Example 6.2 on page 83](#) shows how to configure a server to use the HTTPS transport protocol. The `sec:keyManagers` element specifies the server's own certificate, `bob.pfx`, and the `sec:trustManagers`

element specifies the trusted CA list. For details of how to configure HTTPS on the server side, see ["Configuring HTTPS" on page 49](#).

Example 6.2. Server HTTPS Configuration in Spring

```
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:http="http://cxf.apache.org/transports/http/configuration"
       xmlns:sec="http://cxf.apache.org/configuration/security" ... >

  <httpj:engine-factory id="tls-settings">
    <httpj:engine port="9001">
      <httpj:tlsServerParameters>
        <sec:keyManagers keyPassword="password">
          <sec:keyStore type="pkcs12" password="password" resource="certs/bob.pfx"/>
        </sec:keyManagers>
        <sec:trustManagers>
          <sec:keyStore type="pkcs12" password="password" resource="certs/alice.pfx"/>
        </sec:trustManagers>
      </httpj:tlsServerParameters>
    </httpj:engine>
  </httpj:engine-factory>
  ...
</beans>
```

Policy subject

A transport binding policy must be applied to an endpoint policy subject (see ["Endpoint policy subject" on page 73](#)). For example, given the transport binding policy with ID, `UserNameOverTransport_IPingService_policy`, you could apply the policy to an endpoint binding as follows:

```
<wsdl:binding name="UserNameOverTransport_IPingService" type="i0:IPingService">
  <wsp:PolicyReference URI="#UserNameOverTransport_IPingService_policy"/>
  ...
</wsdl:binding>
```

Syntax

The `TransportBinding` element has the following syntax:

```
<sp:TransportBinding xmlns:sp="..." ... >
  <wsp:Policy xmlns:wsp="...">
    <sp:TransportToken ... >
      <wsp:Policy> ... </wsp:Policy>
    ...
  </sp:TransportToken>
  <sp:AlgorithmSuite ... > ... </sp:AlgorithmSuite>
```

```

    <sp:Layout ... > ... </sp:Layout> ?
    <sp:IncludeTimestamp ... /> ?
    ...
  </wsp:Policy>
  ...
</sp:TransportBinding>

```

Sample policy

[Example 6.3 on page 84](#) shows an example of a transport binding that requires confidentiality and integrity using the HTTPS transport (specified by the `sp:HttpsToken` element) and a 256-bit algorithm suite (specified by the `sp:Basic256` element).

Example 6.3. Example of a Transport Binding

```

<wsp:Policy wsu:Id="UserNameOverTransport_IPingService_policy">
  <wsp:ExactlyOne>
    <wsp:All>
      <sp:TransportBinding xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">

        <wsp:Policy>
          <sp:TransportToken>
            <wsp:Policy>
              <sp:HttpsToken RequireClientCertificate="false"/>
            </wsp:Policy>
          </sp:TransportToken>
          <sp:AlgorithmSuite>
            <wsp:Policy>
              <sp:Basic256/>
            </wsp:Policy>
          </sp:AlgorithmSuite>
          <sp:Layout>
            <wsp:Policy>
              <sp:Lax/>
            </wsp:Policy>
          </sp:Layout>
          <sp:IncludeTimestamp/>
        </wsp:Policy>
      </sp:TransportBinding>
      ...
    <sp:Wss10 xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">
      <wsp:Policy>
        <sp:MustSupportRefKeyIdentifier/>
        <sp:MustSupportRefIssuerSerial/>
      </wsp:Policy>
    </sp:Wss10>
  </wsp:All>

```

```
</wsp:ExactlyOne>
</wsp:Policy>
```

sp:TransportToken

This element has a two-fold effect: it requires a particular type of security token and it indicates how the transport is secured. For example, by specifying the `sp:HttpsToken`, you indicate that the connection is secured by the HTTPS protocol and the security tokens are X.509 certificates.

sp:AlgorithmSuite

This element specifies the suite of cryptographic algorithms to use for signing and encryption. For details of the available algorithm suites, see ["Specifying the Algorithm Suite" on page 111](#).

sp:Layout

This element specifies whether to impose any conditions on the order in which security headers are added to the SOAP message. The `sp:Lax` element specifies that no conditions are imposed on the order of security headers. The alternatives to `sp:Lax` are `sp:Strict`, `sp:LaxTimestampFirst`, or `sp:LaxTimestampLast`.

sp:IncludeTimestamp

If this element is included in the policy, the runtime adds a `wsu:Timestamp` element to the `wsse:Security` header. By default, the timestamp is *not* included.

sp:MustSupportRefKeyIdentifier

This element specifies that the security runtime must be able to process *Key Identifier* token references, as specified in the WS-Security 1.0 specification. A key identifier is a mechanism for identifying a key token, which may be used inside signature or encryption elements. Fuse Services Framework requires this feature.

sp:MustSupportRefIssuerSerial

This element specifies that the security runtime must be able to process *Issuer and Serial Number* token references, as specified in the WS-Security 1.0 specification. An issuer and serial number is a mechanism for identifying a key token, which may be used inside signature or encryption elements. Fuse Services Framework requires this feature.

SOAP Message Protection

Introduction to SOAP Message Protection	87
Basic Signing and Encryption Scenario	89
Specifying an AsymmetricBinding Policy	91
Specifying a SymmetricBinding Policy	97
Specifying Parts of Message to Encrypt and Sign	101
Providing Encryption Keys and Signing Keys	104
Specifying the Algorithm Suite	111

Introduction to SOAP Message Protection

Overview

By applying message protection at the SOAP encoding layer, instead of at the transport layer, you have access to a more flexible range of protection policies. In particular, because the SOAP layer is aware of the message structure, you can apply protection at a finer level of granularity—for example, by encrypting and signing only those headers that actually require protection. This feature enables you to support more sophisticated multi-tier architectures. For example, one plaintext header might be aimed at an intermediate tier (located within a secure intranet), while an encrypted header might be aimed at the final destination (reached through an insecure public network).

Security bindings

As described in the WS-SecurityPolicy specification, one of the following binding types can be used to protect SOAP messages:

- `sp:TransportBinding`—the *transport binding* refers to message protection provided at the transport level (for example, through HTTPS). This binding can be used to secure any message type, not just SOAP, and it is described in detail in the preceding section, ["Transport Layer Message Protection" on page 82](#).
- `sp:AsymmetricBinding`—the *asymmetric binding* refers to message protection provided at the SOAP message encoding layer, where the protection features are implemented using asymmetric cryptography (also known as public key cryptography).
- `sp:SymmetricBinding`—the *symmetric binding* refers to message protection provided at the SOAP message encoding layer, where the protection features are implemented using symmetric cryptography. Examples of symmetric cryptography are the tokens provided by WS-SecureConversation and Kerberos tokens.

Message protection

The following qualities of protection can be applied to part or all of a message:

- Encryption.
- Signing.
- Signing+encryption (sign before encrypting).
- Encryption+signing (encrypt before signing).

These qualities of protection can be arbitrarily combined in a single message. Thus, some parts of a message can be just encrypted, while other parts of the message are just signed, and other parts of the message can be both signed and encrypted. It is also possible to leave parts of the message unprotected.

The most flexible options for applying message protection are available at the SOAP layer (`sp:AsymmetricBinding` or `sp:SymmetricBinding`). The transport layer (`sp:TransportBinding`) only gives you the option of applying protection to the *whole* message.

Specifying parts of the message to protect

Currently, Fuse Services Framework enables you to sign or encrypt the following parts of a SOAP message:

- *Body*—sign and/or encrypt the whole of the `soap:BODY` element in a SOAP message.
- *Header(s)*—sign and/or encrypt one or more SOAP message headers. You can specify the quality of protection for each header individually.
- *Attachments*—sign and/or encrypt all of the attachments in a SOAP message.

The WS-SecurityPolicy specification also defines policies for applying protection to individual XML elements, but this is currently *not* supported in Fuse Services Framework.

Role of configuration

Not all of the details required for message protection are specified using policies. The policies are primarily intended to provide a way of specifying the quality of protection required for a service. Supporting details, such as security tokens, passwords, and so on, must be provided using a separate, product-specific mechanism. In practice, this means that in Fuse Services Framework, some supporting configuration details must be provided in Spring XML configuration files. For details, see ["Providing Encryption Keys and Signing Keys" on page 104](#).

Basic Signing and Encryption Scenario

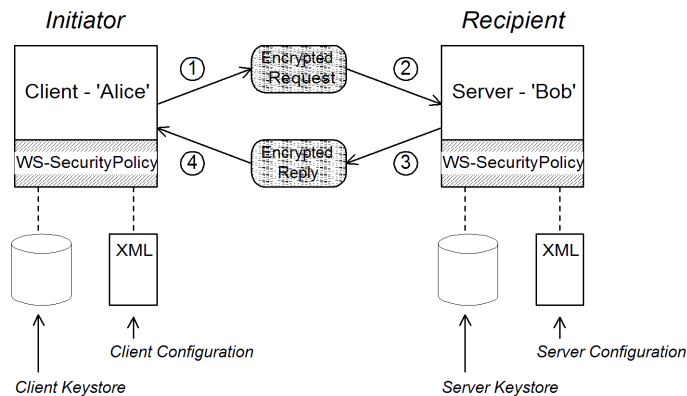
Overview

The scenario described here is a client-server application, where an *asymmetric binding policy* is set up to encrypt and sign the SOAP body of messages that pass back and forth between the client and the server.

Example scenario

Figure 6.1 on page 89 shows an overview of the basic signing and encryption scenario, which is specified by associating an asymmetric binding policy with an endpoint in the WSDL contract.

Figure 6.1. Basic Signing and Encryption Scenario



Scenario steps

When the client in Figure 6.1 on page 89 invokes a synchronous operation on the recipient's endpoint, the request and reply message are processed as follows:

1. As the outgoing request message passes through the WS-SecurityPolicy handler, the handler processes the message in accordance with the policies specified in the client's asymmetric binding policy. In this example, the handler performs the following processing:
 - a. Encrypt the SOAP body of the message using Bob's public key.
 - b. Sign the encrypted SOAP body using Alice's private key.

2. As the incoming request message passes through the server's WS-SecurityPolicy handler, the handler processes the message in accordance with the policies specified in the server's asymmetric binding policy. In this example, the handler performs the following processing:
 - a. Verify the signature using Alice's public key.
 - b. Decrypt the SOAP body using Bob's private key.
3. As the outgoing reply message passes back through the server's WS-SecurityPolicy handler, the handler performs the following processing:
 - a. Encrypt the SOAP body of the message using Alice's public key.
 - b. Sign the encrypted SOAP body using Bob's private key.
4. As the incoming reply message passes back through the client's WS-SecurityPolicy handler, the handler performs the following processing:
 - a. Verify the signature using Bob's public key.
 - b. Decrypt the SOAP body using Alice's private key.

Specifying an AsymmetricBinding Policy

Overview

The asymmetric binding policy implements SOAP message protection using asymmetric key algorithms (public/private key combinations) and does so at the SOAP layer. The encryption and signing algorithms used by the asymmetric binding are similar to the encryption and signing algorithms used by SSL/TLS. A crucial difference, however, is that SOAP message protection enables you to select particular parts of a message to protect (for example, individual headers, body, or attachments), whereas transport layer security can operate only on the *whole* message.

Policy subject

An asymmetric binding policy must be applied to an endpoint policy subject (see ["Endpoint policy subject" on page 73](#)). For example, given the asymmetric binding policy with ID, `MutualCertificate10SignEncrypt_IPingService_policy`, you could apply the policy to an endpoint binding as follows:

```
<wsdl:binding name="MutualCertificate10SignEncrypt_IPingService" type="i0:IPingService">
  <wsp:PolicyReference URI="#MutualCertificate10SignEncrypt_IPingService_policy"/>
  ...
</wsdl:binding>
```

Syntax

The `AsymmetricBinding` element has the following syntax:

```
<sp:AsymmetricBinding xmlns:sp="..." ... >
  <wsp:Policy xmlns:wsp="...">
    (
      <sp:InitiatorToken>
        <wsp:Policy> ... </wsp:Policy>
      </sp:InitiatorToken>
    ) | (
      <sp:InitiatorSignatureToken>
        <wsp:Policy> ... </wsp:Policy>
      </sp:InitiatorSignatureToken>
      <sp:InitiatorEncryptionToken>
        <wsp:Policy> ... </wsp:Policy>
      </sp:InitiatorEncryptionToken>
    )
  (
    <sp:RecipientToken>
      <wsp:Policy> ... </wsp:Policy>
    </sp:RecipientToken>
```

```

) | (
  <sp:RecipientSignatureToken>
    <wsp:Policy> ... </wsp:Policy>
  </sp:RecipientSignatureToken>
  <sp:RecipientEncryptionToken>
    <wsp:Policy> ... </wsp:Policy>
  </sp:RecipientEncryptionToken>
)
  <sp:AlgorithmSuite ... > ... </sp:AlgorithmSuite>
  <sp:Layout ... > ... </sp:Layout> ?
  <sp:IncludeTimestamp ... /> ?
  <sp:EncryptBeforeSigning ... /> ?
  <sp:EncryptSignature ... /> ?
  <sp:ProtectTokens ... /> ?
  <sp:OnlySignEntireHeadersAndBody ... /> ?
  ...
</wsp:Policy>
...
</sp:AsymmetricBinding>

```

Sample policy

[Example 6.4 on page 92](#) shows an example of an asymmetric binding that supports message protection with signatures and encryption, where the signing and encryption is done using pairs of public/private keys (that is, using asymmetric cryptography). This example does not specify *which* parts of the message should be signed and encrypted, however. For details of how to do that, see ["Specifying Parts of Message to Encrypt and Sign" on page 101](#).

Example 6.4. Example of an Asymmetric Binding

```

<wsp:Policy wsu:Id="MutualCertificate10SignEncrypt_IPingService_policy">
  <wsp:ExactlyOne>
    <wsp>All>
      <sp:AsymmetricBinding
        xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">
        <wsp:Policy>
          <sp:InitiatorToken>
            <wsp:Policy>
              <sp:X509Token
                sp:IncludeToken="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy/In
cludeToken/AlwaysToRecipient">
                <wsp:Policy>
                  <sp:WssX509V3Token10/>
                </wsp:Policy>
              </sp:X509Token>
            </wsp:Policy>
          </sp:InitiatorToken>

```

```

        <sp:RecipientToken>
          <wsp:Policy>
            <sp:X509Token
              sp:IncludeToken="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy/In
cludeToken/Never">
              <wsp:Policy>
                <sp:WssX509V3Token10/>
              </wsp:Policy>
            </sp:X509Token>
          </wsp:Policy>
        </sp:RecipientToken>
        <sp:AlgorithmSuite>
          <wsp:Policy>
            <sp:Basic256/>
          </wsp:Policy>
        </sp:AlgorithmSuite>
        <sp:Layout>
          <wsp:Policy>
            <sp:Lax/>
          </wsp:Policy>
        </sp:Layout>
        <sp:IncludeTimestamp/>
        <sp:EncryptSignature/>
        <sp:OnlySignEntireHeadersAndBody/>
      </wsp:Policy>
    </sp:AsymmetricBinding>
    <sp:Wss10 xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">
      <wsp:Policy>
        <sp:MustSupportRefKeyIdentifier/>
        <sp:MustSupportRefIssuerSerial/>
      </wsp:Policy>
    </sp:Wss10>
  </wsp:All>
</wsp:ExactlyOne>
</wsp:Policy>

```

sp:InitiatorToken

The *initiator token* refers to the public/private key-pair owned by the initiator. This token is used as follows:

- The token's private key signs messages sent from initiator to recipient.
- The token's public key verifies signatures received by the recipient.
- The token's public key encrypts messages sent from recipient to initiator.
- The token's private key decrypts messages received by the initiator.

Confusingly, this token is used both by the initiator *and* by the recipient. However, only the initiator has access to the private key so, in this sense, the token can be said to belong to the initiator. In ["Basic Signing and Encryption Scenario" on page 89](#), the initiator token is the certificate, Alice.

This element should contain a nested `wsp:Policy` element and `sp:X509Token` element as shown. The `sp:IncludeToken` attribute is set to `AlwaysToRecipient`, which instructs the runtime to include Alice's public key with every message sent to the recipient. This option is useful, in case the recipient wants to use the initiator's certificate to perform authentication. The most deeply nested element, `wssX509V3Token10` is optional. It specifies what specification version the X.509 certificate should conform to. The following alternatives (or none) can be specified here:

`sp:WssX509V3Token10`

This optional element is a policy assertion that indicates that an X509 Version 3 token should be used.

`sp:WssX509Pkcs7Token10`

This optional element is a policy assertion that indicates that an X509 PKCS7 token should be used.

`sp:WssX509PkiPathV1Token10`

This optional element is a policy assertion that indicates that an X509 PKI Path Version 1 token should be used.

`sp:WssX509V1Token11`

This optional element is a policy assertion that indicates that an X509 Version 1 token should be used.

`sp:WssX509V3Token11`

This optional element is a policy assertion that indicates that an X509 Version 3 token should be used.

`sp:WssX509Pkcs7Token11`

This optional element is a policy assertion that indicates that an X509 PKCS7 token should be used.

`sp:WssX509PkiPathV1Token11`

This optional element is a policy assertion that indicates that an X509 PKI Path Version 1 token should be used.

sp:RecipientToken

The *recipient token* refers to the public/private key-pair owned by the recipient. This token is used as follows:

- The token's public key encrypts messages sent from initiator to recipient.
- The token's private key decrypts messages received by the recipient.
- The token's private key signs messages sent from recipient to initiator.
- The token's public key verifies signatures received by the initiator.

Confusingly, this token is used both by the recipient *and* by the initiator. However, only the recipient has access to the private key so, in this sense, the token can be said to belong to the recipient. In ["Basic Signing and Encryption Scenario" on page 89](#), the recipient token is the certificate, Bob.

This element should contain a nested `wsp:Policy` element and `sp:X509Token` element as shown. The `sp:IncludeToken` attribute is set to `Never`, because there is no need to include Bob's public key in the reply messages.



Note

In Fuse Services Framework, there is never a need to send Bob's or Alice's token in a message, because both Bob's certificate and Alice's certificate are provided at both ends of the connection—see ["Providing Encryption Keys and Signing Keys" on page 104](#).

sp:AlgorithmSuite

This element specifies the suite of cryptographic algorithms to use for signing and encryption. For details of the available algorithm suites, see ["Specifying the Algorithm Suite" on page 111](#).

sp:Layout

This element specifies whether to impose any conditions on the order in which security headers are added to the SOAP message. The `sp:Lax` element specifies that no conditions are imposed on the order of security headers. The alternatives to `sp:Lax` are `sp:Strict`, `sp:LaxTimestampFirst`, or `sp:LaxTimestampLast`.

sp:IncludeTimestamp

If this element is included in the policy, the runtime adds a `wsu:Timestamp` element to the `wsse:Security` header. By default, the timestamp is *not* included.

sp:EncryptBeforeSigning

If a message part is subject to both encryption and signing, it is necessary to specify the order in which these operations are performed. The default order is to sign before encrypting. But if you include this element in your asymmetric policy, the order is changed to encrypt before signing.



Note

Implicitly, this element also affects the order of the decryption and signature verification operations. For example, if the sender of a message signs before encrypting, the receiver of the message must decrypt before verifying the signature.

sp:EncryptSignature

This element specifies that the message signature must be encrypted (by the encryption token, specified as described in ["Providing Encryption Keys and Signing Keys" on page 104](#)). Default is false.



Note

The *message signature* is the signature obtained directly by signing various parts of the message, such as message body, message headers, or individual elements (see ["Specifying Parts of Message to Encrypt and Sign" on page 101](#)). Sometimes the message signature is referred to as the *primary signature*, because the WS-SecurityPolicy specification also supports the concept of an endorsing supporting token, which is used to sign the primary signature. Hence, if an `sp:EndorsingSupportingTokens` element is applied to an endpoint, you can have a chain of signatures: the primary signature, which signs the message itself, and the secondary signature, which signs the primary signature.

For more details about the various kinds of endorsing supporting token, see ["SupportingTokens assertions" on page 123](#).

sp:ProtectTokens

This element specifies that signatures must cover the token used to generate that signature. Default is false.

sp:OnlySignEntireHeadersAndBody

This element specifies that signatures can be applied *only* to an entire body or to entire headers, not to sub-elements of the body or sub-elements of a header. When this option is enabled, you are effectively prevented from using the `sp:SignedElements` assertion (see ["Specifying Parts of Message to Encrypt and Sign" on page 101](#)).

Specifying a SymmetricBinding Policy

Overview

The symmetric binding policy implements SOAP message protection using symmetric key algorithms (shared secret key) and does so at the SOAP layer. Examples of a symmetric binding are the Kerberos protocol and the WS-SecureConversation protocol.



Note

Currently, Fuse Services Framework supports *only* WS-SecureConversation tokens in a symmetric binding.

Policy subject

A symmetric binding policy must be applied to an endpoint policy subject (see ["Endpoint policy subject" on page 73](#)). For example, given the symmetric binding policy with ID, `SecureConversation_MutualCertificate10SignEncrypt_IPingService_policy`, you could apply the policy to an endpoint binding as follows:

```
<wsdl:binding name="SecureConversation_MutualCertificate10SignEncrypt_IPingService"
type="i0:IPingService">
  <wsp:PolicyReference URI="#SecureConversation_MutualCertificate10SignEncrypt_IPingSer
vice_policy"/>
  ...
</wsdl:binding>
```

Syntax

The `SymmetricBinding` element has the following syntax:

```
<sp:SymmetricBinding xmlns:sp="..." ... >
  <wsp:Policy xmlns:wsp="...">
    (
      <sp:EncryptionToken ... >
        <wsp:Policy> ... </wsp:Policy>
      </sp:EncryptionToken>
      <sp:SignatureToken ... >
        <wsp:Policy> ... </wsp:Policy>
      </sp:SignatureToken>
    ) | (
      <sp:ProtectionToken ... >
        <wsp:Policy> ... </wsp:Policy>
```

```

    </sp:ProtectionToken>
  )
  <sp:AlgorithmSuite ... > ... </sp:AlgorithmSuite>
  <sp:Layout ... > ... </sp:Layout> ?
  <sp:IncludeTimestamp ... /> ?
  <sp:EncryptBeforeSigning ... /> ?
  <sp:EncryptSignature ... /> ?
  <sp:ProtectTokens ... /> ?
  <sp:OnlySignEntireHeadersAndBody ... /> ?
  ...
</wsp:Policy>
...
</sp:SymmetricBinding>

```

Sample policy

[Example 6.5 on page 98](#) shows an example of a symmetric binding that supports message protection with signatures and encryption, where the signing and encryption is done using a single symmetric key (that is, using symmetric cryptography). This example does not specify *which* parts of the message should be signed and encrypted, however. For details of how to do that, see ["Specifying Parts of Message to Encrypt and Sign" on page 101](#).

Example 6.5. Example of a Symmetric Binding

```

<wsp:Policy wsu:Id="SecureConversation_MutualCertificate10SignEncrypt_IPingService_policy">
  <wsp:ExactlyOne>
    <wsp:All>
      <sp:SymmetricBinding xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">
        <wsp:Policy>
          <sp:ProtectionToken>
            <wsp:Policy>
              <sp:SecureConversationToken>
                ...
              </sp:SecureConversationToken>
            </wsp:Policy>
          </sp:ProtectionToken>
          <sp:AlgorithmSuite>
            <wsp:Policy>
              <sp:Basic256/>
            </wsp:Policy>
          </sp:AlgorithmSuite>
          <sp:Layout>
            <wsp:Policy>
              <sp:Lax/>
            </wsp:Policy>
          </sp:Layout>
        </wsp:Policy>
      </sp:SymmetricBinding>
    </wsp:All>
  </wsp:ExactlyOne>
</wsp:Policy>

```

```

        </sp:Layout>
        <sp:IncludeTimestamp/>
        <sp:EncryptSignature/>
        <sp:OnlySignEntireHeadersAndBody/>
    </wsp:Policy>
</sp:SymmetricBinding>
<sp:Wss10 xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">
    <wsp:Policy>
        <sp:MustSupportRefKeyIdentifier/>
        <sp:MustSupportRefIssuerSerial/>
    </wsp:Policy>
</sp:Wss10>
    ...
</wsp:All>
</wsp:ExactlyOne>
</wsp:Policy>

```

sp:ProtectionToken

This element specifies a symmetric token to use for both signing and encrypting messages. For example, you could specify a WS-SecureConversation token here.

If you want to use distinct tokens for signing and encrypting operations, use the `sp:SignatureToken` element and the `sp:EncryptionToken` element in place of this element.

sp:SignatureToken

This element specifies a symmetric token to use for signing messages. It should be used in combination with the `sp:EncryptionToken` element.

sp:EncryptionToken

This element specifies a symmetric token to use for encrypting messages. It should be used in combination with the `sp:SignatureToken` element.

sp:AlgorithmSuite

This element specifies the suite of cryptographic algorithms to use for signing and encryption. For details of the available algorithm suites, see ["Specifying the Algorithm Suite" on page 111](#).

sp:Layout

This element specifies whether to impose any conditions on the order in which security headers are added to the SOAP message. The `sp:Lax` element specifies that no conditions are imposed on the order of security headers. The alternatives to `sp:Lax` are `sp:Strict`, `sp:LaxTimestampFirst`, or `sp:LaxTimestampLast`.

sp:IncludeTimestamp

If this element is included in the policy, the runtime adds a `wsu:Timestamp` element to the `wsse:Security` header. By default, the timestamp is *not* included.

sp:EncryptBeforeSigning

When a message part is subject to both encryption and signing, it is necessary to specify the order in which these operations are performed. The default order is to sign before encrypting. But if you include this element in your symmetric policy, the order is changed to encrypt before signing.



Note

Implicitly, this element also affects the order of the decryption and signature verification operations. For example, if the sender of a message signs before encrypting, the receiver of the message must decrypt before verifying the signature.

sp:EncryptSignature

This element specifies that the message signature must be encrypted. Default is false.

sp:ProtectTokens

This element specifies that signatures must cover the token used to generate that signature. Default is false.

sp:OnlySignEntireHeadersAndBody

This element specifies that signatures can be applied *only* to an entire body or to entire headers, not to sub-elements of the body or sub-elements of a header. When this option is enabled, you are effectively prevented from using the `sp:SignedElements` assertion (see ["Specifying Parts of Message to Encrypt and Sign" on page 101](#)).

Specifying Parts of Message to Encrypt and Sign

Overview

Encryption and signing provide two kinds of protection: confidentiality and integrity, respectively. The WS-SecurityPolicy protection assertions are used to specify *which* parts of a message are subject to protection. Details of the protection mechanisms, on the other hand, are specified separately in the relevant binding policy (see x ["Specifying an AsymmetricBinding Policy" on page 91](#), ["Specifying a SymmetricBinding Policy" on page 97](#), and ["Transport Layer Message Protection" on page 82](#)).

The protection assertions described here are really intended to be used in combination with SOAP security, because they apply to features of a SOAP message. Nonetheless, these policies can also be satisfied by a transport binding (such as HTTPS), which applies protection to the *entire* message, rather than to specific parts.

Policy subject

A protection assertion must be applied to a *message policy subject* (see ["Message policy subject" on page 75](#)). In other words, it must be placed inside a `wsdl:input`, `wsdl:output`, or `wsdl:fault` element in a WSDL binding. For example, given the protection policy with ID, `MutualCertificate10SignEncrypt_IPingService_header_Input_policy`, you could apply the policy to a `wsdl:input` message part as follows:

```
<wsdl:operation name="header">
  <soap:operation soapAction="http://InteropBaseAddress/interop/header" style="document"/>

  <wsdl:input name="headerRequest">
    <wsp:PolicyReference
      URI="#MutualCertificate10SignEncrypt_IPingService_header_Input_policy"/>
    <soap:header message="i0:headerRequest_Headers" part="CustomHeader" use="literal"/>
    <soap:body use="literal"/>
  </wsdl:input>
  ...
</wsdl:operation>
```

Protection assertions

The following WS-SecurityPolicy protection assertions are currently supported by Fuse Services Framework:

- SignedParts
- EncryptedParts

The following WS-SecurityPolicy protection assertions are *not* supported by Fuse Services Framework:

- SignedElements
- EncryptedElements
- ContentEncryptedElements
- RequiredElements
- RequiredParts

Syntax

The SignedParts element has the following syntax:

```
<sp:SignedParts xmlns:sp="..." ... >
  <sp:Body />?
  <sp:Header Name="xs:NCName"? Namespace="xs:anyURI" ... />*
  <sp:Attachments />?
  ...
</sp:SignedParts>
```

The EncryptedParts element has the following syntax:

```
<sp:EncryptedParts xmlns:sp="..." ... >
  <sp:Body/>?
  <sp:Header Name="xs:NCName"? Namespace="xs:anyURI" ... />*
  <sp:Attachments />?
  ...
</sp:EncryptedParts>
```

Sample policy

[Example 6.6 on page 102](#) shows a policy that combines two protection assertions: a signed parts assertion and an encrypted parts assertion. When this policy is applied to a message part, the affected message bodies are signed and encrypted. In addition, the message header named CustomHeader is signed.

Example 6.6. Integrity and Encryption Policy Assertions

```
<wsp:Policy wsu:Id="MutualCertificate10SignEncrypt_IPingService_header_input_policy">
  <wsp:ExactlyOne>
    <wsp:All>
      <sp:SignedParts xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">
        <sp:Body/>
        <sp:Header Name="CustomHeader" Namespace="http://InteropBaseAddress/interop"/>
      </sp:SignedParts>
    </wsp:All>
  </wsp:ExactlyOne>
</wsp:Policy>
```

```

        </sp:SignedParts>
        <sp:EncryptedParts xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">
            <sp:Body/>
        </sp:EncryptedParts>
    </wsp:All>
    </wsp:ExactlyOne>
</wsp:Policy>

```

sp:Body

This element specifies that protection (encryption or signing) is applied to the body of the message. The protection is applied to the *entire* message body: that is, the `soap:Body` element, its attributes, and its content.

sp:Header

This element specifies that protection is applied to the SOAP header specified by the header's local name, using the `Name` attribute, and namespace, using the `Namespace` attribute. The protection is applied to the *entire* message header, including its attributes and its content.

sp:Attachments

This element specifies that *all* SOAP with Attachments (SwA) attachments are protected.

Providing Encryption Keys and Signing Keys

Overview

The standard WS-SecurityPolicy policies are designed to specify security *requirements* in some detail: for example, security protocols, security algorithms, token types, authentication requirements, and so on, are all described. But the standard policy assertions do not provide any mechanism for specifying associated security data, such as keys and credentials. WS-SecurityPolicy expects that the requisite security data is provided through a proprietary mechanism. In Fuse Services Framework, the associated security data is provided through Spring XML configuration.

Configuring encryption keys and signing keys

You can specify an application's encryption keys and signing keys by setting properties on a client's request context or on an endpoint context (see ["Add encryption and signing properties to Spring configuration" on page 105](#)). The properties you can set are shown in [Table 6.1 on page 104](#).

Table 6.1. Encryption and Signing Properties

Property	Description
<code>ws-security.signature.properties</code>	The WSS4J properties file/object that contains the WSS4J properties for configuring the signature keystore (which is also used for decrypting) and Crypto objects.
<code>ws-security.signature.username</code>	(Optional) The username or alias of the key in the signature keystore to use. If not specified, the alias set in the properties file is used. If that is also not set, and the keystore only contains a single key, that key will be used.
<code>ws-security.encryption.properties</code>	The WSS4J properties file/object that contains the WSS4J properties for configuring the encryption keystore (which is also used for validating signatures) and Crypto objects.
<code>ws-security.encryption.username</code>	(Optional) The username or alias of the key in the encryption keystore to use. If not specified, the alias set in the properties file is used. If that is also not set, and the keystore only contains a single key, that key will be used.

**Tip**

The names of the preceding properties are not so well chosen, because they do not accurately reflect what they are used for. The key specified by `ws-security.signature.properties` is actually used both for signing *and* decrypting. The key specified by `ws-security.encryption.properties` is actually used both for encrypting *and* for validating signatures.

Add encryption and signing properties to Spring configuration

Before you can use any WS-Policy policies in a Fuse Services Framework application, you must add the policies feature to the default CXF bus. Add the `p:policies` element to the CXF bus, as shown in the following Spring configuration fragment:

```
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:jaxws="http://cxf.apache.org/jaxws"
       xmlns:cxf="http://cxf.apache.org/core"
       xmlns:p="http://cxf.apache.org/policy" ... >

  <cxf:bus>
    <cxf:features>
      <p:policies/>
    </cxf:features>
  </cxf:bus>
  ...
</beans>
```

The following example shows how to add signature and encryption properties to proxies of the specified service type (where the service name is specified by the `name` attribute of the `jaxws:client` element). The properties are stored in WSS4J property files, where `alice.properties` contains the properties for the signature key and `bob.properties` contains the properties for the encryption key.

```
<beans ... >
  <jaxws:client name="{http://InteropBaseAddress/interop}MutualCertificate10SignEncrypt_IP
ingService"
               createdFromAPI="true">
    <jaxws:properties>
      <entry key="ws-security.signature.properties" value="etc/alice.properties"/>
      <entry key="ws-security.encryption.properties" value="etc/bob.properties"/>
    </jaxws:properties>
  </jaxws:client>
  ...
</beans>
```

In fact, although it is not obvious from the property names, each of these keys is used for two distinct purposes on the client side:

- `alice.properties` (that is, the key specified by `ws-security.signature.properties`) is used on the client side as follows:
 - For signing outgoing messages.
 - For decrypting incoming messages.
- `bob.properties` (that is, the key specified by `ws-security.encryption.properties`) is used on the client side as follows:
 - For encrypting outgoing messages.
 - For verifying signatures on incoming messages.

If you find this confusing, see ["Basic Signing and Encryption Scenario" on page 89](#) for a more detailed explanation.

The following example shows how to add signature and encryption properties to a JAX-WS endpoint. The properties file, `bob.properties`, contains the properties for the signature key and the properties file, `alice.properties`, contains the properties for the encryption key (this is the inverse of the client configuration).

```
<beans ... >
  <jaxws:endpoint
    name="{http://InteropBaseAddress/interop}MutualCertificate10SignEncrypt_IPingService"

    id="MutualCertificate10SignEncrypt"
    address="http://localhost:9002/MutualCertificate10SignEncrypt"
    serviceName="interop:PingService10"
    endpointName="interop:MutualCertificate10SignEncrypt_IPingService"
    implementor="interop.server.MutualCertificate10SignEncrypt">

    <jaxws:properties>
      <entry key="ws-security.signature.properties" value="etc/bob.properties"/>
      <entry key="ws-security.encryption.properties" value="etc/alice.properties"/>
    </jaxws:properties>

  </jaxws:endpoint>
  ...
</beans>
```

Each of these keys is used for two distinct purposes on the server side:

- `bob.properties` (that is, the key specified by `ws-security.signature.properties`) is used on the server side as follows:

- For signing outgoing messages.
- For decrypting incoming messages.
- `alice.properties` (that is, the key specified by `ws-security.encryption.properties`) is used on the server side as follows:
 - For encrypting outgoing messages.
 - For verifying signatures on incoming messages.

Define the WSS4J property files

Fuse Services Framework uses WSS4J property files to load the public keys and the private keys needed for encryption and signing. [Table 6.2 on page 107](#) describes the properties that you can set in these files.

Table 6.2. WSS4J Keystore Properties

Property	Description
<code>org.apache.ws.security.crypto.provider</code>	Specifies an implementation of the Crypto interface (see "WSS4J Crypto interface" on page 109). Normally, you specify the default WSS4J implementation of Crypto, <code>org.apache.ws.security.components.crypto.Merlin</code> . <i>The rest of the properties in this table are specific to the Merlin implementation of the Crypto interface.</i>
<code>org.apache.ws.security.crypto.merlin.keystore.provider</code>	(Optional) The name of the JSSE keystore provider to use. The default keystore provider is Bouncy Castle ¹ . You can switch provider to Sun's JSSE keystore provider by setting this property to <code>SunJSSE</code> .
<code>org.apache.ws.security.crypto.merlin.keystore.type</code>	The Bouncy Castle keystore provider supports the following types of keystore: JKS and PKCS12. In addition, Bouncy Castle supports the following proprietary keystore types: BKS and UBER.
<code>org.apache.ws.security.crypto.merlin.keystore.file</code>	Specifies the location of the keystore file to load, where the location is specified relative to the Classpath.
<code>org.apache.ws.security.crypto.merlin.keystore.alias</code>	(Optional) If the keystore type is JKS (Java keystore), you can select a specific key from the keystore by

¹ <http://www.bouncycastle.org/specifications.html>

Property	Description
	specifying its alias. If the keystore contains only one key, there is no need to specify an alias.
<code>org.apache.ws.security.crypto.merlin.keystore.password</code>	The password specified by this property is used for two purposes: to unlock the keystore (keystore password) and to decrypt a private key that is stored in the keystore (private key password). Hence, the keystore password must be same as the private key password.

For example, the `etc/alice.properties` file contains property settings to load the PKCS#12 file, `certs/alice.pfx`, as follows:

```
org.apache.ws.security.crypto.provider=org.apache.ws.security.components.crypto.Merlin
org.apache.ws.security.crypto.merlin.keystore.type=PKCS12
org.apache.ws.security.crypto.merlin.keystore.password=password
org.apache.ws.security.crypto.merlin.keystore.file=certs/alice.pfx
```

The `etc/bob.properties` file contains property settings to load the PKCS#12 file, `certs/bob.pfx`, as follows:

```
org.apache.ws.security.crypto.provider=org.apache.ws.security.components.crypto.Merlin
org.apache.ws.security.crypto.merlin.keystore.password=password
# for some reason, bouncycastle has issues with bob.pfx
org.apache.ws.security.crypto.merlin.keystore.provider=SunJSSE
org.apache.ws.security.crypto.merlin.keystore.type=PKCS12
org.apache.ws.security.crypto.merlin.keystore.file=certs/bob.pfx
```

Programming encryption keys and signing keys

An alternative approach to loading encryption keys and signing keys is to use the properties shown in [Table 6.3 on page 108](#) to specify Crypto objects that load the relevant keys. This requires you to provide your own implementation of the WSS4J Crypto interface, `org.apache.ws.security.components.crypto.Crypto`.

Table 6.3. Properties for Specifying Crypto Objects

Property	Description
<code>ws-security.signature.crypto</code>	Specifies an instance of a Crypto object that is responsible for loading the keys for signing and decrypting messages.

Property	Description
<code>ws-security.encryption.crypto</code>	Specifies an instance of a Crypto object that is responsible for loading the keys for encrypting messages and verifying signatures.

WSS4J Crypto interface

[Example 6.7 on page 109](#) shows the definition of the Crypto interface that you can implement, if you want to provide encryption keys and signing keys by programming. For more information, see the [WSS4J home page](#)².

Example 6.7. WSS4J Crypto Interface

```
// Java
package org.apache.ws.security.components.crypto;

import org.apache.ws.security.WSSecurityException;

import java.io.InputStream;
import java.math.BigInteger;
import java.security.KeyStore;
import java.security.PrivateKey;
import java.security.cert.Certificate;
import java.security.cert.CertificateFactory;
import java.security.cert.X509Certificate;

public interface Crypto {
    X509Certificate loadCertificate(InputStream in)
        throws WSSecurityException;

    X509Certificate[] getX509Certificates(byte[] data, boolean reverse)
        throws WSSecurityException;

    byte[] getCertificateData(boolean reverse, X509Certificate[] certs)
        throws WSSecurityException;

    public PrivateKey getPrivateKey(String alias, String password)
        throws Exception;

    public X509Certificate[] getCertificates(String alias)
        throws WSSecurityException;

    public String getAliasForX509Cert(Certificate cert)
        throws WSSecurityException;
}
```

² <http://ws.apache.org/wss4j/>

```
public String getAliasForX509Cert(String issuer)
throws WSSecurityException;

public String getAliasForX509Cert(String issuer, BigInteger serialNumber)
throws WSSecurityException;

public String getAliasForX509Cert(byte[] skiBytes)
throws WSSecurityException;

public String getDefaultX509Alias();

public byte[] getSKIBytesFromCert(X509Certificate cert)
throws WSSecurityException;

public String getAliasForX509CertThumb(byte[] thumb)
throws WSSecurityException;

public KeyStore getKeyStore();

public CertificateFactory getCertificateFactory()
throws WSSecurityException;

public boolean validateCertPath(X509Certificate[] certs)
throws WSSecurityException;

public String[] getAliasesForDN(String subjectDN)
throws WSSecurityException;
}
```

Specifying the Algorithm Suite

Overview

An algorithm suite is a coherent collection of cryptographic algorithms for performing operations such as signing, encryption, generating message digests, and so on.

For reference purposes, this section describes the algorithm suites defined by the WS-SecurityPolicy specification. Whether or not a particular algorithm suite is available, however, depends on the underlying security provider. Fuse Services Framework security is based on the pluggable Java Cryptography Extension (JCE) and Java Secure Socket Extension (JSSE) layers. By default, Fuse Services Framework is configured with Sun's JSSE provider, which supports the cipher suites described in [Appendix A³](http://java.sun.com/j2se/1.5.0/docs/guide/security/jsse/JSSERefGuide.html#AppA) of Sun's *JSSE Reference Guide*.

Syntax

The AlgorithmSuite element has the following syntax:

```
<sp:AlgorithmSuite xmlns:sp="..." ... >
  <wsp:Policy xmlns:wsp="...">
    (<sp:Basic256 ... /> |
     <sp:Basic192 ... /> |
     <sp:Basic128 ... /> |
     <sp:TripleDes ... /> |
     <sp:Basic256Rsa15 ... /> |
     <sp:Basic192Rsa15 ... /> |
     <sp:Basic128Rsa15 ... /> |
     <sp:TripleDesRsa15 ... /> |
     <sp:Basic256Sha256 ... /> |
     <sp:Basic192Sha256 ... /> |
     <sp:Basic128Sha256 ... /> |
     <sp:TripleDesSha256 ... /> |
     <sp:Basic256Sha256Rsa15 ... /> |
     <sp:Basic192Sha256Rsa15 ... /> |
     <sp:Basic128Sha256Rsa15 ... /> |
     <sp:TripleDesSha256Rsa15 ... /> |
     ...)
    <sp>InclusiveC14N ... /> ?
    <sp:SOAPNormalization10 ... /> ?
    <sp:STRTransform10 ... /> ?
    (<sp:XPath10 ... /> |
     <sp:XPathFilter20 ... /> |
     <sp:AbsXPath ... /> |
     ...)?
```

³ <http://java.sun.com/j2se/1.5.0/docs/guide/security/jsse/JSSERefGuide.html#AppA>

```

    ...
    </wsp:Policy>
    ...
</sp:AlgorithmSuite>

```

The algorithm suite assertion supports a large number of alternative algorithms (for example, Basic256). For a detailed description of the algorithm suite alternatives, see [Table 6.4 on page 112](#).

Algorithm suites

[Table 6.4 on page 112](#) provides a summary of the algorithm suites supported by WS-SecurityPolicy. The column headings refer to different types of cryptographic algorithm, as follows: [Dig] is the digest algorithm; [Enc] is the encryption algorithm; [Sym KW] is the symmetric key-wrap algorithm; [Asym KW] is the asymmetric key-wrap algorithm; [Enc KD] is the encryption key derivation algorithm; [Sig KD] is the signature key derivation algorithm.

Table 6.4. Algorithm Suites

Algorithm Suite	[Dig]	[Enc]	[Sym KW]	[Asym KW]	[Enc KD]	[Sig KD]
Basic256	Sha1	Aes256	KwAes256	KwRsa0aep	PSha1L256	PSha1L192
Basic192	Sha1	Aes192	KwAes192	KwRsa0aep	PSha1L192	PSha1L192
Basic128	Sha1	Aes128	KwAes128	KwRsa0aep	PSha1L128	PSha1L128
TripleDes	Sha1	TripleDes	KwTripleDes	KwRsa0aep	PSha1L192	PSha1L192
Basic256Rsa15	Sha1	Aes256	KwAes256	KwRsa15	PSha1L256	PSha1L192
Basic192Rsa15	Sha1	Aes192	KwAes192	KwRsa15	PSha1L192	PSha1L192
Basic128Rsa15	Sha1	Aes128	KwAes128	KwRsa15	PSha1L128	PSha1L128
TripleDesRsa15	Sha1	TripleDes	KwTripleDes	KwRsa15	PSha1L192	PSha1L192
Basic256Sha256	Sha256	Aes256	KwAes256	KwRsa0aep	PSha1L256	PSha1L192
Basic192Sha256	Sha256	Aes192	KwAes192	KwRsa0aep	PSha1L192	PSha1L192
Basic128Sha256	Sha256	Aes128	KwAes128	KwRsa0aep	PSha1L128	PSha1L128
TripleDesSha256	Sha256	TripleDes	KwTripleDes	KwRsa0aep	PSha1L192	PSha1L192
Basic256Sha256Rsa15	Sha256	Aes256	KwAes256	KwRsa15	PSha1L256	PSha1L192
Basic192Sha256Rsa15	Sha256	Aes192	KwAes192	KwRsa15	PSha1L192	PSha1L192
Basic128Sha256Rsa15	Sha256	Aes128	KwAes128	KwRsa15	PSha1L128	PSha1L128
TripleDesSha256Rsa15	Sha256	TripleDes	KwTripleDes	KwRsa15	PSha1L192	PSha1L192

Types of cryptographic algorithm

The following types of cryptographic algorithm are supported by WS-SecurityPolicy:

- ["Symmetric key signature" on page 113](#)
- ["Asymmetric key signature" on page 113](#)
- ["Digest" on page 114](#)
- ["Encryption" on page 114](#)
- ["Symmetric key wrap" on page 114](#)
- ["Asymmetric key wrap" on page 115](#)
- ["Computed key" on page 115](#)
- ["Encryption key derivation" on page 115](#)
- ["Signature key derivation" on page 116](#)

Symmetric key signature

The symmetric key signature property, [Sym Sig], specifies the algorithm for generating a signature using a symmetric key. WS-SecurityPolicy specifies that the HmacSha1 algorithm is always used.

The HmacSha1 algorithm is identified by the following URI:

```
http://www.w3.org/2000/09/xmldsig#hmac-sha1
```

Asymmetric key signature

The asymmetric key signature property, [Asym Sig], specifies the algorithm for generating a signature using an asymmetric key. WS-SecurityPolicy specifies that the Rsasha1 algorithm is always used.

The Rsasha1 algorithm is identified by the following URI:

```
http://www.w3.org/2000/09/xmldsig#rsa-sha1
```

Digest

The digest property, [Dig], specifies the algorithm used for generating a message digest value. WS-SecurityPolicy supports two alternative digest algorithms: Sha1 and Sha256.

The Sha1 algorithm is identified by the following URI:

```
http://www.w3.org/2000/09/xmldsig#sha1
```

The Sha256 algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmenc#sha256
```

Encryption

The encryption property, [Enc], specifies the algorithm used for encrypting data. WS-SecurityPolicy supports the following encryption algorithms: Aes256, Aes192, Aes128, TripleDes.

The Aes256 algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmenc#aes256-cbc
```

The Aes192 algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmenc#aes192-cbc
```

The Aes128 algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmenc#aes128-cbc
```

The TripleDes algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmenc#tripledes-cbc
```

Symmetric key wrap

The symmetric key wrap property, [Sym KW], specifies the algorithm used for signing and encrypting symmetric keys. WS-SecurityPolicy supports the following symmetric key wrap algorithms: KwAes256, KwAes192, KwAes128, KwTripleDes.

The KwAes256 algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmenc#kw-aes256
```

The KwAes192 algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmlescenc#kw-aes128
```

The KwAes128 algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmlescenc#kw-aes128
```

The KwTripleDes algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmlescenc#tripledes-cbc
```

Asymmetric key wrap

The asymmetric key wrap property, [Asym KW], specifies the algorithm used for signing and encrypting asymmetric keys. WS-SecurityPolicy supports the following asymmetric key wrap algorithms: KwRsa0aep, KwRsa15.

The KwRsa0aep algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmlescenc#rsa-oaep-mgf1p
```

The KwRsa15 algorithm is identified by the following URI:

```
http://www.w3.org/2001/04/xmlescenc#rsa-1_5
```

Computed key

The computed key property, [Comp Key], specifies the algorithm used to compute a derived key. When secure parties communicate with the aid of a shared secret key (for example, when using WS-SecureConversation), it is recommended that a derived key is used instead of the original shared key, in order to avoid exposing too much data for analysis by hostile third parties. WS-SecurityPolicy specifies that the PSha1 algorithm is always used.

The PSha1 algorithm is identified by the following URI:

```
http://docs.oasis-open.org/ws-sx/ws-secureconversation/200512/dk/p_sha1
```

Encryption key derivation

The encryption key derivation property, [Enc KD], specifies the algorithm used to compute a derived encryption key. WS-SecurityPolicy supports the following encryption key derivation algorithms: PSha1L256, PSha1L192, PSha1L128.

The PSha1 algorithm is identified by the following URI (the same algorithm is used for PSha1L256, PSha1L192, and PSha1L128; just the key lengths differ):

http://docs.oasis-open.org/ws-sx/ws-secureconversation/200512/dk/p_sha1

Signature key derivation

The signature key derivation property, [Sig KD], specifies the algorithm used to compute a derived signature key. WS-SecurityPolicy supports the following signature key derivation algorithms: PSha1L192, PSha1L128.

Key length properties

[Table 6.5 on page 116](#) shows the minimum and maximum key lengths supported in WS-SecurityPolicy.

Table 6.5. Key Length Properties

Property	Key Length
Minimum symmetric key length [Min SKL]	128, 192, 256
Maximum symmetric key length [Max SKL]	256
Minimum asymmetric key length [Min AKL]	1024
Maximum asymmetric key length [Max AKL]	4096

The value of the minimum symmetric key length, [Min SKL], depends on which algorithm suite is selected.

Chapter 7. Authentication

This chapter describes how to use policies to configure authentication in a Fuse Services Framework application. Currently, the only credentials type supported in the SOAP layer is the WS-Security UsernameToken.

Introduction to Authentication	118
Specifying an Authentication Policy	119
Providing Client Credentials	127
Authenticating Received Credentials	132

Introduction to Authentication

Overview

In Fuse Services Framework, an application can be set up to use authentication through a combination of policy assertions in the WSDL contract and configuration settings in Spring XML.



Note

Remember, you can also use the HTTPS protocol as the basis for authentication and, in some cases, this might be easier to configure. See ["Authentication Alternatives" on page 50](#).

Steps to set up authentication

In outline, you need to perform the following steps to set up an application to use authentication:

1. Add a supporting tokens policy to an endpoint in the WSDL contract. This has the effect of requiring the endpoint to include a particular type of token (client credentials) in its request messages.
2. On the client side, provide credentials to send by configuring the relevant endpoint in Spring XML.
3. (*Optional*) On the client side, if you decide to provide passwords using a callback handler, implement the callback handler in Java.
4. On the server side, associate a callback handler class with the endpoint in Spring XML. The callback handler is then responsible for authenticating the credentials received from remote clients.

Specifying an Authentication Policy

Overview

If you want an endpoint to support authentication, associate a *supporting tokens policy assertion* with the relevant endpoint binding. There are several different kinds of supporting tokens policy assertions, whose elements all have names of the form **SupportingTokens* (for example, *SupportingTokens*, *SignedSupportingTokens*, and so on). For a complete list, see ["SupportingTokens assertions" on page 123](#).

Associating a supporting tokens assertion with an endpoint has the following effects:

- Messages to or from the endpoint are required to include the specified token type (where the token's direction is specified by the `sp:IncludeToken` attribute).
- Depending on the particular type of supporting tokens element you use, the endpoint might be required to sign and/or encrypt the token.

The supporting tokens assertion implies that the runtime will check that these requirements are satisfied. But the WS-SecurityPolicy policies do *not* define the mechanism for providing credentials to the runtime. You must use Spring XML configuration to specify the credentials (see ["Providing Client Credentials" on page 127](#)).

Syntax

The **SupportingTokens* elements (that is, all elements with the *SupportingTokens* suffix—see ["SupportingTokens assertions" on page 123](#)) have the following syntax:

```
<sp:SupportingTokensElement xmlns:sp="..." ... >
  <wsp:Policy xmlns:wsp="...">
    [Token Assertion]+
    <sp:AlgorithmSuite ... > ... </sp:AlgorithmSuite> ?
    (
      <sp:SignedParts ... > ... </sp:SignedParts> |
      <sp:SignedElements ... > ... </sp:SignedElements> |
      <sp:EncryptedParts ... > ... </sp:EncryptedParts> |
      <sp:EncryptedElements ... > ... </sp:EncryptedElements> |
    ) *
    ...
  </wsp:Policy>
  ...
</sp:SupportingTokensElement>
```

Where *SupportingTokensElement* stands for one of the supporting token elements, **SupportingTokens*. Typically, if you simply want to include a token (or tokens) in the security header, you would include one or more token assertions, `[Token Assertion]`, in the policy. In particular, this is all that is required for authentication.

If the token is of an appropriate type (for example, an X.509 certificate or a symmetric key), you could theoretically also use it to sign or encrypt specific parts of the current message using the `sp:AlgorithmSuite`, `sp:SignedParts`, `sp:SignedElements`, `sp:EncryptedParts`, and `sp:EncryptedElements` elements. This functionality is currently *not* supported by Fuse Services Framework, however.

Sample policy

[Example 7.1 on page 120](#) shows an example of a policy that requires a WS-Security UsernameToken token (which contains username/password credentials) to be included in the security header. In addition, because the token is specified inside an `sp:SignedSupportingTokens` element, the policy requires that the token is signed. This example uses a transport binding, so it is the underlying transport that is responsible for signing the message.

For example, if the underlying transport is HTTPS, the SSL/TLS protocol (configured with an appropriate algorithm suite) is responsible for signing the *entire* message, including the security header that contains the specified token. This is sufficient to satisfy the requirement that the supporting token is signed.

Example 7.1. Example of a Supporting Tokens Policy

```
<wsp:Policy wsu:Id="UserNameOverTransport_IPingService_policy">
  <wsp:ExactlyOne>
    <wsp:All>
      <sp:TransportBinding> ... </sp:TransportBinding>
      <sp:SignedSupportingTokens
        xmlns:sp="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy">
        <wsp:Policy>
          <sp:UsernameToken
            sp:IncludeToken="http://schemas.xmlsoap.org/ws/2005/07/securitypolicy/Inclu
deToken/AlwaysToRecipient">
            <wsp:Policy>
              <sp:WssUsernameToken10/>
            </wsp:Policy>
          </sp:UsernameToken>
        </wsp:Policy>
      </sp:SignedSupportingTokens>
      ...
    </wsp:All>
  </wsp:ExactlyOne>
</wsp:Policy>
```

Where the presence of the `sp:WssUsernameToken10` sub-element indicates that the UsernameToken header should conform to version 1.0 of the WS-Security UsernameToken specification.

Token types

In principle, you can specify any of the WS-SecurityPolicy token types in a supporting tokens assertion. For SOAP-level authentication, however, only the `sp:UsernameToken` token type is relevant.

sp:UsernameToken

In the context of a supporting tokens assertion, this element specifies that a WS-Security UsernameToken is to be included in the security SOAP header. Essentially, a WS-Security UsernameToken is used to send username/password credentials in the WS-Security SOAP header. The sp:UsernameToken element has the following syntax:

```
<sp:UsernameToken sp:IncludeToken="xs:anyURI"? xmlns:sp="..." ... >
  (
    <sp:Issuer>wsa:EndpointReferenceType</sp:Issuer> |
    <sp:IssuerName>xs:anyURI</sp:IssuerName>
  ) ?
  <wst:Claims Dialect="..."> ... </wst:Claims> ?
  <wsp:Policy xmlns:wsp="...">
    (
      <sp:NoPassword ... /> |
      <sp:HashPassword ... />
    ) ?
    (
      <sp:RequireDerivedKeys /> |
      <sp:RequireImpliedDerivedKeys ... /> |
      <sp:RequireExplicitDerivedKeys ... />
    ) ?
    (
      <sp:WssUsernameToken10 ... /> |
      <sp:WssUsernameToken11 ... />
    ) ?
    ...
  </wsp:Policy>
  ...
</sp:UsernameToken>
```

The sub-elements of sp:UsernameToken are all optional and are not needed for ordinary authentication. Normally, the only part of this syntax that is relevant is the sp:IncludeToken attribute.



Note

Currently, in the sp:UsernameToken syntax, only the sp:WssUsernameToken10 sub-element is supported in Fuse Services Framework.

sp:IncludeToken attribute

The value of the sp:IncludeToken must match the WS-SecurityPolicy version from the enclosing policy. The current version is 1.2, but legacy WSDL might use version 1.1. Valid values of the sp:IncludeToken attribute are as follows:

Never

The token **MUST NOT** be included in any messages sent between the initiator and the recipient; rather, an external reference to the token should be used. Valid URI values are:

1.2	http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702/ IncludeToken/Never
1.1	http://schemas.xmlsoap.org/ws/2005/07/securitypolicy/ IncludeToken/Never

Once

The token **MUST** be included in only one message sent from the initiator to the recipient. References to the token **MAY** use an internal reference mechanism. Subsequent related messages sent between the recipient and the initiator may refer to the token using an external reference mechanism. Valid URI values are:

1.2	http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702/ IncludeToken/Once
1.1	http://schemas.xmlsoap.org/ws/2005/07/securitypolicy/ IncludeToken/Once

AlwaysToRecipient

The token **MUST** be included in all messages sent from initiator to the recipient. The token **MUST NOT** be included in messages sent from the recipient to the initiator. Valid URI values are:

1.2	http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702/ IncludeToken/AlwaysToRecipient
1.1	http://schemas.xmlsoap.org/ws/2005/07/securitypolicy/ IncludeToken/AlwaysToRecipient

AlwaysToInitiator

The token **MUST** be included in all messages sent from the recipient to the initiator. The token **MUST NOT** be included in messages sent from the initiator to the recipient. Valid URI values are:

1.2	http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702/ IncludeToken/AlwaysToInitiator
1.1	http://schemas.xmlsoap.org/ws/2005/07/securitypolicy/ IncludeToken/AlwaysToInitiator

Always

The token **MUST** be included in all messages sent between the initiator and the recipient. This is the default behavior. Valid URI values are:

1.2	http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702/ IncludeToken/Always
1.1	http://schemas.xmlsoap.org/ws/2005/07/securitypolicy/ IncludeToken/Always

SupportingTokens assertions

The following kinds of supporting tokens assertions are supported:

- ["sp:SupportingTokens" on page 123.](#)
- ["sp:SignedSupportingTokens" on page 124.](#)
- ["sp:EncryptedSupportingTokens" on page 124.](#)
- ["sp:SignedEncryptedSupportingTokens" on page 124.](#)
- ["sp:EndorsingSupportingTokens" on page 124.](#)
- ["sp:SignedEndorsingSupportingTokens" on page 125.](#)
- ["sp:EndorsingEncryptedSupportingTokens" on page 125.](#)
- ["sp:SignedEndorsingEncryptedSupportingTokens" on page 126.](#)

sp:SupportingTokens

This element requires a token (or tokens) of the specified type to be included in the `wsse:Security` header. No additional requirements are imposed.



Warning

This policy does not explicitly require the tokens to be signed or encrypted. It is normally essential, however, to protect tokens by signing and encryption.

sp:SignedSupportingTokens

This element requires a token (or tokens) of the specified type to be included in the `wsse:Security` header. In addition, this policy requires that the token is signed, in order to guarantee token integrity.



Warning

This policy does not explicitly require the tokens to be encrypted. It is normally essential, however, to protect tokens both by signing and encryption.

sp:EncryptedSupportingTokens

This element requires a token (or tokens) of the specified type to be included in the `wsse:Security` header. In addition, this policy requires that the token is encrypted, in order to guarantee token confidentiality.



Warning

This policy does not explicitly require the tokens to be signed. It is normally essential, however, to protect tokens both by signing and encryption.

sp:SignedEncryptedSupportingTokens

This element requires a token (or tokens) of the specified type to be included in the `wsse:Security` header. In addition, this policy requires that the token is both signed and encrypted, in order to guarantee token integrity and confidentiality.

sp:EndorsingSupportingTokens

An endorsing supporting token is used to sign the message signature (primary signature). This signature is known as an *endorsing signature* or *secondary signature*. Hence, by applying an endorsing supporting tokens policy, you can have a chain of signatures: the primary signature, which signs the message itself, and the secondary signature, which signs the primary signature.



Note

If you are using a transport binding (for example, HTTPS), the message signature is not actually part of the SOAP message, so it is not possible to sign the message signature in this case. If you specify this policy with a transport binding, the endorsing token signs the timestamp instead.



Warning

This policy does not explicitly require the tokens to be signed or encrypted. It is normally essential, however, to protect tokens by signing and encryption.

sp:SignedEndorsingSupportingTokens

This policy is the same as the endorsing supporting tokens policy, except that the tokens are required to be signed, in order to guarantee token integrity.



Warning

This policy does not explicitly require the tokens to be encrypted. It is normally essential, however, to protect tokens both by signing and encryption.

sp:EndorsingEncryptedSupportingTokens

This policy is the same as the endorsing supporting tokens policy, except that the tokens are required to be encrypted, in order to guarantee token confidentiality.



Warning

This policy does not explicitly require the tokens to be signed. It is normally essential, however, to protect tokens both by signing and encryption.

sp:SignedEndorsingEncryptedSupportingTokens

This policy is the same as the endorsing supporting tokens policy, except that the tokens are required to be signed and encrypted, in order to guarantee token integrity and confidentiality.

Providing Client Credentials

Overview

There are essentially two approaches to providing UsernameToken client credentials: you can either set both the username and the password directly in the client's Spring XML configuration; or you can set the username in the client's configuration and implement a callback handler to provide passwords programmatically. The latter approach (by programming) has the advantage that passwords are easier to hide from view.

Client credentials properties

[Table 7.1 on page 127](#) shows the properties you can use to specify WS-Security username/password credentials on a client's request context in Spring XML.

Table 7.1. Client Credentials Properties

Properties	Description
ws-security.username	Specifies the username for UsernameToken policy assertions.
ws-security.password	Specifies the password for UsernameToken policy assertions. If not specified, the password is obtained by calling the callback handler.
ws-security.callback-handler	Specifies the class name of the WSS4J callback handler that retrieves passwords for UsernameToken policy assertions. Note that the callback handler can also handle other kinds of security events.

Configuring client credentials in Spring XML

To configure username/password credentials in a client's request context in Spring XML, set the ws-security.username and ws-security.password properties as follows:

```
<beans ... >
  <jaxws:client name="{NamespaceName}LocalPortName"
    createdFromAPI="true">
    <jaxws:properties>
      <entry key="ws-security.username" value="Alice"/>
      <entry key="ws-security.password" value="abcd!1234"/>
    </jaxws:properties>
  </jaxws:client>
  ...
</beans>
```

If you prefer not to store the password directly in Spring XML (which might potentially be a security hazard), you can provide passwords using a callback handler instead.

Programming a callback handler for passwords

If you want to use a callback handler to provide passwords for the UsernameToken header, you must first modify the client configuration in Spring XML, replacing the `ws-security.password` setting by a `ws-security.callback-handler` setting, as follows:

```
<beans ... >
  <jaxws:client name="{NamespaceName}LocalPortName"
    createdFromAPI="true">
    <jaxws:properties>
      <entry key="ws-security.username" value="Alice"/>
      <entry key="ws-security.callback-handler" value="interop.client.UTPasswordCall
back"/>
    </jaxws:properties>
  </jaxws:client>
  ...
</beans>
```

In the preceding example, the callback handler is implemented by the `UTPasswordCallback` class. You can write a callback handler by implementing the `javax.security.auth.callback.CallbackHandler` interface, as shown in [Example 7.2 on page 128](#).

Example 7.2. Callback Handler for UsernameToken Passwords

```
package interop.client;

import java.io.IOException;
import java.util.HashMap;
import java.util.Map;

import javax.security.auth.callback.Callback;
import javax.security.auth.callback.CallbackHandler;
import javax.security.auth.callback.UnsupportedCallbackException;

import org.apache.ws.security.WSPasswordCallback;

public class UTPasswordCallback implements CallbackHandler {

    private Map<String, String> passwords =
        new HashMap<String, String>();

    public UTPasswordCallback() {
        passwords.put("Alice", "ecila");
        passwords.put("Frank", "invalid-password");
        //for MS clients
    }
}
```



```

        passwords.put("abcd", "dcba");
    }

    public void handle(Callback[] callbacks) throws IOException, UnsupportedCallbackException
    {
        for (int i = 0; i < callbacks.length; i++) {
            WSPasswordCallback pc = (WSPasswordCallback)callbacks[i];

            String pass = passwords.get(pc.getIdentifier());
            if (pass != null) {
                pc.setPassword(pass);
                return;
            }
        }

        throw new IOException();
    }

    // Add an alias/password pair to the callback mechanism.
    public void setAliasPassword(String alias, String password) {
        passwords.put(alias, password);
    }
}

```

The callback functionality is implemented by the `CallbackHandler.handle()` method. In this example, it assumed that the callback objects passed to the `handle()` method are all of [org.apache.ws.security.WSPasswordCallback](http://ws.apache.org/wss4j/apidocs/org/apache/ws/security/WSPasswordCallback.html)¹ type (in a more realistic example, you would check the type of the callback objects).

A more realistic implementation of a client callback handler would probably consist of prompting the user to enter their password.

WSPasswordCallback class

When a `CallbackHandler` is called in a Fuse Services Framework client for the purpose of setting a UsernameToken password, the corresponding `WSPasswordCallback` object has the `USERNAME_TOKEN` usage code.

For more details about the `WSPasswordCallback` class, see [org.apache.ws.security.WSPasswordCallback](http://ws.apache.org/wss4j/apidocs/org/apache/ws/security/WSPasswordCallback.html)².

The `WSPasswordCallback` class defines several different usage codes, as follows:

¹ <http://ws.apache.org/wss4j/apidocs/org/apache/ws/security/WSPasswordCallback.html>

² <http://ws.apache.org/wss4j/apidocs/org/apache/ws/security/WSPasswordCallback.html>

USERNAME_TOKEN

Obtain the password for UsernameToken credentials. This usage code is used both on the client side (to obtain a password to send to the server) and on the server side (to obtain a password in order to compare it with the password received from the client).

On the server side, this code is set in the following cases:

- *Digest password*—if the UsernameToken contains a digest password, the callback must return the corresponding password for the given user name (given by `WSPasswordCallback.getIdentifier()`). Verification of the password (by comparing with the digest password) is done by the WSS4J runtime.
- *Plaintext password*—implemented the same way as the digest password case (since Fuse Services Framework 2.4.0).
- *Custom password type*—if `getHandleCustomPasswordTypes()` is true on `org.apache.ws.security.WSSConfig`, this case is implemented the same way as the digest password case (since Fuse Services Framework 2.4.0). Otherwise, an exception is thrown.

If no Password element is included in a received UsernameToken on the server side, the callback handler is not called (since Fuse Services Framework 2.4.0).

DECRYPT

Need a password to retrieve a private key from a Java keystore, where `WSPasswordCallback.getIdentifier()` gives the alias of the keystore entry. WSS4J uses this private key to decrypt the session (symmetric) key.

SIGNATURE

Need a password to retrieve a private key from a Java keystore, where `WSPasswordCallback.getIdentifier()` gives the alias of the keystore entry. WSS4J uses this private key to produce a signature.

SECRET_KEY

Need a secret key for encryption or signature on the outbound side, or for decryption or verification on the inbound side. The callback handler must set the key using the `setKey(byte[])` method.

SECURITY_CONTEXT_TOKEN

Need the key for a `wsc:SecurityContextToken`, which you provide by calling the `setKey(byte[])` method.

CUSTOM_TOKEN

Need a token as a DOM element. For example, this is used for the case of a reference to a SAML Assertion or SecurityContextToken that is not in the message. The callback handler must set the token using the `setCustomToken(Element)` method.

KEY_NAME

(Obsolete) Since Fuse Services Framework 2.4.0, this usage code is obsolete.

USERNAME_TOKEN_UNKNOWN

(Obsolete) Since Fuse Services Framework 2.4.0, this usage code is obsolete.

UNKNOWN

Not used by WSS4J.

Authenticating Received Credentials

Overview

On the server side, you can verify that received credentials are authentic by registering a callback handler with the Fuse Services Framework runtime. You can either write your own custom code to perform credentials verification or you can implement a callback handler that integrates with a third-party enterprise security system (for example, an LDAP server).

Configuring a server callback handler in Spring XML

To configure a server callback handler that verifies UsernameToken credentials received from clients, set the `ws-security.callback-handler` property in the server's Spring XML configuration, as follows:

```
<beans ... >
  <jaxws:endpoint
    id="UserNameOverTransport"
    address="https://localhost:9001/UserNameOverTransport"
    serviceName="interop:PingService10"
    endpointName="interop:UserNameOverTransport_IPingService"
    implementor="interop.server.UserNameOverTransport"
    depends-on="tls-settings">

    <jaxws:properties>
      <entry key="ws-security.username" value="Alice"/>
      <entry key="ws-security.callback-handler" value="interop.client.UTPasswordCall
back"/>
    </jaxws:properties>

  </jaxws:endpoint>
  ...
</beans>
```

In the preceding example, the callback handler is implemented by the `UTPasswordCallback` class.

Implementing the callback handler to check passwords

To implement a callback handler for checking passwords on the server side, implement the `javax.security.auth.callback.CallbackHandler` interface. The general approach to implementing the `CallbackHandler` interface for a server is similar to implementing a `CallbackHandler` for a client. The interpretation given to the returned password on the server side is different, however: the password from the callback handler is compared against the received client password in order to verify the client's credentials.

For example, you could use the sample implementation shown in [Example 7.2 on page 128](#) to obtain passwords on the server side. On the server side, the WSS4J runtime would compare the password obtained from the

callback with the password in the received client credentials. If the two passwords match, the credentials are successfully verified.

A more realistic implementation of a server callback handler would involve writing an integration with a third-party database that is used to store security data (for example, integration with an LDAP server).

Chapter 8. WS-Trust

WS-Trust provides the necessary security infrastructure for supporting advanced authentication and authorization requirements. In particular, WS-Trust enables you to store security data in a central location (in the Security Token Service) and support various single sign-on scenarios.

Introduction to WS-Trust	136
Basic Scenarios	139
WS-Trust Single Sign-On Demonstration	143
Overview of the Demonstration	144
Configure the Security Token Service	146
Define the Security Policy	152
Configure the Client-STS Connection	155
Configure the Server-Side Interceptor	160
Build and Run the Demonstration	165
Sample Message Exchanges	166
Defining an IssuedToken Policy	173
Creating an STSClient Instance	178

Introduction to WS-Trust

Overview

The WS-Trust standard is based around a centralized security server (the Security Token Service), which is capable of authenticating clients and can issue tokens containing various kinds of authentication and authorization data.

WS-Trust specification

The WS-Trust features of Artix are based on the WS-Trust standard from [Oasis](http://docs.oasis-open.org/ws-sx/ws-trust/v1.4/ws-trust.html)¹:

<http://docs.oasis-open.org/ws-sx/ws-trust/v1.4/ws-trust.html>

Supporting specifications

Apart from the WS-Trust specification itself, several other specifications play an important role in the WS-Trust architecture, as follows:

- [WS-SecurityPolicy 1.2](http://docs.oasis-open.org/ws-sx/ws-securitypolicy/v1.2/ws-securitypolicy.html)²
- [SAML 2.0](http://saml.xml.org/saml-specifications)³
- [Username Token Profile](http://www.oasis-open.org/committees/download.php/16782/wss-v1.1-spec-os-UsernameTokenProfile.pdf)⁴
- [X.509 Token Profile](http://www.oasis-open.org/committees/download.php/16785/wss-v1.1-spec-os-x509TokenProfile.pdf)⁵
- [SAML Token Profile](http://www.oasis-open.org/committees/download.php/16768/wss-v1.1-spec-os-SAMLTOKENProfile.pdf)⁶
- [Kerberos Token Profile](http://www.oasis-open.org/committees/download.php/16788/wss-v1.1-spec-os-KerberosTokenProfile.pdf)⁷

WS-Trust architecture

[Figure 8.1 on page 137](#) shows a general overview of the WS-Trust architecture.

¹ <http://www.oasis-open.org>

² <http://docs.oasis-open.org/ws-sx/ws-securitypolicy/v1.2/ws-securitypolicy.html>

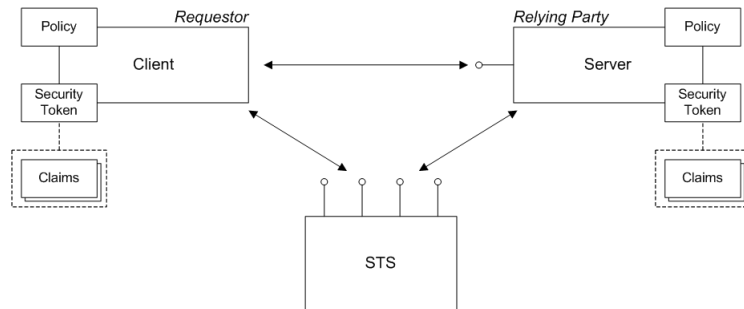
³ <http://saml.xml.org/saml-specifications>

⁴ <http://www.oasis-open.org/committees/download.php/16782/wss-v1.1-spec-os-UsernameTokenProfile.pdf>

⁵ <http://www.oasis-open.org/committees/download.php/16785/wss-v1.1-spec-os-x509TokenProfile.pdf>

⁶ <http://www.oasis-open.org/committees/download.php/16768/wss-v1.1-spec-os-SAMLTOKENProfile.pdf>

⁷ <http://www.oasis-open.org/committees/download.php/16788/wss-v1.1-spec-os-KerberosTokenProfile.pdf>

Figure 8.1. WS-Trust Architecture

Requestor

A *requestor* is an entity that tries to invoke a secure operation over a network connection. In practice, a requestor is typically a Web service client.

Relying party

A *relying party* refers to an entity that has some services or resources that must be secured against unauthorized access. In practice, a relying party is typically a Web service.

Note

This is a term defined by the SAML specification, not by WS-Trust.

Security token

A *security token* is a collection of security data that a requestor sends inside a request (typically embedded in the message header) in order to invoke a secure operation or to gain access to a secure resource. In the WS-Trust framework, the notion of a security token is quite general and can be used to describe any block of security data that might accompany a request.

In principle, WS-Trust can be used with the following kinds of security token:

- SAML token.
- UsernameToken token.

- X.509 certificate token.
- Kerberos token.

SAML token

A SAML token is a particularly flexible kind of security token. The [SAML specification](http://saml.xml.org/saml-specifications)⁸ defines a general-purpose XML schema that enables you to wrap almost any kind of security data and enables you to sign and encrypt part or all of the token.

SAML is a popular choice of token to use in the context of WS-Trust, because SAML has all of the necessary features to support typical WS-Trust authentication scenarios.

Claims

A SAML security token is formally defined to consist of a collection of *claims*. Each claim typically contains a particular kind of security data.

Policy

In WS-Trust scenarios, a *policy* can represent the security configuration of a participant in a secure application. The requestor, the relying party, and the security token service are all configured by policies. For example, a policy can be used to configure what kinds of authentication are supported and required.

Security token service

The *security token service* (STS) lies at the heart of the WS-Trust security architecture. In the WS-Trust standard, the following bindings are defined (not all of which are supported by Apache CXF):

- *Issue binding*—the specification defines this binding as follows: *Based on the credential provided/proven in the request, a new token is issued, possibly with new proof information.*
- *Validate binding*—the specification defines this binding as follows: *The validity of the specified security token is evaluated and a result is returned. The result may be a status, a new token, or both.*
- *Renew binding*—the specification defines this binding as follows: *A previously issued token with expiration is presented (and possibly proven) and the same token is returned with new expiration semantics.*
- *Cancel binding*—the specification defines this binding as follows: *When a previously issued token is no longer needed, the Cancel binding can be used to cancel the token, terminating its use.*

⁸ <http://saml.xml.org/saml-specifications>

Basic Scenarios

Overview

This section describes the basic scenarios supported by WS-Trust, which are closely related to some of the use cases defined in the SAML standard. It is, therefore, worth taking a moment to look at the relationship between the SAML standard and the WS-Trust standard.

SAML architecture

The SAML standard is specified in four distinct parts, as follows:

- *Assertions*—specifies the format of a SAML token, which is a standardized packet of XML containing security data. SAML tokens can contain authentication data (such as username/password, X.509 certificate, and so on), authorization data (such as roles, permissions, privileges), security attributes (such as issuer identity, name and address of subject). A SAML token can also optionally be encrypted and/or signed.
- *Protocol*—describes request and response messages for operations such as issuing, validating, and renewing SAML tokens.
- *Bindings*—maps the abstract SAML protocol to specific network protocols.
- *Profiles*—describes particular use cases for building a security system based on SAML.

WS-Trust and SAML

There are close parallels between the WS-Trust architecture and the SAML architecture. In particular, WS-Trust explicitly relies on and uses SAML *assertions* (that is, SAML tokens). On the other hand, WS-Trust does *not* use any of the *protocol*, *bindings*, or *profiles* components of the SAML standard.

The relationship between WS-Trust and SAML can be quite confusing, in fact, because WS-Trust *does* define an abstract protocol (for communicating with the STS), a binding (to the SOAP protocol), and scenarios that are remarkably similar to some of the SAML scenarios. But these aspects of WS-Trust are defined *independently* of the SAML standard.

Scenarios

SAML defines the following fundamental types of authentication scenario, which are also supported by WS-Trust:

- "Server-vouches scenario" .
- "Bearer scenario" .

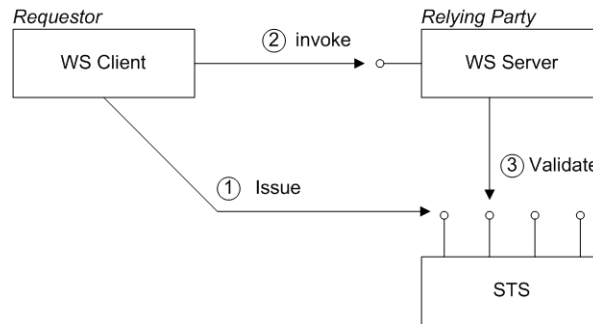
- "Holder-of-key scenario" .

Server-vouches scenario

In the *server-vouches* scenario, the SAML token that the client sends to the server contains insufficient data for the server to trust the incoming message. The server must perform an additional verification step to confirm that the SAML token is trustworthy.

Figure 8.2 on page 140 shows a general outline of a typical server-vouches scenario.

Figure 8.2. Server-Vouches Scenario



Steps in the server-vouches scenario

The server-vouches scenario proceeds as follows:

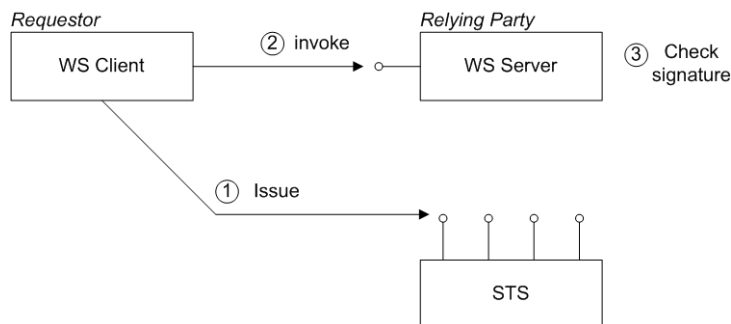
1. Before invoking an operation on the server, the client sends a RequestSecurityToken (RST) message to the Issue binding of the STS. The STS generates a SAML token with subject confirmation type *server-vouches* and returns the SAML token in a RequestSecurityTokenReply (RSTR) message.
2. The client attempts to invoke an operation on the server, with the SAML token embedded in the SOAP header of the request message.
3. Because the server does not have enough information to verify the received SAML token, the server sends the token to the Validate binding of the STS. The STS sends back confirmation of the token's validity in an RSTR message.

Bearer scenario

In the *bearer* scenario, the server automatically trusts the SAML token (after verifying its signature). Thus, in the bearer scenario, *any* client that presents the token can make use of the claims contained in the token (roles, permissions, and so on). It follows that the client must be very careful not to expose the SAML token or to pass it to any untrusted applications. For example, the client/server connection must use encryption, to protect the SAML token from snooping.

Figure 8.3 on page 141 shows a general outline of a typical bearer scenario.

Figure 8.3. Bearer Scenario



Steps in the bearer scenario

The bearer scenario proceeds as follows:

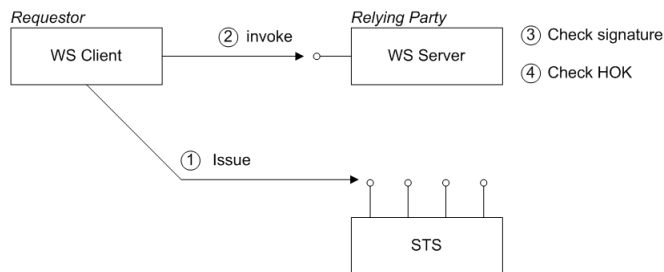
1. Before invoking an operation on the server, the client sends a `RequestSecurityToken` (RST) message to the `Issue` binding of the STS. The STS generates a SAML token with subject confirmation type *bearer*, signs the token using its private key, and then returns the token in a `RequestSecurityTokenReply` (RSTR) message.
2. The client attempts to invoke an operation on the server, with the SAML token embedded in the SOAP header of the request message, where either the SOAP header or the transport connection must be encrypted, to protect the token.
3. The server checks the signature of the SAML token (using a local copy of the STS public key), to ensure that it has not been tampered with.

Holder-of-key scenario

The *holder-of-key* scenario is a refinement of the bearer scenario where, instead of accepting the SAML token when presented by any client, the server attempts to authenticate the client and checks that the client identity matches the holder-of-key identity embedded in the SAML token.

Figure 8.4 on page 142 shows a general outline of a typical holder-of-key scenario.

Figure 8.4. Holder-of-Key Scenario



Steps in the holder-of-key scenario

The bearer scenario proceeds as follows:

1. Before invoking an operation on the server, the client sends a RequestSecurityToken (RST) message to the Issue binding of the STS. The STS generates a SAML token with subject confirmation type *holder-of-key*, embeds the client identity in the token (the holder-of-key identity), signs the token using its private key, and then returns the token in a RequestSecurityTokenReply (RSTR) message.
2. The client attempts to invoke an operation on the server, with the SAML token embedded in the SOAP header of the request message.
3. The server checks the signature of the SAML token (using a local copy of the STS public key), to ensure that it has not been tampered with.
4. The server attempts to authenticate the client (for example, by requiring a client X.509 certificate or by checking WS-Security UsernameToken credentials) and checks that the client's identity matches the holder-of-key identity.

WS-Trust Single Sign-On Demonstration

Overview of the Demonstration	144
Configure the Security Token Service	146
Define the Security Policy	152
Configure the Client-STS Connection	155
Configure the Server-Side Interceptor	160
Build and Run the Demonstration	165
Sample Message Exchanges	166

Overview of the Demonstration

Overview

This section describes how to set up and run a WS-Trust single sign-on (SSO) demonstration. The demonstration is based on the following Apache CXF samples;

`samples/sts_issue_operation`

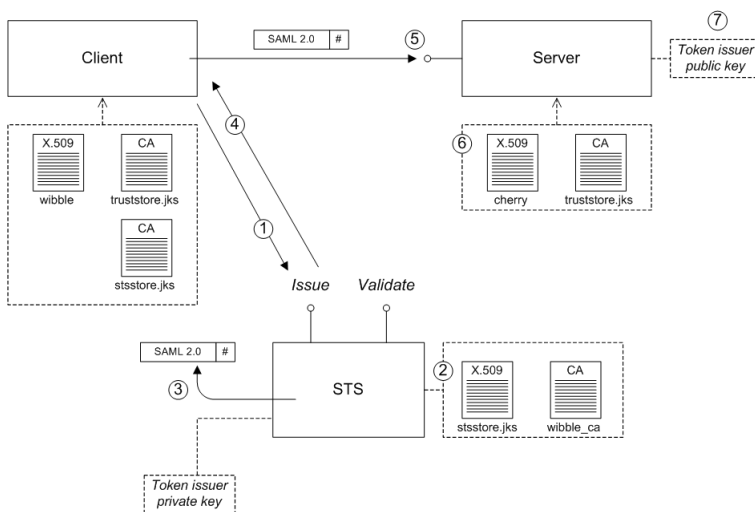
Implements a lightweight STS, supporting only the Issue operation.

`samples/wsd1_first_https`

A basic Web service client/server sample, with TLS enabled on the client/server connection.

The instructions in this section explain how to modify these two samples, in order to obtain a demonstration of WS-Trust single sign-on. [Figure 8.6 on page 146](#) gives an overview of the WS-Trust single sign-on scenario implemented by this demonstration.

Figure 8.5. WS-Trust Single Sign-On Scenario



Steps for single sign-on

The WS-Trust single sign-on scenario shown in [Figure 8.6 on page 146](#) proceeds as follows:

1. Before invoking an operation on the server for the first time, the client initiates SSO login by contacting the Issue binding on the STS.
2. The client presents its own X.509 certificate, `wibble`, during the TLS handshake. On the STS side of the connection, the JAX-WS endpoint verifies the client certificate using the CA certificate, `wibble_ca`, and on the client side, the client verifies the STS certificate using a stored copy of the STS certificate, `stsstore.jks`.
3. The STS creates a SAML 2.0 token and the token issuer private key is used to sign the SAML token (where the signature is shown as # in the figure). This enables relying parties (WS servers) to verify the integrity of the SAML token.
4. The STS replies to the client, sending back the signed SAML token.
5. The client initiates a connection to the server, in order to invoke an operation.
6. During the TLS handshake, the client presents its own X.509 certificate, `wibble`. On the server side of the connection, the JAX-WS endpoint verifies the client certificate using the CA certificate, `wibble_ca`, and on the client side, the client verifies the server certificate also using the CA certificate, `wibble_ca`.

After the connection is established, the client sends an invocation request to the server, which includes the SAML token embedded in a SOAP header.

7. The server tests the integrity of the received SAML token using the token issuer public key (which complements the token issuer private key used by the STS). If this test is successful, it proves that the SAML token has not been modified or corrupted since it was issued by the STS.

Demonstration parts

To configure and run the WS-Trust single sign-on demonstration, follow the instructions in these demonstration parts:

1. ["Configure the Security Token Service" on page 146.](#)
2. ["Define the Security Policy" on page 152.](#)
3. ["Configure the Client-STS Connection" on page 155.](#)
4. ["Configure the Server-Side Interceptor" on page 160.](#)
5. ["Build and Run the Demonstration" on page 165.](#)
6. ["Sample Message Exchanges" on page 166.](#)

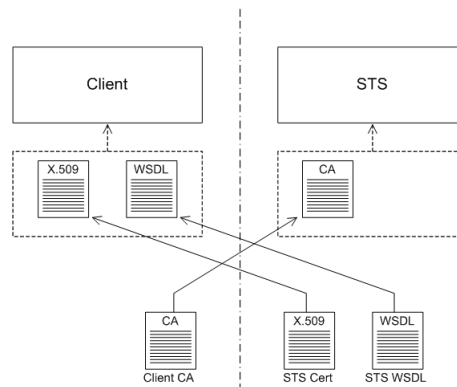
Configure the Security Token Service

Overview

The Fuse Services Framework samples include a demonstration implementation of a security token service. The functionality of this sample security token service is limited: it supports only the `Issue` operation and always returns a signed SAML token (which can conform either to SAML 1.1 or SAML 2.0).

In general, before a client can invoke operations on a given STS, a certain amount of setup preparation is required, as illustrated in [Figure 8.6 on page 146](#). The STS must provide its own X.509 certificate, enabling the client to authenticate the STS and to confirm signatures made by the STS, and the STS must provide a copy of its WSDL contract, which is a standard requirement for contacting any Web service. Assuming the client identifies itself using an X.509 certificate, it is also necessary for the client to provide a copy of its trusted CA certificate to the STS.

Figure 8.6. Installing Requisite Certificates and WSDL File



Location of the STS example

The sample WS-Trust security token service is located in the following sample directory:

```
CxfInstallDir/samples/sts_issue_operation
```

Checklist

In general, whenever you are integrating a WS client with a security token service, you will need to perform some or all of the following actions:

- *Install client's trusted CA certificate in STS*—normally, the connection between the client and the STS would be secured using SSL/TLS. If the STS decides to authenticate the client using an X.509 certificate (as in this demonstration), the STS will be configured to require a client certificate during the SSL/TLS handshake. The STS must have a copy of the client's trusted CA certificate in order to verify the signature of the received client certificate.



Note

Alternatively, an STS might be configured to authenticate the client using a mechanism such as UsernameToken credentials. In this case, an administrator would need to add the client's username/password to the STS's store of security data.

- *Install STS certificate in client*—likewise, in order for the client to verify that it is talking to the authentic STS, you need to install the STS certificate (or the CA certificate that signed the STS certificate) into the client's trust store.
- *Install STS WSDL in client*—just as with any other Web service, the client needs a copy of the STS's WSDL file in order to open a connection to the STS.
- *Customize address in STS WSDL*—often enough (and also in this demonstration), the SOAP address advertised in a WSDL file does not correspond to the actual address that the Web service is listening on. Because the client reads the STS endpoint's address from the STS WSDL file, it is essential to make sure that this address is accurate.

Steps to configure the STS

Perform the following steps to configure the security token service:

1. Configure the Maven Jetty plug-in with a secure SSL port.

The result of building the `sts_issue_operation` sample is a WAR file, which can then be deployed into any Web container. You can also deploy the WAR file into a Jetty container, which can conveniently be started up using the Maven Jetty plug-in. By default, however, the Jetty configuration in the `sts_issue_operation` POM file exposes an *insecure* HTTP port. To change the Jetty configuration to use a *secure* HTTPS port, edit the `pom.xml` file in the `samples/sts_issue_operation` directory, adding the `connectors` element as shown in the following extract:

```
<?xml version="1.0"?>
<project ... >
  ...
  <build>
```

```

...
<plugins>
...
<plugin>
  <groupId>org.mortbay.jetty</groupId>
  <artifactId>maven-jetty-plugin</artifactId>
  <version>6.1.24</version>
  <configuration>
    <contextPath>/${project.build.finalName}</contextPath>
    <connectors>
      <connector implementation="org.mortbay.jetty.security.SslSocketCon
nector">
        <port>8181</port>
        <keystore>src/main/resources/stsstore.jks</keystore>
        <password>stsspass</password>
        <keyPassword>stskpass</keyPassword>
        <truststore>src/main/resources/stsstore.jks</truststore>
        <trustPassword>stsspass</trustPassword>
      </connector>
    </connectors>
  </configuration>
</plugin>
</plugins>
</build>
</project>

```



Note

If you prefer to use Jetty version 7 or later, you should note that the package name for the `SslSocketConnector` class has changed, because Jetty 7 is now hosted by the Eclipse foundation.

2. Import the client's trusted CA certificate into the STS trust store file.

Copy the `cacert.pem` file from the `samples/wsd1_first_https/certs` directory to the `samples/sts_issue_operation/src/main/resources` directory. Install this CA cert into the STS's trust store, `stsstore.jks`, by opening a new command prompt, changing directory to `samples/sts_issue_operation/src/main/resources`, and entering the following command:

```
keytool -import -file cacert.pem -alias wibble_ca -keystore stsstore.jks -storepass stsspass
```

This command uses the Java `keytool` utility to install the client's CA certificate, `cacert.pem`, into the STS trust store, `stsstore.jks`, and assigns the certificate alias, `wibble_ca`, to identify the new entry in the trust store.



Tip

A convenient way of viewing and manipulating Java keystore files is to use the free [Portecle](http://sourceforge.net/projects/portecle/)⁹ tool, which implements a visual frontend to the `keytool` utility.

3. Add the client's trusted CA to the list of trusted certificates in the STS.

The demonstration STS checks received certificates using its `CertificateVerifier` class. This class does *not* automatically trust all of the certificates present in the STS trust store, `stsstore.jks`, however. In order to permit authentication using the client's trusted CA, you must explicitly add the certificate alias to the certificate verifier's list of trusted certificates.

Edit the `beans.xml` file located in the following directory:

```
CxfInstallDir/samples/sts_issue_operation/src/main/webapp/WEB-INF
```

Look for the bean with the ID, `certificateVerifierConfig`, and add the `wibble_ca` certificate alias to the list of trusted aliases, as shown in the following fragment:

```
<beans ...>
  ...
  <bean id="certificateVerifierConfig"
    class="demo.sts.provider.cert.CertificateVerifierConfig">
    <property name="storePath" value="/stsstore.jks"/>
    <property name="storePwd" value="stsspass"/>
    <!-- if false exception for self-signed cert will be thrown -->
    <property name="verifySelfSignedCert" value="true"/>
    <property name="trustCertAliases">
      <list>
        <value>myclientkey</value>
        <value>wibble_ca</value>
      </list>
    </property>
    <property name="keySignAlias" value="mystskey"/>
    <property name="keySignPwd" value="stskpass"/>
  </bean>
</beans>
```

⁹ <http://sourceforge.net/projects/portecle/>

4. Install the STS certificate in the client.

Create a new `sts` directory and a new `sts/certs` directory under the `samples/wsd1_first_https` directory. Copy the `stsstore.jks` file from this directory:

```
CxfInstallDir/samples/sts_issue_operation/src/main/resources/
```

To this directory:

```
CxfInstallDir/samples/wsd1_first_https/sts/certs/
```

5. Install the STS WSDL in the client.

Create a new `sts/wsd1` directory under the `samples/wsd1_first_https` directory. Copy *all* of the files (`.wsdl` files and `.xsd` files) from this directory:

```
CxfInstallDir/samples/sts_issue_operation/src/main/webapp/WEB-INF/wsd1/
```

To this directory:

```
CxfInstallDir/samples/wsd1_first_https/sts/wsd1/
```

6. Customize the SOAP address in the client copy of the STS WSDL.

Edit the `ws-trust-1.4-service.wsdl` file from the `wsd1_first_https/sts/wsd1/` directory. Scroll down to the end of this file, where the `wsdl:service` element is defined. Change the value of the `location` attribute in the `soap:address` element as shown in the following fragment:

```
<wsdl:definitions ... >
  ...
  <wsdl:service name="SecurityTokenServiceProvider">
    <wsdl:port binding="tns:SecurityTokenServiceSOAP" name="SecurityTokenServiceSOAP">
      <soap:address location="https://localhost:8181/sts/SecurityTokenService"/>
    </wsdl:port>
  </wsdl:service>
</wsdl:definitions>
```

The new address, `https://localhost:8181/sts/SecurityTokenService`, is consistent with the Jetty container configuration, as specified in step 1.

**Note**

Don't forget to specify the secure protocol scheme, `https`, instead of `http` in this address.

7. To build and run the STS server using Maven, open a new command prompt, change directory to `samples/sts_issue_operation/`, and enter the following Maven command:

```
mvn jetty:run
```

The first time that you run this command, Maven will build the project before running STS in the Jetty container.

Define the Security Policy

Overview

To enable single sign-on for the client-server connection, you need to add a suitable security policy to the `hello_world` WSDL contract. The policy used in this example follows the general pattern of an authentication policy and has two major parts:

- A *TransportBinding* policy—this policy specifies that the basic protective features (such as message encryption) are provided by the transport layer, that is SSL/TLS. For more details, see [Transport Layer Message Protection on page 82](#).
- An *IssuedToken* policy—this policy enables the single sign-on scenario, involving the security token service (STS).

IssuedToken policy

The IssuedToken policy is a special case of an authentication token policy. Instead of supplying an authentication token directly, the client is required to call out to the STS, to obtain an authentication token (usually a SAML token). The presence of the IssuedToken policy in the WSDL contract *automatically* triggers the client to implement single sign-on semantics, where the client requests a remote STS to issue a token, which then gets embedded in the outgoing request to the server.

For details of how to specify an IssuedToken policy, see ["Defining an IssuedToken Policy" on page 173](#).

Sample security policy

[Example 8.1 on page 152](#) shows the security policy for single sign-on, which is applied to the client-server connection.

Example 8.1. Sample Security Policy for Single Sign-On

```
<wsdl:definitions ... >
  ...
  <wsp:Policy wsu:Id="STS_SAML_Token_policy"
    xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy"
    xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-
utility-1.0.xsd">
    <sp:TransportBinding xmlns:sp="http://docs.oasis-open.org/ws-sx/ws-securitypo
licy/200702">
      <wsp:Policy>
        <sp:TransportToken>
          <wsp:Policy>
            <sp:HttpsToken/>
```



```

        </wsp:Policy>
      </sp:TransportToken>
    <sp:Layout>
      <wsp:Policy>
        <sp:Strict/>
      </wsp:Policy>
    </sp:Layout>
    <sp:AlgorithmSuite>
      <wsp:Policy>
        <sp:Basic128/>
      </wsp:Policy>
    </sp:AlgorithmSuite>
  </wsp:Policy>
</sp:TransportBinding>
<sp:SignedSupportingTokens xmlns:sp="http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702">
  <wsp:Policy>
    <!-- Put IssuedToken element in here -->
    <sp:IssuedToken
      sp:IncludeToken="http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702/IncludeToken/AlwaysToRecipient">
      <sp:RequestSecurityTokenTemplate>
        <trust:TokenType
          xmlns:trust="http://docs.oasis-open.org/ws-sx/ws-trust/200512"
          >urn:oasis:names:tc:SAML:2.0:assertion</trust:TokenType>
        <!-- The demo STS signs the SAML token by default -->
      </sp:RequestSecurityTokenTemplate>
    </wsp:Policy>
    <!-- No extra policies needed in this demo. -->
  </wsp:Policy>
</sp:IssuedToken>
</wsp:Policy>
</sp:SignedSupportingTokens>
</wsp:Policy>
</wsdl:definitions>

```

Steps to add the security policy

Perform the following steps to add the single sign-on security policy to the `hello_world` WSDL contract:

1. Edit the `hello_world.wsdl` file from the `wsdl_first_https/wsdl/` directory. Add the single sign-on policy shown in [Example 8.1 on page 152](#) as a child of the `wsdl:definitions` element.
2. Continue editing the `hello_world.wsdl` file, in order to add a policy reference to the WSDL port. Search for the `SOAPService wsdl:service` element and then add the `wsp:PolicyReference` element as a child of the `wsdl:port` element, as shown in the following WSDL fragment:

```

<wsdl:definitions ... >
  ...
  <wsdl:service name="SOAPService">
    <wsdl:port binding="tns:Greeter_SOAPBinding" name="SoapPort">
      <wsp:PolicyReference xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy"
        URI="#STS_SAML_Token_policy"/>
      <soap:address location="https://localhost:9001/SoapContext/SoapPort"/>
    </wsdl:port>
  </wsdl:service>
  ...
</wsdl:definitions>

```

By inserting the `wsp:PolicyReference` element at this point, you are associating the WSDL port with the security policy referenced by the URI attribute value, `#STS_SAML_Token_policy`, (which matches the `wsu:Id` attribute of the single sign-on security policy).

3. The server requires a separate copy of the WSDL file, which *omits* the `IssuedToken` policy. Copy `hello_world.wsdl` to `hello_world_server.wsdl` (in the same directory). Edit the new `hello_world_server.wsdl` file and delete the `sp:SignedSupportingTokens` element from the policy, so that the content of the `hello_world_server.wsdl` file now has the following outline:

```

<wsdl:definitions ... >
  ...
  <wsp:Policy wsu:Id="STS_SAML_Token_policy"
    xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy"
    xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-
utility-1.0.xsd">
    <sp:TransportBinding xmlns:sp="http://docs.oasis-open.org/ws-sx/ws-securitypo
licy/200702">
      ...
    </sp:TransportBinding>
    <!-- sp:SignedSupportingTokens element is omitted in server copy of the WSDL -->
  </wsp:Policy>
</wsdl:definitions>

```



Note

If you completely omit the `wsp:Policy` element from the server's copy of the WSDL file, this would implicitly disable the auto-installation of the WSS4J interceptors. When you run the demonstration, the server would be unable to parse the security header and would therefore return a *mustUnderstand* fault.

Configure the Client-STS Connection

Overview

In order to retrieve a token from the STS, the client creates a connection to the STS, which must be configured in the same way as a regular client proxy. The code for instantiating this client proxy and invoking the STS is implemented by the following special class:

```
org.apache.cxf.ws.security.trust.STSClient
```

A client that supports single sign-on must explicitly create an instance of the `STSClient` class. When the client parses a WS-SecurityPolicy containing an IssuedToken policy, the Apache CXF runtime automatically looks for an `STSClient` instance to obtain a token from the STS. In other words, the IssuedToken policy and the `STSClient` object work hand-in-hand: the presence of an IssuedToken policy automatically triggers the `STSClient` object to interact with the STS.

Adding properties to a client proxy using XML

An interesting point about the sample client code is that it illustrates how you can use XML to configure a client proxy that has already been created in Java. In other words, this example answers the question: what do you do, when the client proxy is created in Java, but you want to specify some of its properties in XML?

First of all, consider the typical approach for instantiating a client proxy in Java, using the generated Greeter stub code, as follows:

```
// Java
import org.apache.hello_world_soap_http.Greeter;
import org.apache.hello_world_soap_http.SOAPService;
...
// Instantiate 'Greeter' client proxy
SOAPService ss = new SOAPService(wsdlURL, SERVICE_NAME);
Greeter port = ss.getPort(PORT_NAME, Greeter.class);
```

Now, in the XML configuration, you cannot use the `jaxws:client` element in the normal way to instantiate and configure the client proxy, because the Greeter client proxy already exists. It turns out, however, that the `jaxws:client` element supports a special syntax that enables you to inject properties into an existing instance, as shown in the following XML fragment:

```
<beans ...>
  ...
  <jaxws:client name="{http://apache.org/hello_world_soap_http}SoapPort"
               createdFromAPI="true">
    <!-- Set jaxws:properties, and so on -->
    ...
  </jaxws:client>
```

```
...
</beans>
```

The special syntax for modifying an existing client proxy uses the following attributes:

name

The full `QName` of the WSDL port associated with the existing client proxy (or possibly proxies).

createdFromAPI

When `true`, indicates that the client proxy was already created in Java code, and that this `jaxws:client` element should only be used to inject properties into the existing client proxy instance (or instances).

Steps to configure the STSClient

Perform the following steps to configure the STSClient:

1. Specify the `ws-security.sts.client` property on the client proxy. This property is used to reference an `org.apache.cxf.ws.security.trust.STSClient` instance, which is responsible for connecting to the STS. This property *must* be set, if the effective security policy contains an `IssuedToken` policy.

Edit the `wibbleClient.xml` file from the `wsdl_first_https/src/demo/hw_https/client` directory. Add the following `jaxws:client` element as a child of the `beans` element:

```
<beans ...>
...
<jaxws:client name="{http://apache.org/hello_world_soap_http}SoapPort" created
FromAPI="true">
  <jaxws:properties>
    <entry key="ws-security.sts.client" value-ref="default.sts-client" />
  </jaxws:properties>
</jaxws:client>
...
</beans>
```

2. Create the `STSClient` bean as follows. Continue editing the `wibbleClient.xml` file. Add the following `STSClient` bean definition to the XML file as shown:

```
<beans ...>
...
<bean name="default.sts-client"
class="org.apache.cxf.ws.security.trust.STSClient">
  <constructor-arg ref="cxf"/>
  <property name="wsdlLocation" value="sts/wsdl/ws-trust-1.4-service.wsdl"/>
  <property name="serviceName"
value="{http://docs.oasis-open.org/ws-sx/ws-trust/200512/wsdl}SecurityTokenService
Provider"/>
```

```

    <property name="endpointName"
      value="{http://docs.oasis-open.org/ws-sx/ws-trust/200512/wsd1}SecurityTokenSer
viceSOAP"/>
  </bean>
  ...
</beans>

```

Notice how the STSClient constructor requires a reference to the root Bus object (identified by the string, cxf, in the constructor-arg element) and the wsdlLocation attribute points to the the client's copy of the STS WSDL contract.

3. Secure the client-STS connection with SSL/TLS. Continue editing the wibbleClient.xml file. Add the following http:conduit element as a child of the beans element:

```

<beans ...>
  ...
  <http:conduit name="{http://docs.oasis-open.org/ws-sx/ws-trust/200512/wsd1}SecurityTokenSer
viceSOAP.http-conduit">
    <http:tlsClientParameters disableCNCheck="true">
      <sec:trustManagers>
        <sec:keyStore type="JKS" password="stsspass"
          file="sts/certs/stsstore.jks"/>
      </sec:trustManagers>
      <sec:keyManagers keyPassword="password">
        <sec:keyStore type="JKS" password="password"
          file="certs/wibble.jks"/>
      </sec:keyManagers>
      <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>

```

Notice how the STSClient trust store is configured to use the sts/certs/stsstore.jks keystore file, enabling the STSClient to authenticate the remote STS.

The name attribute of http:conduit follows the format, *WSDLPortQName*.http-conduit. Because *WSDLPortQName* matches the name of the STS WSDL port, these settings are automatically applied to the client proxy for the client-STS connection. For more details about the SSL/TLS security settings, see ["Security for HTTP-Compatible Bindings" on page 13](#).

4. Enable policy support and logging as follows. Continue editing the wibbleClient.xml file. Add the following cxf:bus element as a child of the beans element:

```

<beans ...>
  ...
  <cxf:bus xmlns:cxf="http://cxf.apache.org/core">
    <cxf:features>
      <p:policies xmlns:p="http://cxf.apache.org/policy"/>
      <cxf:logging/>
    </cxf:features>
  </cxf:bus>
  ...
</beans>

```



Note

It is essential to include the `<p:policies>` feature in the client's XML configuration. Otherwise, the policies in the WSDL file *would have no effect whatsoever*.

5. Add the requisite XML schema locations. Continue editing the `wibbleClient.xml` file. To support the `jaxws`, `cxf`, and `p` namespace prefixes, add the highlighted schema locations and define the `jaxws` namespace prefix, as follows:

```

<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:sec="http://cxf.apache.org/configuration/security"
  xmlns:http="http://cxf.apache.org/transports/http/configuration"
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  xsi:schemaLocation="
    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://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://cxf.apache.org/core http://cxf.apache.org/schemas/core.xsd
    http://cxf.apache.org/policy http://cxf.apache.org/schemas/policy.xsd">
  ...

```

6. Add the requisite Maven dependencies. In order to use `WS-SecurityPolicy`, you need to ensure that the requisite JARs are included on the classpath. For the Maven build system, this requires you to include additional dependencies in the POM file. Edit the `wsdl_first_https/pom.xml` file and add dependencies on the `cxf-rt-ws-security` artifact and on the `cxf-rt-ws-policy` artifact as highlighted in the following fragment:

```
<project ...>
  ...
  <dependencies>
    ...
    <dependency>
      <groupId>org.apache.cxf</groupId>
      <artifactId>cxf-rt-ws-security</artifactId>
      <version>2.4.0-fuse-00-27</version>
    </dependency>
    <dependency>
      <groupId>org.apache.cxf</groupId>
      <artifactId>cxf-rt-ws-policy</artifactId>
      <version>2.4.0-fuse-00-27</version>
    </dependency>
  </dependencies>
</project>
```

Configure the Server-Side Interceptor

Overview

This section describes how to customize the server configuration, so that it is capable of checking and confirming the signature that appears in the SAML tokens received from the client.

Auto-installation of WSS4J interceptors

The underlying implementation of WS-Security and WS-Trust is provided by a set of WSS4J interceptors, based on the open source [Apache WSS4J](http://ws.apache.org/wss4j/)¹⁰ security toolkit. These interceptors are automatically installed into a server endpoint, if the endpoint is associated with a WS-SecurityPolicy policy in the server's WSDL file. In the absence of WSDL policies, it is also possible to install the WSS4J interceptors explicitly—see the [WS-Security](http://cxf.apache.org/docs/ws-security.html)¹¹ page.

Confirming the SAML token signature

In this example, the SAML token issued by the STS is signed by the STS certificate. The issued SAML token also declares the subject confirmation method to be bearer, which implies that the server should trust the bearer of the SAML token, only if the token's signature is confirmed.

Hence, the server must be configured so that it can confirm incoming signatures. The details of this type of configuration are given in [Providing Encryption Keys and Signing Keys on page 104](#) and, essentially, it consists of setting the `ws-security.encryption.properties` property in the relevant `jaxws:endpoint` element. Confusingly, although this property is named *encryption*, it is also used for confirming signatures. The `ws-security.encryption.properties` property accesses the STS certificate, which is used to confirm the token's signature.

Limitations of the server implementation

Because the current example is intended mainly to illustrate how to integrate a client with WS-Trust, *the server side has not been fully implemented*. Although the server (as configured here) is capable of understanding and processing security headers from the client, the server is *not* properly configured to enforce any authentication or authorization policies. In particular, if a client fails to send an authentication token, the server will happily process the request in any case.

A properly implemented server would be augmented by policies that require a particular kind of authentication token and the server would check the value of the received token, as appropriate.

¹⁰ <http://ws.apache.org/wss4j/>

¹¹ <http://cxf.apache.org/docs/ws-security.html>

Steps to configure the server-side interceptor

Perform the following steps to configure the server-side interceptor:

1. Comment out the Java code for instantiating the Web service endpoint (in this example, it is more convenient to instantiate the endpoint in XML, because it enables you to specify all of the endpoint's properties in one place).

Edit the `Server.java` file from the `wsdl_first_https/src/demo/hw_https/server` directory. Look for the lines of Java code that instantiate the Web service endpoint (highlighted below) and enclose them between `/*` and `*/`, so that the lines are commented out as shown.

```
package demo.hw_https.server;
...
public class Server {

    protected Server() throws Exception {
        System.out.println("Starting Server");

        SpringBusFactory bf = new SpringBusFactory();
        URL busFile = Server.class.getResource("CherryServer.xml");
        Bus bus = bf.createBus(busFile.toString());
        bf.setDefaultBus(bus);

        /*
            Object implementor = new GreeterImpl();
            String address = "https://localhost:9001/SoapContext/SoapPort";
            Endpoint.publish(address, implementor);
        */
    }
    ...
}
```

2. Create the Web service endpoint in XML. Edit the `CherryServer.xml` file from the `wsdl_first_https/src/demo/hw_https/server` directory. Add the following `jaxws:endpoint` element as a child of the `beans` element to instantiate the Web service endpoint.

```
<beans ...>
...
<jaxws:endpoint id="server"
    endpointName="s:SoapPort"
    serviceName="s:SOAPService"
    implementor="demo.hw_https.server.GreeterImpl"
    address="https://localhost:9001/SoapContext/SoapPort"
    wsdlLocation="wsdl/hello_world_server.wsdl"
    xmlns:s="http://apache.org/hello_world_soap_http" >
</jaxws:properties>
```

```

        <entry key="ws-security.encryption.properties" value="sts/sts.properties" />
    </jaxws:properties>
</jaxws:endpoint>
...
</beans>

```

Notice how the `wsdlLocation` attribute points at the `hello_world_server.wsdl` file, which is the copy of the WSDL contract that excludes the `IssuedToken` policy. The `ws-security.encryption.properties` property points at the file, `sts/sts.properties`, which will be defined in a later step.

3. Enable policy support and logging as follows. Continue editing the `CherryServer.xml` file. Add the following `cxf:bus` element as a child of the `beans` element:

```

<beans ...>
    ...
    <cxf:bus xmlns:cxf="http://cxf.apache.org/core">
        <cxf:features>
            <p:polices xmlns:p="http://cxf.apache.org/policy"/>
            <cxf:logging/>
        </cxf:features>
    </cxf:bus>
    ...
</beans>

```

4. Add the requisite XML schema locations. Continue editing the `CherryServer.xml` file. To support the `jaxws` and `cxf` namespace prefixes, add the highlighted schema locations and define the `jaxws` namespace prefix, as follows:

```

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

```

- Now define the `sts.properties` file, which specifies the WSS4J properties for accessing the STS certificate. In the `wsdl_first_https/sts` directory, use your favorite text editor to create the file, `sts.properties`, containing the following property settings:

```
org.apache.ws.security.crypto.provider=org.apache.ws.security.components.crypto.Merlin
org.apache.ws.security.crypto.merlin.keystore.type=JKS
org.apache.ws.security.crypto.merlin.keystore.file=sts/certs/stsstore.jks
org.apache.ws.security.crypto.merlin.keystore.password=stsspass
org.apache.ws.security.crypto.merlin.keystore.alias=mystskey
```

For more details on these WSS4J properties, see [Providing Encryption Keys and Signing Keys on page 104](#).

- Add Maven instructions to copy the `sts.properties` file into the `target/classes/sts` directory (so that the `sts.properties` file gets included in the WAR package). Edit the `pom.xml` file from the `wsdl_first_https/` directory and search for the `copyxmlfiles` target of the Maven antrun plug-in. Under configuration tasks, add the highlighted lines as shown in the following fragment:

```
<project ...>
  ...
  <build>
    ...
    <plugins>
      <plugin>
        <artifactId>maven-antrun-plugin</artifactId>
        <executions>
          <execution>
            <id>copyxmlfiles</id>
            ...
            <configuration>
              <tasks>
                <copy file="${basedir}/src/demo/hw_https/server/CherryServ
er.xml" todir="${basedir}/target/classes/demo/hw_https/server" />
                <copy file="${basedir}/src/demo/hw_https/client/WibbleCli
ent.xml" todir="${basedir}/target/classes/demo/hw_https/client" />
                <copy file="${basedir}/src/demo/hw_https/client/InsecureCli
ent.xml" todir="${basedir}/target/classes/demo/hw_https/client" />
                <copy todir="${basedir}/target/classes/certs">
                  <fileset dir="${basedir}/certs" />
                </copy>
                <copy todir="${basedir}/target/classes/sts">
                  <fileset dir="${basedir}/sts" />
                </copy>
              </tasks>
            </configuration>
          </execution>
        </executions>
      </plugin>
    </plugins>
```

```
</build>  
...  
</project>
```



Note

This step is only necessary, because this `wsdl_first_https` Maven project is set up in a slightly unconventional way. Normally, in a Maven project, you put all of your resource files under `src/main/resources/`, which Maven automatically copies into the target package.

Build and Run the Demonstration

Steps to run the demonstration

Perform the following steps to build and run the demonstration:

1. If you have not already done so, start up the STS as follows. Open a new command prompt, change directory to `samples/sts_issue_operation`, and enter the following Maven command:

```
mvn jetty:run
```

2. Build and start up the WS server. Open a new command prompt, change directory to `samples/wsd1_first_https`, and enter the following Maven command:

```
mvn -Pserver
```



Note

The server is programmed to time out after five minutes.

3. Build and run the WS client. Open a new command prompt, change directory to `samples/wsd1_first_https`, and enter the following Maven command:

```
mvn -Psecure.client
```

Sample Message Exchanges

Overview

When the preceding demonstration runs successfully, you can see the following message exchanges logged to the console window (assuming you enabled logging by including the `cxfr:logging` feature in the client and server configuration).

Outbound message to STS

To obtain a SAML security token issued by the security token service, the client sends the following RequestSecurityToken (RST) message to the security token service:

```
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <soap:Body>
    <wst:RequestSecurityToken xmlns:wst="http://docs.oasis-open.org/ws-sx/ws-
trust/200512">
      <wst:SecondaryParameters>
        <trust:TokenType xmlns:trust="http://docs.oasis-open.org/ws-sx/ws-
trust/200512"
          >urn:oasis:names:tc:SAML:2.0:assertion</trust:TokenType>
        </wst:SecondaryParameters>
        <wst:RequestType>http://docs.oasis-open.org/ws-sx/ws-trust/200512/Issue</wst:Re
questType>
        <wst:KeyType>http://docs.oasis-open.org/ws-sx/ws-trust/200512/Symmet
ricKey</wst:KeyType>
        <wst:KeySize>256</wst:KeySize>
        <wst:Entropy>
          <wst:BinarySecret Type="http://docs.oasis-open.org/ws-sx/ws-
trust/200512/Nonce"
            >IwPnJynT6rioPQRuo0q2vj+lEM/xvDy+ZXkRR7dxG1k=</wst:BinarySecret>
          </wst:Entropy>
          <wst:ComputedKeyAlgorithm>http://docs.oasis-open.org/ws-sx/ws-
trust/200512/CK/PSHA1</wst:ComputedKeyAlgorithm>
        </wst:RequestSecurityToken>
      </soap:Body>
    </soap:Envelope>
```

Inbound message from STS

The security token service sends back the following RequestSecurityTokenResponse (RSTR) message, containing a signed SAML token, `saml2:Assertion`, back to the client:

```
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <soap:Body>
```

```

    <RequestSecurityTokenResponseCollection
      xmlns="http://docs.oasis-open.org/ws-sx/ws-trust/200512"
      xmlns:ns2="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-
utility-1.0.xsd"
      xmlns:ns3="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-
secect-1.0.xsd"
      xmlns:ns4="http://www.w3.org/2005/08/addressing">
      <RequestSecurityTokenResponse>
        <TokenType>http://docs.oasis-open.org/wss/oasis-wss-saml-token-profile-
1.1#SAMLV2.0</TokenType>
        <RequestedSecurityToken>
          <saml2:Assertion xmlns:saml2="urn:oasis:names:tc:SAML:2.0:assertion"
            ID="_181835fb981efecaf71d80ecd5fc3c74"
            IssueInstant="2011-05-09T09:36:37.359Z" Version="2.0">
            <saml2:Issuer>http://www.sopera.de/SAML2</saml2:Issuer>
            <saml2:Subject>
              <saml2:NameID
                Format="urn:oasis:names:tc:SAML:1.1:nameid-format:transient"/>

              <saml2:SubjectConfirmation
                Method="urn:oasis:names:tc:SAML:2.0:cm:bearer"/>
            </saml2:Subject>
            <saml2:Conditions NotBefore="2011-05-09T08:36:37.359Z"
              NotOnOrAfter="2011-05-09T10:36:37.359Z"/>
            <saml2:AuthnStatement AuthnInstant="2011-05-09T09:36:37.515Z">
              <saml2:AuthnContext>
                <saml2:AuthnContextClassRef>ac:classes:X509</saml2:AuthnCon
textClassRef>

                </saml2:AuthnContext>
              </saml2:AuthnStatement>
            <Signature:Signature xmlns:Signature="http://www.w3.org/2000/09/xm
ldsig#"

              xmlns="http://www.w3.org/2000/09/xmldsig#">
              <SignedInfo>
                <CanonicalizationMethod
                  Algorithm="http://www.w3.org/2001/10/xml-exc-c14n#" />
                <SignatureMethod
                  Algorithm="http://www.w3.org/2000/09/xmldsig#rsa-sha1"/>

                <Reference URI="#_181835fb981efecaf71d80ecd5fc3c74">
                  <Transforms>
                    <Transform
                      Algorithm="http://www.w3.org/2000/09/xmldsig#en
veloped-signature"/>

                    <Transform
                      Algorithm="http://www.w3.org/2001/10/xml-exc-
c14n#" />

                    </Transforms>
                  <DigestMethod Algorithm="http://www.w3.org/2000/09/xm

```

```

ldsig#sha1"/>
                                <DigestValue>Gpzf8TjPATPsQDAm2ojNdEpht1A=</DigestValue>

                                </Reference>
                                </SignedInfo>
                                <SignatureValue>jsbIP1Z25q4Qedn60Sid4QcV4cs6+lgwB+jDiImwM
MEoyzp1BjwQWB+1SibHfa9rtmmTszLdmeTqxSXiAy2CeVZcIDk1UAfySAhDrrmR5N6lJMJsQgU4o1ysLsZMK
wtR2FL+eya7hJ9e4UtQVH1K0a7Cx1rvl4Dr8u8FuN5Myg=</SignatureValue>
                                <KeyInfo>
                                <X509Data>
                                <X509Subject
Name>1.2.840.113549.1.9.1=#160b737473407374732e636f6d,CN=www.sts.com,OU=IT
                                Department,O=Sample STS -- NOT FOR
                                PRODUCTION,L=Baltimore,ST=Maryland,C=US</X509Sub
jectName>
                                <X509Certificate>MIID5jCCA0+gAwIBAgIJAPahVdM2UPibMA0GC
SgGSib3DQEBBQUAMIGpMQswCQYDVQQGEwJVUzER
                                MA8GA1UECBMITWFyewxhbmQxEjAQBgNVBAcTCUJhbHRpbw9yZTEp
MCCGA1UEChMU2FtcGx1IFNU
                                UyAtLSBOT1QgRk9SIFBST0RVQ1RJT04xFjAUBgNVBASTDULUIER
lcGFydG1lbnQxFDASBgNVBAMT
                                C3d3dy5zdHMuY29tMRowGAYJKoZIhvcNAQkBFgtzdHNAc3RzLmN
vbTAeFw0xMTAyMDkxODM4MTNa
                                Fw0yMTAyMDYxODM4MTNaMIGpMQswCQYDVQQGEwJVUzERMA8GA1UECB
MITWFyewxhbmQxEjAQBgNV
                                BAcTCUJhbHRpbw9yZTEpMCCGA1UEChMU2FtcGx1IFNUUyAtLS
BOT1QgRk9SIFBST0RVQ1RJT04x
                                FjAUBgNVBASTDULUIERlcGFydG1lbnQxFDASBgNVBAMTC3d3dy5zdH
MuY29tMRowGAYJKoZIhvcN
                                AQkBFgtzdHNAc3RzLmNvbTCBnzANBGMkqhkig9w0BAQEFAAOB
jQAwGyKCGYEAo+f8gs4WcteLdSPW
                                Pm8+ciyEz7zVmA7kcCGFQQv100smxRViWJ1x+yniT5Uu86Ur
AQjxRJyANBomQrirfE7KPrnCM6iV
                                OsGDEntuIZAf7DFPnrv5p++jAZQuR3vm4ZHX
FOFTXmI+/FD5AqLfNi17xiTxZCDYyDdD39CNFTrB
                                2PkCAwEAAa0CARIwggEOMB0GA1UdDgQWBBrA0A38holQIbJM
FW7m5ZSw+iVDHDCB3gYDVR0jBIHW
                                MIHTgBBrA0A38holQIbJMFW7m5ZSw+iVDHKGBr6SBrDCBqTEL
MAKGA1UEBhMCMVVMxETAPBgNVBAGT
                                CE1hcnlsYW5kMRIwEAYDVQQHEw1CYWx0aW1vcmlKTAnBgNVBAo
TIFNhbXBsZSBTVFMgLS0gTk9U
                                IEZPUiBQUk9EVUNUSU90MRYwFAYDVQQLEw1JVCBEZXhcnRtZW50MR
QwEgYDVQQDEwt3d3cuc3Rz
                                LmNvbTEaMBGCSqGSib3DQE
JARYLc3RzQHN0cy5jb22CCQD2oVXTNlD4mzAMBgNVHRMEBTADAQH/
                                MA0GCSqGSib3DQEBBQUAA4GBACp9yK1I9r++pyFT0yr
caV1m1Sub6urJH+GxQLBaTnTsaPLuzq2g
                                IsJHpwk5XggB+IDe69iKKeb74Vt8a0e5usIwVASgi9ckqCwd
fTqYu6KG9BlezqHZdExnIG2v/cD/

```



```

tificate>
    3NkKr70/a7DjlbE6FZ4G1nrOfVJkjmeAa6txtYm1Dm/f</X509Cer
    </X509Data>
    </KeyInfo>
    </Signature:Signature>
  </saml2:Assertion>
</RequestedSecurityToken>
<RequestedAttachedReference>
  <ns3:SecurityTokenReference>
    <ns3:KeyIdentifier>_181835fb981efecaf71d80ecd5fc3c74</ns3:KeyIden
tifier>
    </ns3:SecurityTokenReference>
  </RequestedAttachedReference>
  <RequestedUnattachedReference>
    <ns3:SecurityTokenReference>
      <ns3:KeyIdentifier>_181835fb981efecaf71d80ecd5fc3c74</ns3:KeyIden
tifier>
    </ns3:SecurityTokenReference>
  </RequestedUnattachedReference>
</RequestSecurityTokenResponse>
</RequestSecurityTokenResponseCollection>
</soap:Body>
</soap:Envelope>

```

Outbound message to server

The client now embeds the signed SAML token, `saml2:Assertion`, in the WS-Security header, `wsse:Security`, when it invokes the `greetMe` operation on the server:

```

<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <soap:Header>
    <wsse:Security
      xmlns:wsse="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-
secext-1.0.xsd"
      xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-
utility-1.0.xsd"
      soap:mustUnderstand="1">
        <saml2:Assertion xmlns:saml2="urn:oasis:names:tc:SAML:2.0:assertion"
          ID="_181835fb981efecaf71d80ecd5fc3c74" IssueInstant="2011-05-09T09:36:37.359Z"
          Version="2.0">
          <saml2:Issuer>http://www.sopera.de/SAML2</saml2:Issuer>
          <saml2:Subject>
            <saml2:NameID Format="urn:oasis:names:tc:SAML:1.1:nameid-format:transi
ent"/>
            <saml2:SubjectConfirmation Method="urn:oasis:names:tc:SAML:2.0:cm:bear
er"/>
          </saml2:Subject>

```

```

        <saml2:Conditions NotBefore="2011-05-09T08:36:37.359Z"
            NotOnOrAfter="2011-05-09T10:36:37.359Z"/>
        <saml2:AuthnStatement AuthnInstant="2011-05-09T09:36:37.515Z">
            <saml2:AuthnContext>
                <saml2:AuthnContextClassRef>ac:classes:X509</saml2:AuthnContext
ClassRef>
                </saml2:AuthnContext>
            </saml2:AuthnStatement>
            <Signature xmlns="http://www.w3.org/2000/09/xmldsig#"
                xmlns:Signature="http://www.w3.org/2000/09/xmldsig#">
                <SignedInfo>
                    <CanonicalizationMethod Algorithm="http://www.w3.org/2001/10/xml-
exc-c14n#"/>
                    <SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#rsa-
sha1"/>
                    <Reference URI="#_181835fb981efecaf71d80ecd5fc3c74">
                        <Transforms>
                            <Transform
                                Algorithm="http://www.w3.org/2000/09/xmldsig#enveloped-
signature"/>
                            <Transform Algorithm="http://www.w3.org/2001/10/xml-exc-
c14n#"/>
                        </Transforms>
                    <DigestMethod Algorithm="http://www.w3.org/2000/09/xmldsig#sha1"/>
                        <DigestValue>Gpzf8TjPATPsQDAm2ojNdEpht1A=</DigestValue>
                    </Reference>
                </SignedInfo>
                <SignatureValue>jsbIP1Z25q4Qedn60Sid4QcV4cs6+lgwB+jDiImwM
MEoyzp1BjWQWB+1SIbHfa9rtmmTszLdmeTqxSxiAy2CeVZcIDk1UAfySAhDrrmR5N6lJMJqsQgU4o1ysLsZMK
wtR2FL+eya7hJ9e4UtQVH1K0a7Cx1rvl4Dr8u8FuN5Myg=</SignatureValue>
                <KeyInfo>
                    <X509Data>
                        <X509Subject
Name>1.2.840.113549.1.9.1=#160b737473407374732e636f6d,CN=www.sts.com,OU=IT
                        Department,O=Sample STS -- NOT FOR
                        PRODUCTION,L=Baltimore,ST=Maryland,C=US</X509SubjectName>
                        <X509Certificate>MIID5jCCA0+gAwIBAgIJAPahVdM2UPibMA0GCSqGS
Ib3DQEBBQUAMIGpMQswCQYDVQQGEwJVUzER
                        MA8GA1UECBMITWFyZWxhbmQxEjAQBGNVBACTCUJhbHRpbW9yZTEpMC
cGA1UEChMgU2FtcGx1IFNU
                        UyAtLSB0T1QgRk9SIFBST0RVQ1RJT04xFjAUBGNVBASTDUlUIERlcGFy
dG1lbnQxFDASBgNVBAMT
                        C3d3dy5zdHMuY29tMRowGAYJKoZIhvcNAQkBFgtzdHNAc3RzLmNvbTAe
Fw0xMTAyMDkxODM4MTNa
                        Fw0yMTAyMDYxODM4MTNAMIGpMQswCQYDVQQGEwJVUzERMA8GA1UECBMITW
FyZWxhbmQxEjAQBGNV
                        BAcTCUJhbHRpbW9yZTEpMCCGA1UEChMgU2FtcGx1IFNUUyAtLSB0T1QgRk9SIF
BST0RVQ1RJT04x

```

```

RowGAYJKoZIhvcN          FjAUBgNVBASTDUIUERlcGFydG11bnQxFDASBgNVBAMTC3d3dy5zdHMuY29tM
CgYEAo+f8gs4wcteLdSPW    AQkBFgtzdHNAC3RzLmNvbTCBnzANBghkhiG9w0BAQEFAA0BjQAwgYk
BomQrirfE7KPrnCm6iV      Pm8+ciyEz7zVmA7kcCGFQQv100smxRviWJ1x+yniT5Uu86UrAQjxRJyAN
FOFTXmI+/FD5AqLfNi17xiTxZCDYyDdD39CNFTTrB    OsGDEntuIZAf7DFPnrV5p++jAZQuR3vm4ZHX
HDCB3gYDVR0jBIHW        2PkCAwEAAaOCARIwggEOMB0GA1UdDgQWBBrA0A38ho1QIbJMFw7m5ZSw+iVD
MCVVMxETAPBgNVBAGT      MIHTgBRA0A38ho1QIbJMFw7m5ZSw+iVDHKGBR6SBrDCBqTELMakGA1UEBh
sZSBTVFMgLS0gTk9U       CE1hcnlsYW5kMRIwEAYDVQQHEw1CYWx0aw1vcuXKTANBgNVBAoTIFNhbXB
DVQDEwt3d3cuc3Rz        IEZPUiBQOk9EVUNUSU90MRyWfAYDVQQLEw1JVCBEZXBhcncRtZW50MRQwEgY
BgNVHRMEBTADAQH/        LmNvbTEaMBGCSqGSIb3DQEJARYLc3RzQHN0cy5jb22CCQD2oVXTNld4mzAM
JH+GxQLBaTnTsaPLuzq2g   MA0GCSqGSIb3DQEBBQUAA4GBACp9yK1I9r++pyFT0yrcaV1m1Sub6ur
fTqYu6KG9BleZqHZdExnIG2v/cD/    IsJHpwk5XggB+IDe69iKKeB74Vt8a0e5usIWVASgi9ckqCwd
3NkKr70/a7DjlbE6FZ4G1nr0fVJkjmeAa6txtYm1Dm/f</X509Certificate>

    </X509Data>
    </KeyInfo>
  </Signature:Signature>
</saml2:Assertion>
</wsse:Security>
</soap:Header>
<soap:Body>
  <greetMe xmlns="http://apache.org/hello_world_soap_http/types">
    <requestType>JBLOGGS</requestType>
  </greetMe>
</soap:Body>
</soap:Envelope>

```

Inbound message from server

When the server receives the preceding SOAP request, the `soap:mustUnderstand="1"` attribute setting ensures that the server *must* process the security header. In addition, the presence of a signature in the SAML token means that the server *must* confirm the signature.

After successfully processing the security header, the server sends back the following reply to the client:

```

<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <soap:Header/>

```

```
<soap:Body>
  <greetMeResponse xmlns="http://apache.org/hello_world_soap_http/types">
    <responseType>Hello JBLOGGS</responseType>
  </greetMeResponse>
</soap:Body>
</soap:Envelope>
```

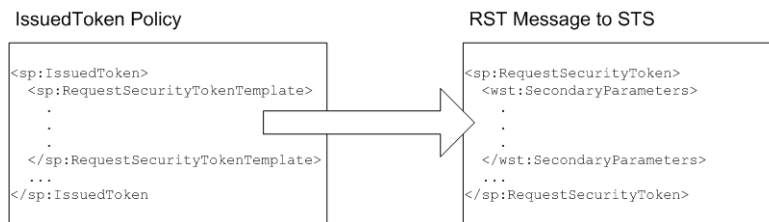
Defining an IssuedToken Policy

Overview

In general, an IssuedToken policy can appear in any context where a regular token can appear. When an `sp:IssuedToken` element appears in a policy, it indicates that the application must call out to the STS to obtain a token (as specified in the `sp:IssuedToken` element) and then use the token in the current context.

Fundamentally, an IssuedToken policy consists of two parts: one part is aimed at the client, specifying how the client must use the IssuedToken; and another part is aimed at the STS, specifying what type of token to issue and how the token should be constructed. The part that is aimed at the STS is put inside the special element, `sp:RequestSecurityTokenTemplate`. All of the children of this element will be copied directly into the body of the RequestSecurityToken (RST) message that is sent to the STS when the client asks the STS to issue a token, as shown in [Figure 8.7 on page 173](#).

Figure 8.7. Injecting Parameters into the Outgoing RequestSecurityToken Message



Namespaces

A typical IssuedToken policy is defined using elements from the following schema namespaces:

Table 8.1. XML Namespaces used with IssuedToken

Prefix	XML Namespace
wsp:	http://schemas.xmlsoap.org/ws/2004/09/policy
sp:	http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702
wst:	http://docs.oasis-open.org/ws-sx/ws-trust/200512

Sample policy

The following example shows a minimal IssuedToken policy, where the client requests the STS to issue a SAML 2.0 token (specified by the value of the `trust:TokenType` element). The issued token will be included in the client's request header (indicated by setting the `sp:IncludeToken` attribute to `AlwaysToRecipient`).

```
<sp:IssuedToken
  sp:IncludeToken="http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702/Include
Token/AlwaysToRecipient">
  <sp:RequestSecurityTokenTemplate>
    <trust:TokenType
      xmlns:trust="http://docs.oasis-open.org/ws-sx/ws-trust/200512"
      >urn:oasis:names:tc:SAML:2.0:assertion</trust:TokenType>
    </sp:RequestSecurityTokenTemplate>
  <wsp:Policy>
    <!-- No extra policies needed in this demo. -->
  </wsp:Policy>
</sp:IssuedToken>
```

In an example such as this, where the IssuedToken policy contains few settings, it is assumed that the STS is already configured with sensible default properties.

Syntax

The IssuedToken element is defined with the following syntax:

```
<sp:IssuedToken
  sp:IncludeToken="xs:anyURI"?
  xmlns:sp="..." ... >
  (
    <sp:Issuer>wsa:EndpointReferenceType</sp:Issuer> |
    <sp:IssuerName>xs:anyURI</sp:IssuerName>
  ) ?
  <wst:Claims Dialect="..."> ... </wst:Claims> ?
  <sp:RequestSecurityTokenTemplate TrustVersion="xs:anyURI"? >
    <wst:TokenType>...</wst:TokenType> ?
    <wsp:AppliesTo>...</wsp:AppliesTo> ?
    <!-- Many other WS-Trust elements allowed here -->
    ...
  </sp:RequestSecurityTokenTemplate>
  <wsp:Policy xmlns:wsp="...">
    (
      <sp:RequireDerivedKeys ... /> |
      <sp:RequireImpliedDerivedKeys ... /> |
      <sp:RequireExplicitDerivedKeys ... />
    ) ?
    <sp:RequireExternalReference ... /> ?
```

```

    <sp:RequireInternalReference ... /> ?
    ...
  </wsp:Policy>
  ...
</sp:IssuedToken>

```

XML elements

An IssuedToken policy is specified using the following XML elements:

sp:IssuedToken

The element containing the IssuedToken policy assertion. The `IncludeToken` attribute specifies whether the issued token is meant to be included in the security header of the client's request messages. The allowed values of this attribute are given in ["sp:IncludeToken attribute" on page 122](#).

On the client side, the presence of this assertion signals that the client should attempt to obtain a token by contacting a remote STS.

sp:IssuedToken/sp:Issuer

(Optional) Contains a reference to the issuer of the token, of `wsa:EndpointReferenceType` type.

There is no need to specify the issuer's endpoint reference using `sp:Issuer` in Fuse Services Framework, because the issuer endpoint (that is, the STS address) is taken directly from the STS WSDL file instead. See [Step 2 on page 156](#).

sp:IssuedToken/sp:IssuerName

(Optional) The name of the issuing service (that is, the STS), in the format of a URI.

sp:IssuedToken/wst:Claims

(Optional) Specifies the required claims that the issued token must contain in order to satisfy the IssuedToken policy assertion.

sp:IssuedToken/sp:RequestSecurityTokenTemplate

(Required) This element contains a list of WS-Trust policy assertions to be included in the outgoing RequestSecurityToken (RST) message that is sent to the STS. In other words, this element enables you to modify the Issue query that is sent to the STS to obtain the issued token. This element can contain any of the WS-Trust assertions that are valid children of the `sp:RequestSecurityToken` element, as specified by WS-Trust.

sp:IssuedToken/sp:RequestSecurityTokenTemplate/wst:TokenType

(Optional) The type of security token to issue, specified as a URI string. Although theoretically optional, this assertion is almost always specified. [Table 8.2 on page 176](#) shows the list of standard token type URIs for the following token types: SAML 1.1, SAML 2.0, UsernameToken, X.509v3 single certificate, X.509v1 single certificate, X.509 PKI certificate chain, X.509 PKCS7, and Kerberos.

Table 8.2. Token Type URIs

Token Type URI
http://docs.oasis-open.org/wss/oasis-wss-saml-token-profile-1.1#SAMLV1.1
http://docs.oasis-open.org/wss/oasis-wss-saml-token-profile-1.1#SAMLV2.0
http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-username-token-profile-1.0#UsernameToken
http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-x509-token-profile-1.0#X509v3
http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-x509-token-profile-1.0#X509v1
http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-x509-token-profile-1.0#X509PKIPathv1
http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-x509-token-profile-1.0#PKCS7
http://docs.oasisopen.org/wss/oasiswss-kerberos-tokenprofile-1.1#Kerberosv5APREQSHA1

Consult the documentation for your third-party STS to find out what token types it supports. An STS can also support custom token types not listed in the preceding table.

`sp:IssuedToken/sp:RequestSecurityTokenTemplate/wsp:AppliesTo`

(Optional) This `WS-PolicyAttachment` assertion can be specified as an alternative to or in addition to the `wst:TokenType` assertion (in the latter case, it takes precedence over the `wst:TokenType` assertion). It is used to specify the policy scope for which this token is required. In practice, this enables you to refer to a service or group of services for which this token is issued. The STS can then be configured to specify what kind of token to issue for that service (or services).

`sp:IssuedToken/sp:RequestSecurityTokenTemplate/wst:OtherElements`

(Optional) You can optionally include many other WS-Trust assertions in the RST template. The purpose of these assertions is to specify exactly what the content of the issued token should be and whether it is signed, encrypted, and so on. In practice, however, the details of the issued token are often configured in the STS, which makes it unnecessary to include all of these details in the RST template.

For a full list of of WS-Trust assertions that can be included in the RST template, see the [OASIS WS-Trust 1.4 Specification](#)¹².

`sp:IssuedToken/sp:Policy`

(Required) This element must be included in the `IssuedToken`, even if it is empty.

`sp:IssuedToken/sp:Policy/sp:RequireDerivedKeys`

(Optional) Only applicable when using WS-SecureConversation. The WS-SecureConversation specification enables you to establish a *security context* (analogous to a session), which is used for sending multiple secured messages to a given service. Normally, if you use the straightforward approach of authenticating every message sent to a particular service, the authentication adds a considerable overhead to secure communications. If you know in advance that a client will be sending multiple messages to a Web service,

¹² <http://docs.oasis-open.org/ws-sx/ws-trust/v1.4/ws-trust.html>

however, it makes sense to establish a security context between the client and the server, in order to cut the overheads of secure communication. This is the basic idea of WS-SecureConversation.

When a security context is established, the client and the server normally establish a shared secret. In order to prevent the shared secret being discovered, it is better to avoid using the shared secret directly and use it instead to generate the keys needed for encryption, signing, and so on—that is, to generate *derived keys*.

When the `sp:RequireDerivedKeys` policy assertion is included, the use of derived keys is enabled in WS-SecureConversation and both *implied derived keys* and *explicit derived keys* are allowed.



Note

Only one of `sp:RequireDerivedKeys`, `sp:RequireImpliedDerivedKeys`, or `sp:RequireExplicitDerivedKeys`, can be included in `sp:IssuedToken`.

`sp:IssuedToken/sp:Policy/sp:RequireImpliedDerivedKeys`

(Optional) When the `sp:RequireImpliedDerivedKeys` policy assertion is included, the use of derived keys is enabled in WS-SecureConversation, but only *implied derived keys* are allowed.

`sp:IssuedToken/sp:Policy/sp:RequireExplicitDerivedKeys`

(Optional) When the `sp:RequireExplicitDerivedKeys` policy assertion is included, the use of derived keys is enabled in WS-SecureConversation, but only *explicit derived keys* are allowed.

`sp:IssuedToken/sp:Policy/sp:RequireExternalReference`

(Optional) When included, requires that external references to the issued token must be enabled.

`sp:IssuedToken/sp:Policy/sp:RequireInternalReference`

(Optional) When included, requires that internal references to the issued token must be enabled, where an internal reference uses one of the elements, `wsse:Reference`, `wsse:KeyIdentifier`, or `wsse:Embedded`.

Creating an STSClient Instance

Overview

Whenever an IssuedToken policy is configured on a WSDL port, you must also configure the client to connect to an STS server to obtain a token. The code for connecting to the STS and obtaining a token is implemented by the following class:

```
org.apache.cxf.ws.security.trust.STSClient
```

The client must explicitly create an STSClient instance to manage the client-STS connection. You can do this in either of the following ways:

- *Direct configuration*—the client proxy is configured with the `ws-security.sts.client` property, which contains a reference to an STSClient instance.
- *Indirect configuration*—no change is made to the client proxy definition, but if the Apache CXF runtime finds an appropriately named STSClient bean in the bean registry, it will automatically inject that STSClient bean into the client proxy.

In addition to creating an STSClient instance, it is usually also necessary to enable SSL/TLS security on the STS proxy, as described in ["Configure the Client-STS Connection" on page 155](#).

Direct configuration

In the case of direct configuration, your JAX-WS client proxy references an STSClient instance directly, by setting the `ws-security.sts.client` property on the client proxy. The value of `ws-security.sts.client` must be a reference to an STSClient instance.

For example, the following XML configuration shows how to instantiate a JAX-WS client proxy that references the STSClient with bean ID equal to `default.sts-client` (the bean ID is the same as the value of the name attribute):

```
<beans ...>
  ...
  <jaxws:client
    id="helloWorldProxy"
    serviceClass="org.apache.hello_world_soap_http.Greeter"
    address="https://localhost:9001/SoapContext/SoapPort">
    <jaxws:properties>
      <entry key="ws-security.sts.client"
        value-ref="default.sts-client" />
    </jaxws:properties>
  </jaxws:client>
  ...
```

```

<bean name="default.sts-client"
  class="org.apache.cxf.ws.security.trust.STSClient">
  <constructor-arg ref="cxf"/>
  <property name="wsdlLocation" value="sts/wsdl/ws-trust-1.4-service.wsdl"/>
  <property name="serviceName"
    value="{http://docs.oasis-open.org/ws-sx/ws-trust/200512/wsdl}SecurityTokenService
Provider"/>
  <property name="endpointName"
    value="{http://docs.oasis-open.org/ws-sx/ws-trust/200512/wsdl}SecurityTokenSer
viceSOAP"/>
</bean>
...
</beans>

```

Indirect configuration

In the case of indirect configuration, there is no need to set any property on the JAX-WS client proxy. Implicitly, if the IssuedToken policy assertion is applied to the relevant WSDL port, the runtime automatically searches for an STSClient bean named, *WSDLPortQName.sts-client*. To configure the STSClient bean indirectly, perform the following steps:

1. Define an STSClient bean, whose name attribute has the value, *WSDLPortQName.sts-client*.
2. Set *abstract="true"* on the bean element. This prevents Spring from instantiating the bean. The reason for this is that the runtime is responsible for the lifecycle of the STSClient object.
3. Set the relevant properties of the STSClient bean (typically, the *wsdlLocation*, *serviceName*, and *endpointName* properties). After the STSClient is instantiated in Java, the properties specified in XML will be injected into the STSClient instance.

For example, the following XML configuration creates a JAX-WS client proxy, which is associated with the *{http://apache.org/hello_world_soap_http}SoapPort* port (this is specified in an annotation on the service class, *Greeter*). When the client proxy needs to fetch an issued token for the first time, the runtime automatically creates an STSClient instance, searches for the bean named *WSDLPortQName.sts-client*, and injects the properties from that bean into the STSClient instance.

```

<beans ...>
...
<jaxws:client
  id="helloWorldProxy"
  serviceClass="org.apache.hello_world_soap_http.Greeter"
  address="https://localhost:9001/SoapContext/SoapPort"
/>
...
<bean name="{http://apache.org/hello_world_soap_http}SoapPort.sts-client"
  class="org.apache.cxf.ws.security.trust.STSClient"

```

```

    abstract="true">
    <constructor-arg ref="cxf"/>
    <property name="wsdlLocation" value="sts/wsdl/ws-trust-1.4-service.wsdl"/>
    <property name="serviceName"
        value="{http://docs.oasis-open.org/ws-sx/ws-trust/200512/wsdl}SecurityTokenService
Provider"/>
    <property name="endpointName"
        value="{http://docs.oasis-open.org/ws-sx/ws-trust/200512/wsdl}SecurityTokenSer
viceSOAP"/>
    </bean>
    ...
</beans>

```

Appendix A. ASN.1 and Distinguished Names

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.

ASN.1	182
Distinguished Names	183

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.

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 Fuse Services Framework, 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 A.1 on page 184](#) shows a selection of the attribute types that you are most likely to encounter:

Table A.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
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 A.1 on page 184](#)). For example:

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
```


Appendix B. OpenSSL Utilities

The openssl program consists of a large number of utilities that have been combined into one program. This appendix describes how you use the openssl program with Fuse Services Framework when managing X.509 certificates and private keys.

Using OpenSSL Utilities	188
Utilities Overview	189
The x509 Utility	191
The req Utility	193
The rsa Utility	195
The ca Utility	197
The s_client Utility	199
The s_server Utility	202
The OpenSSL Configuration File	206
Configuration Overview	207
[req] Variables	208
[ca] Variables	209
[policy] Variables	210
Example openssl.cnf File	212

Using OpenSSL Utilities

Utilities Overview	189
The x509 Utility	191
The req Utility	193
The rsa Utility	195
The ca Utility	197
The s_client Utility	199
The s_server Utility	202

Utilities Overview

The OpenSSL package

This section describes a version of the **openssl** program that is available with Eric Young's OpenSSL package, which can be downloaded from the OpenSSL Web site, <http://www.openssl.org>. OpenSSL is a publicly available implementation of the SSL protocol. Consult [Appendix C on page 215](#) for information about the copyright terms of OpenSSL.



Note

For complete documentation of the OpenSSL utilities, consult the documentation at the OpenSSL web site <http://www.openssl.org/docs>.

Command syntax

An **openssl** command line takes the following form:

```
openssl utility arguments
```

For example:

```
openssl x509 -in OrbixCA -text
```

The openssl utilities

This appendix describes the following **openssl** utilities:

x509	Manipulates X.509 certificates.
req	Creates and manipulates certificate signing requests, and self-signed certificates.
rsa	Manipulates RSA private keys.
ca	Implements a Certification Authority (CA).
s_client	Implements a generic SSL/TLS client.
s_server	Implements a generic SSL/TLS server.

The -help option

To get a list of the arguments associated with a particular command, use the `-help` option as follows:

```
openssl utility -help
```

For example:

```
openssl x509 -help
```

The x509 Utility

Purpose of the x509 utility

In Fuse Services Framework the **x509** utility is primarily used for:

- Printing text details of certificates you wish to examine.
- Converting certificates to different formats.

Options

The options supported by the openssl **x509** utility are as follows:

-inform <i>arg</i>	- input format - default PEM (one of DER, NET or PEM)
-outform <i>arg</i>	- output format - default PEM (one of DER, NET or PEM)
-keyform <i>arg</i>	- private key format - default PEM
-CAform <i>arg</i>	- CA format - default PEM
-CAkeyform <i>arg</i>	- CA key format - default PEM
-in <i>arg</i>	- input file - default stdin
-out <i>arg</i>	- output file - default stdout
-serial	- print serial number value
-hash	- print serial number value
-subject	- print subject DN
-issuer	- print issuer DN
-startdate	- notBefore field
-enddate	- notAfter field
-dates	- both Before and After dates
-modulus	- print the RSA key modulus
-fingerprint	- print the certificate fingerprint
-noout	- no certificate output
-days <i>arg</i>	- Time till expiry of a signed certificate - def 30 days

-signkey arg	- self sign cert with arg
-x509toreq	- output a certification request object
-req	- input is a certificate request, sign and output
-CA arg	- set the CA certificate, must be PEM format
-CAkey arg	- set the CA key, must be PEM format. If missing it is assumed to be in the CA file
-CAcreateserial	- create serial number file if it does not exist
-CAserial	- serial file
-text	- print the certificate in text form
-C	- print out C code forms
-md2/-md5/-sha1/ -mdc2	- digest algorithm used when signing certificates

Using the x509 utility

To print the text details of an existing PEM-format X.509 certificate, use the **x509** utility as follows:

```
openssl x509 -in MyCert.pem -inform PEM -text
```

To print the text details of an existing DER-format X.509 certificate, use the **x509** utility as follows:

```
openssl x509 -in MyCert.der -inform DER -text
```

To change a certificate from PEM format to DER format, use the **x509** utility as follows:

```
openssl x509 -in MyCert.pem -inform PEM -outform DER -out MyCert.der
```


The req Utility

Purpose of the req utility

The **req** utility is used to generate a self-signed certificate or a certificate signing request (CSR). A CSR contains details of a certificate issued by a CA. When creating a CSR, the **req** utility prompts you for the necessary information to produce a certificate request file and an encrypted private key file. The certificate request is then submitted to a CA for signing.

If the `-nodes` (no DES) parameter is not supplied to `req`, you are prompted for a pass phrase which is used to protect the private key.



Note

It is important to specify a validity period (using the `-days` parameter). If the certificate expires, applications using that certificate will not be authenticated successfully.

Options

The options supported by the `openssl req` utility are as follows:

<code>-inform arg</code>	input format - one of DER TXT PEM
<code>-outform</code>	arg output format - one of DER TXT PEM
<code>-in arg</code>	input file
<code>-out arg</code>	output file
<code>-text</code>	text form of request
<code>-noout</code>	do not output REQ
<code>-verify</code>	verify signature on REQ
<code>-modulus</code>	RSA modulus
<code>-nodes</code>	do not encrypt the output key
<code>-key file</code>	use the private key contained in file
<code>-keyform arg</code>	key file format
<code>-keyout arg</code>	file to send the key to
<code>-newkey rsa:bits</code>	generate a new RSA key of 'bits' in size

-newkey dsa:file	generate a new DSA key, parameters taken from CA in 'file'
-[digest]	Digest to sign with (md5, sha1, md2, mdc2)
-config file	request template file
-new	new request
-x509	output an x509 structure instead of a certificate req. (Used for creating self signed certificates)
-days	number of days an x509 generated by -x509 is valid for
-asn1-kludge	by default, the req command generates the correct PKCS#10 format for certificate requests that contain no attributes. However, certain CAs only accept requests containing no attributes in an invalid form: this option produces this invalid format.

Using the req Utility

To create a self-signed certificate with an expiry date a year from now, the **req** utility is used to create the certificate `CA_cert.pem` and the corresponding encrypted private key file `CA_pk.pem`, as follows:

```
openssl req -config ssl_conf_path_name -days 365
            -out CA_cert.pem -new -x509 -keyout CA_pk.pem
```

This following command creates the certificate request `MyReq.pem` and the corresponding encrypted private key file `MyEncryptedKey.pem`:

```
openssl req -config ssl_conf_path_name -days 365
            -out MyReq.pem -new -keyout MyEncryptedKey.pem
```

The rsa Utility

Purpose of the rsa utility

The **rsa** command is a useful utility for examining and modifying RSA private key files. Generally RSA keys are stored encrypted with a symmetric algorithm using a user-supplied pass phrase. The OpenSSL **req** command prompts the user for a pass phrase to encrypt the private key. By default, **req** uses the triple DES algorithm. The **rsa** command can be used to change the password that protects the private key and also to convert the format of the private key. Any **rsa** command that involves reading an encrypted **rsa** private key will prompt for the PEM pass phrase used to encrypt it.

Options

The options supported by the openssl **rsa** utility are as follows:

-inform arg	input format - one of DER NET PEM
-outform arg	output format - one of DER NET PEM
-in arg	input file
-out arg	output file
-des	encrypt PEM output with cbc des
-des3	encrypt PEM output with ede cbc des using 168 bit key
-text	print the key in text
-noout	do not print key out
-modulus	print the RSA key modulus

Using the rsa Utility

Converting a private key to PEM format from DER format requires using the **rsa** utility as follows:

```
openssl rsa -inform DER -in MyKey.der -outform PEM -out MyKey.pem
```

Changing the pass phrase that is used to encrypt the private key requires using the **rsa** utility as follows:

```
openssl rsa -inform PEM -in MyKey.pem -outform PEM -out MyKey.pem -des3
```

Removing encryption from the private key (which is not recommended) requires using the **rsa** command utility as follows:

```
openssl rsa -inform PEM -in MyKey.pem -outform PEM -out MyKey2.pem
```



Note

Do not specify the same file for the `-in` and `-out` parameters, because this can corrupt the file.

The ca Utility

Purpose of the ca utility

You can use the **ca** utility to create X.509 certificates by signing existing signing requests. It is imperative that you check the details of a certificate request before signing. Your organization should have a policy with respect to issuing certificates.

The **ca** utility is used to sign certificate requests thereby creating a valid X.509 certificate which can be returned to the request submitter. It can also be used to generate Certificate Revocation Lists (CRLS). For information on the **ca** -policy and -name options, refer to ["The OpenSSL Configuration File" on page 206](#).

Creating a new CA

To create a new CA using the openssl **ca** utility, two files (serial and index.txt) must be created in the location specified by the openssl configuration file that you are using.

Options

The options supported by the openssl **ca** utility are as follows:

-verbose	- Talk alot while doing things
-config file	- A config file
-name arg	- The particular CA definition to use
-gencrl	- Generate a new CRL
-crl days days	- Days is when the next CRL is due
-crl hours hours	- Hours is when the next CRL is due
-days arg	- number of days to certify the certificate for
-md arg	- md to use, one of md2, md5, sha or sha1
-policy arg	- The CA 'policy' to support
-keyfile arg	- PEM private key file
-key arg	- key to decode the private key if it is encrypted
-cert	- The CA certificate
-in file	- The input PEM encoded certificate request(s)
-out file	- Where to put the output file(s)

-outdir dir	- Where to put output certificates
-infile....	- The last argument, requests to process
-spkac file	- File contains DN and signed public key and challenge
-preservedN	- Do not re-order the DN
-batch	- Do not ask questions
-msie_hack	- msie modifications to handle all thos universal strings

Most of the above parameters have default values as defined in openssl.cnf.

Using the **ca** Utility

Converting a private key to PEM format from DER format requires the **ca** utility. To sign the supplied CSR MyReq.pem to be valid for 365 days and to create a new X.509 certificate in PEM format, use the **ca** utility as follows:

```
openssl ca -config ssl_conf_path_name -days 365
           -in MyReq.pem -out MyNewCert.pem
```

The s_client Utility

Purpose of the s_client utility

You can use the **s_client** utility to debug an SSL/TLS server. Using the **s_client** utility, you can negotiate an SSL/TLS handshake under controlled conditions, accompanied by extensive logging and error reporting.

Options

The options supported by the openssl **s_client** utility are as follows:

- | | |
|----------------------|--|
| -connect host[:port] | - Specify the host and (optionally) port to connect to. Default is local host and port 4433. |
| -cert certname | - Specifies the certificate to use, if one is requested by the server. |
| -certform format | - The certificate format, which can be either PEM or DER. Default is PEM. |
| -key keyfile | - File containing the client's private key. Default is to extract the key from the client certificate. |
| -keyform format | - The private key format, which can be either PEM or DER. Default is PEM. |
| -pass arg | - The private key password. |
| -verify depth | - Maximum server certificate chain length. |
| -CApath directory | - Directory to use for server certificate verification. |
| -CAfile file | - File containing trusted CA certificates. |
| -reconnect | - Reconnects to the same server five times using the same session ID. |
| -pause | - Pauses for one second between each read and write call. |
| -showcerts | - Display the whole server certificate chain. |
| -prexit | - Print session information when the program exits. |
| -state | - Prints out the SSL session states. |

-debug	- Log debug data, including hex dump of messages.
-msg	- Show all protocol messages with hex dump.
-nbio_test	- Tests non-blocking I/O.
-nbio	- Turns on non-blocking I/O.
-crlf	- Translates a line feed (LF) from the terminal into CR+LF, as required by some servers.
-ign_eof	- Inhibits shutting down the connection when end of file is reached in the input.
-quiet	- Inhibits printing of session and certificate information; implicitly turns on -ign_eof as well.
-ssl2, -ssl3, -tls1, -no_ssl2, -no_ssl3, -no_tls1	- These options enable/disable the use of certain SSL or TLS protocols.
-bugs	- Enables workarounds to several known bugs in SSL and TLS implementations.
-cipher cipherlist	- Specifies the cipher list sent by the client. The server should use the first supported cipher from the list sent by the client.
-starttls protocol	- Send the protocol-specific message(s) to switch to TLS for communication, where the protocol can be either smtp or pop3.
-engine id	- Specifies an engine, by it's unique id string.
-rand file(s)	- A file or files containing random data used to seed the random number generator, or an EGD socket. The file separator is ; for MS-Windows, , for OpenVMS, and : for all other platforms.

Using the **s_client** utility

Before running the **s_client** utility, there must be an active SSL/TLS server to connect to. For example, you can have an **s_server** test server running on the local host, listening on port 9000. To run the **s_client** test client, open a command prompt and enter the following:


```
openssl s_client -connect localhost:9000 -ssl3 -cert clientcert.pem
```

Where **clientcert.pem** is a file containing the client's X.509 certificate in PEM format. When you enter the command, you are prompted to enter the pass phrase for the **clientcert.pem** file.

The s_server Utility

Purpose of the s_server utility

You can use the **s_server** utility to debug an SSL/TLS client. By entering **openssl s_server** at the command line, you can run a simple SSL/TLS server that listens for incoming SSL/TLS connections on a specified port. The server can be configured to provide extensive logging and error reporting.

Options

The options supported by the openssl **s_server** utility are as follows:

- | | |
|---|---|
| -accept port | - Specifies the IP port to listen for incoming connections. Default is port 4433. |
| -context id | - Sets the SSL context id (any string value). |
| -cert certname | - Specifies the certificate to use for the server. |
| -certform format | - The certificate format, which can be either PEM or DER. Default is PEM. |
| -key keyfile | - File containing the server's private key. Default is to extract the key from the server certificate. |
| -keyform format | - The private key format, which can be either PEM or DER. Default is PEM. |
| -pass arg | - The private key password. |
| -dcert filename, -dkey keyname | - Specifies an additional certificate and private key, enabling the server to have multiple credentials. |
| -dcertform format, -dkeyform format, -dpass arg | - Specifies additional certificate format, private key format, and passphrase respectively. |
| -nocert | - If this option is set, no certificate is used. |
| -dhparam filename | - The DH parameter file to use. |
| -no_dhe | - If this option is set, no DH parameters will be loaded, effectively disabling the ephemeral DH cipher suites. |

-no_tmp_rsa	- Certain export cipher suites sometimes use a temporary RSA key. This option disables temporary RSA key generation.
-verify depth, -Verify depth	- Maximum client certificate chain length. With the -Verify option, the client must supply a certificate or an error occurs.
-CApath directory	- Directory to use for client certificate verification.
-CAfile file	- File containing trusted CA certificates.
-state	- Prints out the SSL session states.
-debug	- Log debug data, including hex dump of messages.
-msg	- Show all protocol messages with hex dump.
-nbio_test	- Tests non-blocking I/O.
-nbio	- Turns on non-blocking I/O.
-crlf	- Translates a line feed (LF) from the terminal into CR+LF, as required by some servers.
-quiet	- Inhibits printing of session and certificate information; implicitly turns on -ign_eof as well.
-ssl2, -ssl3, -tls1, -no_ssl2, -no_ssl3, -no_tls1	- These options enable/disable the use of certain SSL or TLS protocols.
-bugs	- Enables workarounds to several known bugs in SSL and TLS implementations.
-hack	- Enables a further workaround for some some early Netscape SSL code.
-cipher cipherlist	- Specifies the cipher list sent by the client. The server should use the first supported cipher from the list sent by the client.
-www	- Sends a status message back to the client when it connects. The status message is in HTML format.
-WWW	- Emulates a simple web server, where pages are resolved relative to the current directory.

-HTTP	- Emulates a simple web server, where pages are resolved relative to the current directory.
-engine id	- Specifies an engine, by it's unique id string.
-id_prefix_arg	- Generate SSL/TLS session IDs prefixed by arg.
-rand file(s)	- A file or files containing random data used to seed the random number generator, or an EGD socket. The file separator is ; for MS-Windows, , for OpenVMS, and : for all other platforms.

Connected commands

When an SSL client is connected to the test server, you can enter any of the following single letter commands on the server side:

q	End the current SSL connection but still accept new connections.
Q	End the current SSL connection and exit.
r	Renegotiate the SSL session.
R	Renegotiate the SSL session and request a client certificate.
P	Send some plain text down the underlying TCP connection. This should cause the client to disconnect due to a protocol violation.
S	Print out some session cache status information.

Using the s_server utility

To use the **s_server** utility to debug SSL clients, start the test server with the following command:

```
openssl s_server -accept 9000 -cert servercert.pem
```

Where the test server listens on the IP port 9000 and servercert.pem is a file containing the server's X.509 certificate in PEM format.

The **s_server** utility also provides a convenient way to test a secure Web browser. If you start the **s_server** utility with the **-www** switch, the test server functions as a simple Web server, serving up pages from the current directory; for example:

```
openssl s_server -accept 9000 -cert servercert.pem -www
```

The OpenSSL Configuration File

Configuration Overview	207
[req] Variables	208
[ca] Variables	209
[policy] Variables	210
Example openssl.cnf File	212

Configuration Overview

Overview

A number of OpenSSL commands (for example, **req** and **ca**) take a `-config` parameter that specifies the location of the openssl configuration file. This section provides a brief description of the format of the configuration file and how it applies to the **req** and **ca** commands. An example configuration file is shown at the end of this section.

Structure of the OpenSSL configuration file

The `openssl.cnf` configuration file consists of a number of sections that specify a series of default values that are used by the openssl commands.

[req] Variables

Overview of the variables

The req section contains the following variables:

```
default_bits = 1024
default_keyfile = privkey.pem
distinguished_name = req_distinguished_name
attributes = req_attributes
```

default_bits configuration variable

The `default_bits` variable is the default RSA key size that you want to use. Other possible values are 512, 2048, and 4096.

default_keyfile configuration variable

The `default_keyfile` variable is the default name for the private key file created by req.

distinguished_name configuration variable

The `distinguished_name` variable specifies the section in the configuration file that defines the default values for components of the distinguished name field. The `req_attributes` variable specifies the section in the configuration file that defines defaults for certificate request attributes.

[ca] Variables

Choosing the CA section

You can configure the file `openssl.cnf` to support a number of CAs that have different policies for signing CSRs. The `-name` parameter to the `ca` command specifies which CA section to use; for example:

```
openssl ca -name MyCa ...
```

This command refers to the CA section `[MyCa]`. If `-name` is not supplied to the `ca` command, the CA section used is the one indicated by the `default_ca` variable. In the ["Example openssl.cnf File" on page 212](#), this is set to `CA_default` (which is the name of another section listing the defaults for a number of settings associated with the `ca` command). Multiple different CAs can be supported in the configuration file, but there can be only one default CA.

Overview of the variables

Possible `[ca]` variables include the following

```
dir: The location for the CA database
    The database is a simple text database containing the
    following tab separated fields:

status: A value of 'R' - revoked, 'E' -expired or 'V' valid
issued date: When the certificate was certified
revoked date: When it was revoked, blank if not revoked
serial number: The certificate serial number
certificate: Where the certificate is located
CN: The name of the certificate
certs: Where the issued certificates are kept
```

The `serial number` field should be unique, as should the `CN/status` combination. The `ca` utility checks these at startup.

[policy] Variables

Choosing the policy section

The policy variable specifies the default policy section to use if the `-policy` argument is not supplied to the `ca` command. The CA policy section of a configuration file identifies the requirements for the contents of a certificate request which must be met before it is signed by the CA.

There are two policy sections defined in the ["Example openssl.cnf File" on page 212](#) : `policy_match` and `policy_anything`.

Example policy section

The `policy_match` section of the example `openssl.cnf` file specifies the order of the attributes in the generated certificate as follows:

```
countryName
stateOrProvinceName
organizationName
organizationalUnitName
commonName
emailAddress
```

The match policy value

Consider the following value:

```
countryName = match
```

This means that the country name must match the CA certificate.

The optional policy value

Consider the following value:

```
organisationalUnitName = optional
```

This means that the `organisationalUnitName` does not have to be present.

The supplied policy value

Consider the following value:

```
commonName = supplied
```

This means that the `commonName` must be supplied in the certificate request.

Example openssl.cnf File

Listing

The following shows the contents of an example openssl.cnf configuration file:

```
#####
# openssl example configuration file.
# This is mostly used for generation of certificate requests.
#####
[ ca ]
default_ca= CA_default # The default ca section
#####

[ CA_default ]
dir=/opt/iona/OrbixSSL1.0c/certs # Where everything is kept

certs=$dir # Where the issued certs are kept
crl_dir= $dir/crl # Where the issued crl are kept
database= $dir/index.txt # database index file
new_certs_dir= $dir/new_certs # default place for new certs
certificate=$dir/CA/OrbixCA # The CA certificate
serial= $dir/serial # The current serial number
crl= $dir/crl.pem # The current CRL
private_key= $dir/CA/OrbixCA.pk # The private key
RANDFILE= $dir/.rand # private random number file
default_days= 365 # how long to certify for
default_crl_days= 30 # how long before next CRL
default_md= md5 # which message digest to use
preserve= no # keep passed DN ordering

# A few different ways of specifying how closely the request should
# conform to the details of the CA

policy= policy_match

# For the CA policy

[policy_match]
countryName= match
stateOrProvinceName= match
organizationName= match
organizationalUnitName= optional
commonName= supplied
emailAddress= optional

# For the 'anything' policy
# At this point in time, you must list all acceptable 'object'
```

```

# types

[ policy_anything ]
countryName = optional
stateOrProvinceName= optional
localityName= optional
organizationName = optional
organizationalUnitName = optional
commonName= supplied
emailAddress= optional

[ req ]
default_bits = 1024
default_keyfile= privkey.pem
distinguished_name = req_distinguished_name
attributes = req_attributes

[ req_distinguished_name ]
countryName= Country Name (2 letter code)
countryName_min= 2
countryName_max = 2
stateOrProvinceName= State or Province Name (full name)
localityName = Locality Name (eg, city)
organizationName = Organization Name (eg, company)
organizationalUnitName = Organizational Unit Name (eg, section)
commonName = Common Name (eg. YOUR name)
commonName_max = 64
emailAddress = Email Address
emailAddress_max = 40

[ req_attributes ]
challengePassword = A challenge password
challengePassword_min = 4
challengePassword_max = 20
unstructuredName= An optional company name

```


Appendix C. Licenses

This appendix contains the text of licenses that are relevant to Fuse Services Framework.

OpenSSL License	216
-----------------------	-----

OpenSSL License

The licence agreement for using the OpenSSL command line utility shipped with Fuse Services Framework SSL/TLS is as follows:

LICENSE ISSUES

=====

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.

OpenSSL License

/* =====

* Copyright (c) 1998-1999 The OpenSSL Project. All rights reserved.

*

* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:

*

* 1. Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.

*

* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in
* the documentation and/or other materials provided with the
* distribution.

*

* 3. All advertising materials mentioning features or use of this
* software must display the following acknowledgment:

* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

*

* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
* endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.

*

* 5. Products derived from this software may not be called "OpenSSL"
* nor may "OpenSSL" appear in their names without prior written


```

* permission of the OpenSSL Project.
*
* 6. Redistributions of any form whatsoever must retain the following
* acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit (http://www.openssl.org/)"
*
* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS" AND ANY
* EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
* PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR
* ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
* SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
* LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
* STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
* ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
* OF THE POSSIBILITY OF SUCH DAMAGE.
* =====
*
* This product includes cryptographic software written by Eric Young
* (eay@cryptsoft.com). This product includes software written by Tim
* Hudson (tjh@cryptsoft.com).
*
*/

```

Original SSLeay License

```

/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
* All rights reserved.
*
* This package is an SSL implementation written
* by Eric Young (eay@cryptsoft.com).
* The implementation was written so as to conform with Netscapes SSL.
*
* This library is free for commercial and non-commercial use as long as
* the following conditions are aheared to. The following conditions
* apply to all code found in this distribution, be it the RC4, RSA,
* lhash, DES, etc., code; not just the SSL code. The SSL documentation
* included with this distribution is covered by the same copyright terms
* except that the holder is Tim Hudson (tjh@cryptsoft.com).
*

```

```

* Copyright remains Eric Young's, and as such any Copyright notices in
* the code are not to be removed.
* If this package is used in a product, Eric Young should be given attribution
* as the author of the parts of the library used.
* This can be in the form of a textual message at program startup or
* in documentation (online or textual) provided with the package.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the copyright
* notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in the
* documentation and/or other materials provided with the distribution.
* 3. All advertising materials mentioning features or use of this software
* must display the following acknowledgement:
* "This product includes cryptographic software written by
* Eric Young (eay@cryptsoft.com)"
* The word 'cryptographic' can be left out if the routines from the library
* being used are not cryptographic related :-).
* 4. If you include any Windows specific code (or a derivative thereof) from
* the apps directory (application code) you must include an acknowledgement:
* "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"
*
* THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS" AND
* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
* ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
* SUCH DAMAGE.
*
* The licence and distribution terms for any publically available version or
* derivative of this code cannot be changed. i.e. this code cannot simply be
* copied and put under another distribution licence
* [including the GNU Public Licence.]
*/

```

Index

Symbols

[ca] Variables, 209
[policy] Variables, 210
[req] Variables, 208

A

Abstract Syntax Notation One (see ASN.1)
administration
 OpenSSL command-line utilities, 36
ASN.1, 22, 181
 attribute types, 184
 AVA, 184
 OID, 184
ASN.1:
 RDN, 185
attribute value assertion, 184
authentication
 own certificate, specifying, 59
 SSL/TLS, 51
 mutual, 53
 trusted CA list, 56
AVA, 184

B

Basic Encoding Rules (see BER)
BER, 182

C

CA, 22
 choosing a host, 27
 commercial CAs, 26
 index file, 38
 list of trusted, 29
 multiple CAs, 28
 private CAs, 27
 private key, creating, 39
 security precautions, 27
 self-signed, 39

 serial file, 38
 trusted list, 56
ca utility, 197
CA, setting up, 37
CAs, 37
certificate signing request, 42, 45
 signing, 42, 47
certificates
 chaining, 28
 creating and signing, 44
 importing and exporting, 31
 own, specifying, 59
 peer, 28
 PKCS#12 file, 30
 public key, 22
 security handshake, 51, 53
 self-signed, 28, 39
 signing, 22, 42, 47
 signing request, 42, 45
 trusted CA list, 56
 X.509, 22
chaining of certificates, 28
configuration file, 206
CSR, 42, 45

D

DER, 182
Distinguished Encoding Rules (see DER)
distinguished names
 definition, 183
DN
 definition, 183
 string representation, 183

I

index file, 38

M

multiple CAs, 28
mutual authentication, 53

O

OpenSSL, 27

openssl

- configuration file, 206

- utilities, 188

OpenSSL command-line utilities, 36

openssl.cnf example file, 212

P

peer certificate, 28

PKCS#12 files

- creating, 31, 44

- definition, 30

- importing and exporting, 31

- viewing, 31

private key, 39

public keys, 22

R

RDN, 185

relative distinguished name, 185

req utility, 193

req Utility command, 193

root certificate directory, 29

rsa utility, 195

rsa Utility command, 195

S

security handshake

- SSL/TLS, 51, 53

self-signed CA, 39

self-signed certificate, 28

serial file, 38

signing certificates, 22

SSL/TLS

- security handshake, 51, 53

SSLeay, 27

T

target authentication, 51

target only, 51

trusted CA list policy, 56

trusted CAs, 29

X

X.500, 181

X.509 certificate

- definition, 22

x509 utility, 191