Java Platform, Standard Edition Security Developer's Guide





Java Platform, Standard Edition Security Developer's Guide, Release 9

E68624-04

Copyright © 1993, 2017, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience	xvi
Documentation Accessibility	xvi
Related Documents	xvi
Conventions	xvi
General Security	
Java Security Overview	1-1
Introduction to Java Security	1-1
Java Language Security and Bytecode Verification	1-2
Basic Security Architecture	1-3
Security Providers	1-3
File Locations	1-5
Java Cryptography	1-6
Public Key Infrastructure	1-7
Key and Certificate Storage	1-7
Public Key Infrastructure Tools	1-8
Authentication	1-9
Secure Communication	1-10
SSL, TLS, and DTLS Protocols	1-11
Simple Authentication and Security Layer (SASL)	1-11
Generic Security Service API and Kerberos	1-12
Access Control	1-12
Permissions	1-12
Security Policy	1-13
Access Control Enforcement	1-14
XML Signature	1-15
Additional Information about Java Security	1-16
Java Security Classes Summary	1-16
Deprecated Security APIs Marked for Removal	1-18
Security Tools Summary	1-19
Built-In Providers	1-20



Security Architecture	1-20
Standard Algorithm Names	1-20
Permissions in the Java Development Kit (JDK)	1-21
Permission Descriptions and Risks	1-22
NIO-Related Targets	1-23
Methods and the Required Permissions	1-23
java.lang.SecurityManager Method Permission Checks	1-48
Default Policy Implementation and Policy File Syntax	1-50
Default Policy Implementation	1-51
Default Policy File Locations	1-51
Modifying the Policy Implementation	1-52
Policy File Syntax	1-53
Policy File Examples	1-58
Property Expansion in Policy Files	1-60
Windows Systems, File Paths, and Property Expansion	1-62
Path-Name Canonicalization	1-63
General Expansion in Policy Files	1-65
API for Privileged Blocks	1-66
Using the doPrivileged API	1-66
3	
What It Means to Have Privileged Code	1-71
	1-71 1-73
What It Means to Have Privileged Code Reflection Troubleshooting Security	1-73 1-73
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guiden	1-73 1-73
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guidentroduction to Java Cryptography Architecture	1-73 1-73
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guiden	1-73 1-73 e 2-1
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guid ntroduction to Java Cryptography Architecture JCA Design Principles	1-73 1-73 e 2-1 2-2
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guidentroduction to Java Cryptography Architecture JCA Design Principles Provider Architecture	1-73 1-73 e 2-1 2-2 2-3
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guid ntroduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers	1-73 1-73 e 2-1 2-2 2-3 2-3 2-5
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guidentroduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers How Providers Are Actually Implemented	1-73 1-73 e 2-1 2-2 2-3 2-3 2-5 2-7
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guidentroduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers How Providers Are Actually Implemented Keystores	1-73 1-73 e 2-1 2-2 2-3 2-3 2-5 2-7
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guid Introduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers How Providers Are Actually Implemented Keystores Engine Classes and Algorithms	1-73 1-73 e 2-1 2-2 2-3 2-3 2-5 2-7 2-7
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guidentroduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers How Providers Are Actually Implemented Keystores Engine Classes and Algorithms Core Classes and Interfaces	1-73 1-73 e 2-1 2-2 2-3 2-3 2-5 2-7 2-7 2-8
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guidentroduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers How Providers Are Actually Implemented Keystores Engine Classes and Algorithms Core Classes and Interfaces The Provider Class	1-73 1-73 e 2-1 2-2 2-3 2-3 2-5 2-7 2-7 2-7 2-8 2-9
What It Means to Have Privileged Code Reflection Froubleshooting Security Java Cryptography Architecture (JCA) Reference Guidentroduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers How Providers Are Actually Implemented Keystores Engine Classes and Algorithms Core Classes and Interfaces The Provider Class How Provider Implementations Are Requested and Supplied	1-73 1-73 1-73 e 2-1 2-2 2-3 2-3 2-5 2-7 2-7 2-8 2-9 2-10
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guidentroduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers How Providers Are Actually Implemented Keystores Engine Classes and Algorithms Core Classes and Interfaces The Provider Class How Provider Implementations Are Requested and Supplied Installing Providers	1-73 1-73 1-73 e 2-1 2-2 2-3 2-3 2-5 2-7 2-7 2-7 2-8 2-9 2-10 2-12
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guid Introduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers How Providers Are Actually Implemented Keystores Engine Classes and Algorithms Core Classes and Interfaces The Provider Class How Provider Implementations Are Requested and Supplied Installing Providers Provider Class Methods	1-73 1-73 1-73 e 2-1 2-2 2-3 2-5 2-7 2-7 2-7 2-8 2-9 2-10 2-12 2-12
What It Means to Have Privileged Code Reflection Troubleshooting Security Java Cryptography Architecture (JCA) Reference Guid Introduction to Java Cryptography Architecture JCA Design Principles Provider Architecture Cryptographic Service Providers How Providers Are Actually Implemented Keystores Engine Classes and Algorithms Core Classes and Interfaces The Provider Class How Providers Implementations Are Requested and Supplied Installing Providers Provider Class Methods The Security Class	1-73 1-73 1-73 e 2-1 2-2 2-3 2-3 2-5 2-7 2-7 2-7 2-8 2-9 2-10 2-12 2-12 2-13



Creating a SecureRandom Object	2-16
Seeding or Re-Seeding the SecureRandom Object	2-16
Using a SecureRandom Object	2-17
Generating Seed Bytes	2-17
The MessageDigest Class	2-17
Creating a MessageDigest Object	2-17
Updating a Message Digest Object	2-18
Computing the Digest	2-18
The Signature Class	2-18
Signature Object States	2-19
Creating a Signature Object	2-19
Initializing a Signature Object	2-20
Signing with a Signature Object	2-20
Verifying with a Signature Object	2-21
The Cipher Class	2-21
Other Cipher-based Classes	2-29
The Cipher Stream Classes	2-30
The SealedObject Class	2-32
The Mac Class	2-34
Key Interfaces	2-35
The KeyPair Class	2-37
Key Specification Interfaces and Classes	2-37
The KeySpec Interface	2-37
The KeySpec Subinterfaces	2-37
The EncodedKeySpec Class	2-38
Generators and Factories	2-39
The KeyFactory Class	2-39
The SecretKeyFactory Class	2-40
The KeyPairGenerator Class	2-42
The KeyGenerator Class	2-44
The KeyAgreement Class	2-45
Key Management	2-47
The KeyStore Class	2-48
Algorithm Parameters Classes	2-52
The AlgorithmParameterSpec Interface	2-52
The AlgorithmParameters Class	2-53
The AlgorithmParameterGenerator Class	2-54
The CertificateFactory Class	2-55
How the JCA Might Be Used in a SSL/TLS Implementation	2-56
Cryptographic Strength Configuration	2-58
Jurisdiction Policy File Format	2-61



Standard Names	2-67
Packaging Your Application	2-67
Additional JCA Code Samples	2-68
Computing a MessageDigest Object	2-68
Generating a Pair of Keys	2-69
Generating and Verifying a Signature Using Generated Keys	2-71
Generating/Verifying Signatures Using Key Specifications and KeyFactory	2-71
Generating Random Numbers	2-73
Determining If Two Keys Are Equal	2-74
Reading Base64-Encoded Certificates	2-74
Parsing a Certificate Reply	2-75
Using Encryption	2-75
Using Password-Based Encryption	2-76
Sample Programs for Diffie-Hellman Key Exchange, AES/GCM, and HMAC-	
SHA256	2-77
Diffie-Hellman Key Exchange between 2 Parties	2-78
Diffie-Hellman Key Exchange between 3 Parties	2-81
AES/GCM Example	2-83
HMAC-SHA256 Example	2-85
HMAC-SHA256 Example	
·	
HMAC-SHA256 Example	
HMAC-SHA256 Example How to Implement a Provider in the Java Cryptography Arcl	nitecture
HMAC-SHA256 Example How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document	nitecture 3-1
HMAC-SHA256 Example How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology	nitecture 3-1 3-1
HMAC-SHA256 Example How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers	nitecture 3-1 3-1 3-1
HMAC-SHA256 Example How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes	3-1 3-1 3-2 3-2
HMAC-SHA256 Example How to Implement a Provider in the Java Cryptography Arcle Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider	3-1 3-1 3-2 3-2 3-5
HMAC-SHA256 Example How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code	3-1 3-1 3-2 3-2 3-5
HMAC-SHA256 Example How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and	3-1 3-1 3-1 3-2 3-5 3-5
HMAC-SHA256 Example How to Implement a Provider in the Java Cryptography Arcle Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations	3-1 3-1 3-2 3-5 3-6
How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations Step 2: Give your Provider a Name Step 3: Write Your Master Class, a Subclass of Provider Step 3.1: Create a Provider That Uses String Objects to Register Its	3-1 3-1 3-2 3-5 3-6 3-7
How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations Step 2: Give your Provider a Name Step 3: Write Your Master Class, a Subclass of Provider Step 3.1: Create a Provider That Uses String Objects to Register Its Services	3-1 3-1 3-2 3-5 3-5 3-6 3-7 3-7
How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations Step 2: Give your Provider a Name Step 3: Write Your Master Class, a Subclass of Provider Step 3.1: Create a Provider That Uses String Objects to Register Its Services Step 3.2: Create a Provider That Uses Provider. Service	3-1 3-1 3-1 3-2 3-5 3-5 3-7 3-7 3-8
How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations Step 2: Give your Provider a Name Step 3: Write Your Master Class, a Subclass of Provider Step 3.1: Create a Provider That Uses String Objects to Register Its Services Step 3.2: Create a Provider That Uses Provider. Service Step 3.3: Specify Additional Information for Cipher Implementations	3-1 3-1 3-2 3-5 3-5 3-6 3-7 3-7 3-10 3-12
How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations Step 2: Give your Provider a Name Step 3: Write Your Master Class, a Subclass of Provider Step 3.1: Create a Provider That Uses String Objects to Register Its Services Step 3.2: Create a Provider That Uses Provider.Service Step 3.3: Specify Additional Information for Cipher Implementations Step 4: Create a Module Declaration for Your Provider	3-1 3-1 3-1 3-2 3-5 3-5 3-7 3-7 3-10 3-12 3-14
How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations Step 2: Give your Provider a Name Step 3: Write Your Master Class, a Subclass of Provider Step 3.1: Create a Provider That Uses String Objects to Register Its Services Step 3.2: Create a Provider That Uses Provider.Service Step 3.3: Specify Additional Information for Cipher Implementations Step 4: Create a Module Declaration for Your Provider Step 5: Compile Your Code	3-1 3-1 3-2 3-5 3-5 3-7 3-7 3-12 3-14 3-15
How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations Step 2: Give your Provider a Name Step 3: Write Your Master Class, a Subclass of Provider Step 3.1: Create a Provider That Uses String Objects to Register Its Services Step 3.2: Create a Provider That Uses Provider.Service Step 3.3: Specify Additional Information for Cipher Implementations Step 4: Create a Module Declaration for Your Provider Step 5: Compile Your Code Step 6: Place Your Provider in a JAR File	3-1 3-1 3-1 3-2 3-5 3-5 3-7 3-7 3-8 3-10 3-12 3-14 3-15
How to Implement a Provider in the Java Cryptography Arcl Who Should Read This Document Notes on Terminology Introduction to Implementing Providers Engine Classes and Corresponding Service Provider Interface Classes Steps to Implement and Integrate a Provider Step 1: Write your Service Implementation Code Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations Step 2: Give your Provider a Name Step 3: Write Your Master Class, a Subclass of Provider Step 3.1: Create a Provider That Uses String Objects to Register Its Services Step 3.2: Create a Provider That Uses Provider.Service Step 3.3: Specify Additional Information for Cipher Implementations Step 4: Create a Module Declaration for Your Provider Step 5: Compile Your Code	3-1 3-1 3-2 3-5 3-5 3-7 3-7 3-12 3-14 3-15

How to Make Applications Exempt from Cryptographic Restrictions



2-63

	Step 7.2: Sign Your Provider	3-18
	Step 8: Prepare for Testing	3-18
	Step 8.1: Configure the Provider	3-18
	Step 8.2: Set Provider Permissions	3-20
	Step 9: Write and Compile Your Test Programs	3-21
	Step 10: Run Your Test Programs	3-21
	Step 11: Apply for U.S. Government Export Approval If Required	3-23
	Step 12: Document Your Provider and Its Supported Services	3-24
	Step 12.1: Indicate Whether Your Implementation is Cloneable for Message Digests and MACs	3-24
	Step 13: Make Your Class Files and Documentation Available to Clients	3-26
	Further Implementation Details and Requirements	3-26
	Alias Names	3-26
	Service Interdependencies	3-27
	Default Initialization	3-29
	Default Key Pair Generator Parameter Requirements	3-29
	The Provider.Service Class	3-30
	Signature Formats	3-31
	DSA Interfaces and their Required Implementations	3-32
	RSA Interfaces and their Required Implementations	3-34
	Diffie-Hellman Interfaces and their Required Implementations	3-36
	Interfaces for Other Algorithm Types	3-37
	Algorithm Parameter Specification Interfaces and Classes	3-38
	Key Specification Interfaces and Classes Required by Key Factories	3-41
	Secret-Key Generation	3-46
	Adding New Object Identifiers	3-46
	Ensuring Exportability	3-47
	Sample Code for MyProvider	3-48
4	JDK Providers Documentation	
	Introduction to JDK Providers	4-2
	Import Limits on Cryptographic Algorithms	4-3
	Cipher Transformations	4-3
	SecureRandom Implementations	4-3
	The SunPKCS11 Provider	4-4
	The SUN Provider	4-5
	The SunRsaSign Provider	4-7
	The SunJSSE Provider	4-8
	The SunJCE Provider	4-15
	The SunJGSS Provider	4-21
	The SunSASL Provider	4-21



The XMLDSIg Provider	4-21
The SunPCSC Provider	4-22
The SunMSCAPI Provider	4-23
The SunEC Provider	4-24
The OracleUcrypto Provider	4-25
The Apple Provider	4-26
The JdkLDAP Provider	4-27
The JdkSASL Provider	4-27
PKCS#11 Reference Guide	
SunPKCS11 Provider	5-1
SunPKCS11 Requirements	5-2
SunPKCS11 Configuration	5-2
Accessing Network Security Services (NSS)	5-7
Troubleshooting PKCS#11	5-10
Disabling PKCS#11 Providers and/or Individual PKCS#11 Mechanisms	5-10
Application Developers	5-11
Token Login	5-11
Token Keys	5-12
Delayed Provider Selection	5-13
JAAS KeyStoreLoginModule	5-14
Tokens as JSSE Keystore and Trust Stores	5-15
Using keytool and jarsigner with PKCS#11 Tokens	5-15
Policy Tool	5-16
Provider Developers	5-17
Provider Services	5-17
Parameter Support	5-18
SunPKCS11 Provider Supported Algorithms	5-18
SunPKCS11 Provider KeyStore Requirements	5-20
Example Provider	5-22
Java Authentication and Authorization Service (JAAS)
JAAS Reference Guide	6-1
JAAS Tutorials	6-1
Java Authentication and Authorization Service (JAAS): LoginModule Devel Guide	loper's 6-1
Introduction to LoginModule	6-2
Steps to Implement a LoginModule	6-4
	0
Step 1: Understand the Authentication Technology	6-4



Step 3: Implement the Abstract LoginModule Methods	6-4
Step 4: Choose or Write a Sample Application	6-8
Step 5: Compile the LoginModule and Application	6-9
Step 6: Prepare for Testing	6-9
Step 7: Test Use of the LoginModule	6-10
Step 8: Document Your LoginModule Implementation	6-11
Step 9: Make LoginModule JAR File and Documents Available	6-12
Java Generic Security Services (Java GSS-API)	
Java GSS-API and JAAS Tutorials for Use with Kerberos	7-1
Single Sign-on Using Kerberos in Java	7-1
Java GSS Advanced Security Programming	7-1
The Kerberos 5 GSS-API Mechanism	7-1
Java Secure Socket Extension (JSSE) Reference Guide	
Introduction to JSSE	8-1
JSSE Features and Benefits	8-2
JSSE Standard API	8-3
SunJSSE Provider	8-4
JSSE Related Documentation	8-4
Terms and Definitions	8-5
Secure Sockets Layer (SSL) Protocol Overview	8-8
Why Use SSL?	8-9
How SSL Works	8-10
Cryptographic Processes	8-10
Secret-Key Cryptography	8-11
Public-Key Cryptography	8-11
Comparison Between Secret-Key and Public-Key Cryptography	8-12
Public Key Certificates	8-12
Cryptographic Hash Functions	8-13
Message Authentication Code	8-13
Digital Signatures	8-13
The SSL Handshake	8-13
The SSL Protocol	8-14
Handshaking Again (Renegotiation)	8-16
Cipher Suite Choice and Remote Entity Verification	8-17
Client-Driven OCSP and OCSP Stapling	8-17
Client-Driven OCSP and Certificate Revocation	8-18



OCSP Stapling and Certificate Revocation	8-20
Setting Up a Java Client to Use OCSP Stapling	8-22
Setting Up a Java Server to Use OCSP Stapling	8-22
OCSP Stapling Configuration Properties	8-22
JSSE Classes and Interfaces	8-25
JSSE Core Classes and Interfaces	8-25
SocketFactory and ServerSocketFactory Classes	8-26
SSLSocketFactory and SSLServerSocketFactory Classes	8-26
Obtaining an SSLSocketFactory	8-26
SSLSocket and SSLServerSocket Classes	8-27
Obtaining an SSLSocket	8-27
SSLEngine Class	8-27
Creating an SSLEngine Object	8-29
Generating and Processing SSL/TLS Data	8-30
Datagram Transport Layer Security (DTLS) Protocol	8-33
Creating an SSLEngine Object for DTLS	8-42
Generating and Processing DTLS Data	8-43
Understanding SSLEngine Operation Statuses	8-45
Dealing With Blocking Tasks	8-49
Shutting Down a SSL/TLS/DTLS Connection	8-50
SSLSession and ExtendedSSLSession	8-51
HttpsURLConnection Class	8-52
Setting the Assigned SSLSocketFactory	8-52
Setting the Assigned HostnameVerifier	8-52
Support Classes and Interfaces	8-53
The SSLContext Class	8-54
The TrustManager Interface	8-56
The TrustManagerFactory Class	8-56
The X509TrustManager Interface	8-58
X509ExtendedTrustManager Class	8-61
The KeyManager Interface	8-64
The KeyManagerFactory Class	8-64
The X509KeyManager Interface	8-65
The X509ExtendedKeyManager Class	8-66
Relationship Between a TrustManager and a KeyManager	8-67
Secondary Support Classes and Interfaces	8-67
The SSLParameters Class	8-67
The SSLSessionContext Interface	8-68
The SSLSessionBindingListener Interface	8-68
The SSLSessionBindingEvent Class	8-68
The HandShakeCompletedListener Interface	8-69



The HandShakeCompletedEvent Class	8-69
The HostnameVerifier Interface	8-69
The X509Certificate Class	8-69
The AlgorithmConstraints Interface	8-70
The StandardConstants Class	8-70
The SNIServerName Class	8-70
The SNIMatcher Class	8-70
The SNIHostName Class	8-71
Customizing JSSE	8-72
How to Specify a java.lang.System Property	8-78
How to Specify a java.security.Security Property	8-78
Customizing the X509Certificate Implementation	8-79
Specifying an Alternative HTTPS Protocol Implementation	8-79
Customizing the Provider Implementation	8-80
Registering the Cryptographic Provider Statically	8-80
Registering the Cryptographic Service Provider Dynamically	8-80
Provider Configuration	8-81
Configuring the Preferred Provider for Specific Algorithms	8-81
Customizing the Default Keystores and Truststores, Store Types, and Store Passwords	8-82
Customizing the Default Key Managers and Trust Managers	8-84
Disabled and Restricted Cryptographic Algorithms	8-85
Customizing the Encryption Algorithm Providers	8-86
Customizing Size of Ephemeral Diffie-Hellman Keys	8-86
Customizing Maximum Fragment Length Negotiation (MFLN) Extension	8-87
Configuring the Maximum and Minimum Packet Size	8-88
Transport Layer Security (TLS) Renegotiation Issue	8-88
Phased Approach to Fixing This Issue	8-88
Description of the Phase 2 Fix	8-89
Workarounds and Alternatives to SSL/TLS Renegotiation	8-91
TLS Implementation Details	8-92
Description of the Phase 1 Fix	8-92
Allow Unsafe Server Certificate Change in SSL/TLS Renegotiations	8-93
Hardware Acceleration and Smartcard Support	8-93
Configuring JSSE to Use Smartcards as Keystores and Truststores	8-94
Multiple and Dynamic Keystores	8-94
Kerberos Cipher Suites	8-95
Kerberos Requirements	8-96
Peer Identity Information	8-97
Security Manager	8-97
Additional Keystore Formats (PKCS12)	8-98



Server Name Indication (SNI) Extension	8-98
TLS Application Layer Protocol Negotiation	8-100
Setting up ALPN on the Client	8-101
Setting up Default ALPN on the Server	8-102
Setting up Custom ALPN on the Server	8-103
Determining Negotiated ALPN Value during Handshaking	8-105
ALPN Related Classes and Methods	8-107
Troubleshooting JSSE	8-108
Configuration Problems	8-108
CertificateException While Handshaking	8-108
Runtime Exception: SSL Service Not Available	8-109
Runtime Exception: "No available certificate corresponding to the SSL cipher suites which are enabled"	8-109
Runtime Exception: No Cipher Suites in Common	8-110
Socket Disconnected After Sending ClientHello Message	8-110
SunJSSE Cannot Find a JCA Provider That Supports a Required Algorithm and Causes a NoSuchAlgorithmException	8-112
FailedDownloadException Thrown When Trying to Obtain Application Resources from Web Server over SSL	8-112
IllegalArgumentException When RC4 Cipher Suites are Configured for DTLS	8-113
Debugging Utilities	8-113
Debugging SSL/TLS Connections	8-115
Code Examples	8-131
Converting an Unsecure Socket to a Secure Socket	8-131
Running the JSSE Sample Code	8-134
Creating a Keystore to Use with JSSE	8-140
Using the Server Name Indication (SNI) Extension	8-144
Typical Client-Side Usage Examples	8-144
Typical Server-Side Usage Examples	8-145
Working with Virtual Infrastructures	8-145
Standard Names	8-150
Provider Pluggability	8-150
JSSE Cipher Suite Parameters	8-150
Java PKI Programmers Guide	
PKI Programmers Guide Overview	9-1
Introduction to Public Key Certificates	9-2
X.509 Certificates and Certificate Revocation Lists (CRLs)	9-3
Core Classes and Interfaces	9-7
Basic Certification Path Classes	9-8



9

	The CertificateFactory Class	9-9
	The CertPathParameters Interface	9-11
	Certification Path Validation Classes	9-11
	The CertPathValidator Class	9-11
	The CertPathValidatorResult Interface	9-12
	Certification Path Building Classes	9-13
	The CertPathBuilder Class	9-13
	The CertPathBuilderResult Interface	9-14
	Certificate/CRL Storage Classes	9-15
	The CertStore Class	9-15
	The CertStoreParameters Interface	9-16
	The CertSelector and CRLSelector Interfaces	9-17
	PKIX Classes	9-22
	The TrustAnchor Class	9-22
	The PKIXParameters Class	9-23
	The PKIXCertPathValidatorResult Class	9-25
	The PolicyNode Interface and PolicyQualifierInfo Class	9-26
	The PKIXBuilderParameters Class	9-27
	The PKIXCertPathBuilderResult Class	9-28
	The PKIXCertPathChecker Class	9-29
	Using PKIXCertPathChecker in Certificate Path Validation	9-34
	Implementing a Service Provider	9-38
	Steps to Implement and Integrate a Provider	9-39
	Service Interdependencies	9-41
	Certification Path Parameter Specification Interfaces	9-41
	Certification Path Result Specification Interfaces	9-42
	Certification Path Exception Classes	9-42
	Appendix A: Standard Names	9-42
	Appendix B: CertPath Implementation in SUN Provider	9-43
	Appendix C: OCSP Support	9-46
	Appendix D: CertPath Implementation in JdkLDAP Provider	9-48
	Appendix E: Disabling Cryptographic Algorithms	9-49
10	Java SASL API Programming and Deployment Guide	
	Java SASL API Overview	10-2
	Creating the Mechanisms	10-3
	Passing Input to the Mechanisms	10-3
	Using the Mechanisms	10-4
	Using the Negotiated Security Layer	10-5

The CertPath Class



9-8

How SASL Mechanisms are Installed and Selected	10-6
The SunSASL Provider	10-7
The SunSASL Provider Client Mechanisms	10-7
The SunSASL Provider Server Mechanisms	10-9
Debugging and Monitoring	10-10
The JdkSASL Provider	10-11
The JdkSASL Provider Client Mechanism	10-11
The JdkSASL Provider Server Mechanism	10-13
Implementing a SASL Security Provider	10-14
XML Digital Signature	
Java XML Digital Signature API Specification	11-1
Acknowledgements	11-1
Requirements	11-2
API Dependencies	11-3
Non-Goals	11-3
Package Overview	11-3
Service Providers	11-4
DOM Mechanism Requirements	11-5
Open API Issues	11-6
Programming Examples	11-6
XML Digital Signature API Overview and Tutorial	11-15
Package Hierarchy	11-16
Service Providers	11-17
Introduction to XML Signatures	11-18
Example of an XML Signature	11-18
XML Digital Signature API Examples	11-20
Validate Example	11-20
GenEnveloped Example	11-26
Security API Specification	

13 Deprecated Security APIs Marked for Removal



- 14 Security Tools
- 15 Security Tutorials



Preface

This guide provides information about the Java security technology, tools, and implementations of commonly used security algorithms, mechanisms, and protocols on the Java Platform, Standard Edition (Java SE).

Audience

This document is intended for experienced developers who build applications using the comprehensive Java security framework. It is also intended for the user or administrator with a a set of tools to securely manage applications.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documents

See Oracle JDK 9 Documentation for other JDK 9 guides.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.



1

General Security

Java Security Overview introduces you to cryptography, public key infrastructure, authentication, secure communication, access control, and XML signatures.

Security Architecture in the JDK 8 documentation provides an overview of the motivation of major security features, an introduction to security classes and their usage, a discussion of the impact of the security architecture on code, and thoughts on writing security-sensitive code.

Java Security Standard Algorithm Names Specification describes the set of standard names for algorithms, certificate and keystore types that Java SE requires and uses.

Permissions in the Java Development Kit (JDK) describes the built-in JDK permission types and discusses the risks of granting each permission.

Troubleshooting Security lists options for the java.security.debug system property that enable you to monitor security access.

Java Security Overview

Java security includes a large set of APIs, tools, and implementations of commonly-used security algorithms, mechanisms, and protocols. The Java security APIs span a wide range of areas, including cryptography, public key infrastructure, secure communication, authentication, and access control. Java security technology provides the developer with a comprehensive security framework for writing applications, and also provides the user or administrator with a set of tools to securely manage applications.

Introduction to Java Security

The JDK is designed with a strong emphasis on security. At its core, the Java language itself is type-safe and provides automatic garbage collection, enhancing the robustness of application code. A secure class loading and verification mechanism ensures that only legitimate Java code is executed. The Java security architecture includes a large set of application programming interfaces (APIs), tools, and implementations of commonly-used security algorithms, mechanisms, and protocols.

The Java security APIs span a wide range of areas. Cryptographic and public key infrastructure (PKI) interfaces provide the underlying basis for developing secure applications. Interfaces for performing authentication and access control enable applications to guard against unauthorized access to protected resources.

The APIs allow for multiple interoperable implementations of algorithms and other security services. Services are implemented in *providers*, which are plugged into the JDK through a standard interface that makes it easy for applications to obtain security services without having to know anything about their implementations. This allows developers to focus on how to integrate security into their applications, rather than on how to actually implement complex security mechanisms.



The JDK includes a number of providers that implement a core set of security services. It also allows for additional custom providers to be installed. This enables developers to extend the platform with new security mechanisms.

The JDK is divided into modules. Modules that contain security APIs include the following:

Table 1-1 Modules That Contain Security APIs

Module	Description
java.base	Defines the foundational APIs of Java SE. Contained packages include java.security, javax.crypto, javax.net.ssl, and javax.security.auth.
java.security.jgss	Defines the Java binding of the IETF Generic Security Services API (GSS-API). This module also contains GSS-API mechanisms including Kerberos v5 and SPNEGO.
java.security.sasl	Defines Java support for the IETF Simple Authentication and Security Layer (SASL). This module also contains SASL mechanisms including DIGEST-MD5, CRAM-MD5, and NTLM.
java.smartcardio	Defines the Java Smart Card I/O API.
java.xml.crypto	Defines the API for XML cryptography.
jdk.security.auth	Provides implementations of the javax.security.auth.* interfaces and various authentication modules.
jdk.security.jgss	Defines Java extensions to the GSS-API and an implementation of the SASL GSS-API mechanism.

Java Language Security and Bytecode Verification

The Java language is designed to be type-safe and easy to use. It provides automatic memory management, garbage collection, and range-checking on arrays. This reduces the overall programming burden placed on developers, leading to fewer subtle programming errors and to safer, more robust code.

A compiler translates Java programs into a machine-independent bytecode representation. A bytecode verifier is invoked to ensure that only legitimate bytecodes are executed in the Java runtime. It checks that the bytecodes conform to the Java Language Specification and do not violate Java language rules or namespace restrictions. The verifier also checks for memory management violations, stack underflows or overflows, and illegal data typecasts. Once bytecodes have been verified, the Java runtime prepares them for execution.

In addition, the Java language defines different access modifiers that can be assigned to Java classes, methods, and fields, enabling developers to restrict access to their class implementations as appropriate. The language defines four distinct access levels:

private: Most restrictive modifier; access is not allowed outside the particular class in which the private member (a method, for example) is defined.



- protected: Allows access to any subclass or to other classes within the same package.
- Package-private: If not specified, then this is the default access level; allows access to classes within the same package.
- public: No longer guarantees that the element is accessible everywhere; accessibility depends upon whether the package containing that element is exported by its defining module and whether that module is readable by the module containing the code that is attempting to access it.

Basic Security Architecture

The JDK defines a set of APIs spanning major security areas, including cryptography, public key infrastructure, authentication, secure communication, and access control. The APIs allow developers to easily integrate security into their application code.

The APIs are designed around the following principles:

Implementation independence

Applications do not need to implement security themselves. Rather, they can request security services from the JDK. Security services are implemented in providers (see the section Security Providers), which are plugged into the JDK via a standard interface. An application may rely on multiple independent providers for security functionality.

Implementation interoperability

Providers are interoperable across applications. Specifically, an application is not bound to a specific provider if it does not rely on default values from the provider.

Algorithm extensibility

The JDK includes a number of built-in providers that implement a basic set of security services that are widely used today. However, some applications may rely on emerging standards not yet implemented, or on proprietary services. The JDK supports the installation of custom providers that implement such services.

Security Providers

The <code>java.security.Provider</code> class encapsulates the notion of a security provider in the Java platform. It specifies the provider's name and lists the security services it implements. Multiple providers may be configured at the same time and are listed in order of preference. When a security service is requested, the highest priority provider that implements that service is selected.

Applications rely on the relevant getInstance method to request a security service from an underlying provider.

For example, message digest creation represents one type of service available from providers. To request an implementation of a specific message digest algorithm, call the method <code>java.security.MessageDigest.getInstance</code>. The following statement requests a SHA-256 message digest implementation without specifying a provider name:

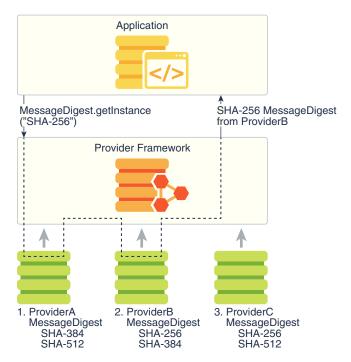
```
MessageDigest md = MessageDigest.getInstance("SHA-256");
```

The following figure illustrates how this statement obtains a SHA-256 message digest implementation. The providers are searched in preference order, and the



implementation from the first provider supplying that particular algorithm, ProviderB, is returned.

Figure 1-1 Request SHA-256 Message Digest Implementation Without Specifying Provider



You can optionally request an implementation from a specific provider by specifying the provider's name. The following statement requests a SHA-256 message digest implementation from a specific provider, ProviderC:

```
MessageDigest md = MessageDigest.getInstance("SHA-256", "ProviderC");
```

The following figure illustrates how this statement requests a SHA-256 message digest implementation from a specific provider, ProviderC. In this case, the implementation from that provider is returned, even though a provider with a higher preference order, ProviderB, also supplies a SHA-256 implementation.



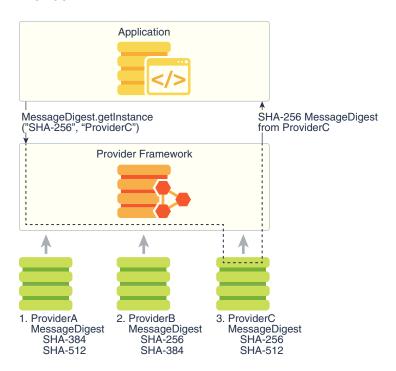


Figure 1-2 Request SHA-256 Message Digest Implementation from Specific Provider

For more information about cryptographic services, such as message digest algorithms, see the section Java Cryptography.

Oracle's implementation of the Java platform includes a number of built-in default providers that implement a basic set of security services that can be used by applications. Note that other vendor implementations of the Java platform may include different sets of providers that encapsulate vendor-specific sets of security services. The term built-in default providers refers to the providers available in Oracle's implementation.

File Locations

The following table lists locations of some security-related files and tools.



Table 1-2 Java security files and tools

File Name or Tool Name	Location	Description
java.security	<pre><java-home>/conf/security</java-home></pre>	Certain aspects of Java security, such as configuring the providers, may be customized by setting Security Properties. You may set Security Properties statically in the java.security file. Security Properties may also be set dynamically by calling appropriate methods of the Security class (in the java.security package).
java.policy	<pre><java-home>/conf/security</java-home></pre>	This is the default system policy file; see Security Policy.
Cryptographic policy directory	<pre><java-home>/conf/security/ policy</java-home></pre>	This directory contains sets of jurisdiction policy files; see Cryptographic Strength Configuration.
cacerts	<pre><java-home>/lib/security</java-home></pre>	The cacerts file represents a system-wide keystore with Certificate Authority (CA) and other trusted certificates. For information about configuring and managing this file, see keytool in Java Platform, Standard Edition Tools Reference.
keytool, jarsigner, policytool Windows only: kinit, klist, ktab	<java-home>/bin</java-home>	For more information about security-related tools, see Security Tools and Commands in Java Platform, Standard Edition Tools Reference.

Java Cryptography

The Java cryptography architecture is a framework for accessing and developing cryptographic functionality for the Java platform.

It includes APIs for a large variety of cryptographic services, including the following:

- Message digest algorithms
- Digital signature algorithms
- Symmetric bulk and stream encryption
- Asymmetric encryption
- Password-based encryption (PBE)
- Elliptic Curve Cryptography (ECC)
- Key agreement algorithms



- Key generators
- Message Authentication Codes (MACs)
- Secure Random Number Generators

For historical (export control) reasons, the cryptography APIs are organized into two distinct packages:

- The java.security and java.security.* packages contains classes that are not subject to export controls (like Signature and MessageDigest)
- The javax.crypto package contains classes that are subject to export controls (like Cipher and KeyAgreement)

The cryptographic interfaces are provider-based, allowing for multiple and interoperable cryptography implementations. Some providers may perform cryptographic operations in software; others may perform the operations on a hardware token (for example, on a smart card device or on a hardware cryptographic accelerator). Providers that implement export-controlled services must be digitally signed by a certificate issued by the Oracle JCE Certificate Authority.

The Java platform includes built-in providers for many of the most commonly used cryptographic algorithms, including the RSA, DSA, and ECDSA signature algorithms, the AES encryption algorithm, the SHA-2 message digest algorithms, and the Diffie-Hellman (DH) and Elliptic Curve Diffie-Hellman (ECDH) key agreement algorithms. Most of the built-in providers implement cryptographic algorithms in Java code.

The Java platform also includes a built-in provider that acts as a bridge to a native PKCS#11 (v2.x) token. This provider, named SunPKCS11, allows Java applications to seamlessly access cryptographic services located on PKCS#11-compliant tokens.

On Windows, the Java platform includes a built-in provider that acts as a bridge to the native Microsoft CryptoAPI. This provider, named SunMSCAPI, allows Java applications to seamlessly access cryptographic services on Windows through the CryptoAPI.

Public Key Infrastructure

Public Key Infrastructure (PKI) is a term used for a framework that enables secure exchange of information based on public key cryptography. It allows identities (of people, organizations, etc.) to be bound to digital certificates and provides a means of verifying the authenticity of certificates. PKI encompasses keys, certificates, public key encryption, and trusted Certification Authorities (CAs) who generate and digitally sign certificates.

The Java platform includes APIs and provider support for X.509 digital certificates and Certificate Revocation Lists (CRLs), as well as PKIX-compliant certification path building and validation. The classes related to PKI are located in the <code>java.security</code> and <code>java.security.cert</code> packages.

Key and Certificate Storage

The Java platform provides for long-term persistent storage of cryptographic keys and certificates via key and certificate stores. Specifically, the <code>java.security.KeyStore</code> class represents a *key store*, a secure repository of cryptographic keys and/or trusted certificates (to be used, for example, during certification path validation), and the <code>java.security.cert.CertStore</code> class represents a *certificate store*, a public and



potentially vast repository of unrelated and typically untrusted certificates. A CertStore may also store CRLs.

KeyStore and CertStore implementations are distinguished by types. The Java platform includes the standard PKCS11 and PKCS12 key store types (whose implementations are compliant with the corresponding PKCS specifications from RSA Security). It also contains a proprietary file-based key store type called JKS (which stands for Java Key Store), and a type called DKS (Domain Key Store) which is a collection of keystores that are presented as a single logical keystore.

The Java platform includes a special built-in key store, cacerts, that contains a number of certificates for well-known, trusted CAs. The keytool utility is able to list the certificates included in cacerts. See keytool in *Java Platform, Standard Edition Tools Reference*.

The SunPKCS11 provider mentioned in the section Java Cryptography includes a PKCS11 KeyStore implementation. This means that keys and certificates residing in secure hardware (such as a smart card) can be accessed and used by Java applications via the KeyStore API. Note that smart card keys may not be permitted to leave the device. In such cases, the <code>java.security.Key</code> object returned by the KeyStore API may simply be a reference to the key (that is, it would not contain the actual key material). Such a Key object can only be used to perform cryptographic operations on the device where the actual key resides.

The Java platform also includes an LDAP certificate store type (for accessing certificates stored in an LDAP directory), as well as an in-memory Collection certificate store type (for accessing certificates managed in a <code>java.util.Collection</code> object).

Public Key Infrastructure Tools

There are two built-in tools for working with keys, certificates, and key stores:

- keytool creates and manages key stores. Use it to perform the following tasks:
 - Create public/private key pairs
 - Display, import, and export X.509 v1, v2, and v3 certificates stored as files
 - Create X.509 certificates
 - Issue certificate (PKCS#10) requests to be sent to CAs
 - Create certificates based on certificate requests
 - Import certificate replies (obtained from the CAs sent certificate requests)
 - Designate public key certificates as trusted
 - Accept a password and store it securely as a secret key
- jarsigner signs JAR files and verifies signatures on signed JAR files. The Java ARchive (JAR) file format enables the bundling of multiple files into a single file. Typically, a JAR file contains the class files and auxiliary resources associated with applets and applications.

To digitally sign code, perform the following:

- 1. Use keytool to generate or import appropriate keys and certificates into your key store (if they are not there already).
- 2. Use the jar tool to package the code in a JAR file.



3. Use the jarsigner tool to sign the JAR file. The jarsigner tool accesses a key store to find any keys and certificates needed to sign a JAR file or to verify the signature of a signed JAR file.

Note:

jarsigner can optionally generate signatures that include a timestamp. Systems (such as Java Plug-in) that verify JAR file signatures can check the timestamp and accept a JAR file that was signed while the signing certificate was valid rather than requiring the certificate to be current. (Certificates typically expire annually, and it is not reasonable to expect JAR file creators to re-sign deployed JAR files annually.)

See keytool and jarsigner in Java Platform, Standard Edition Tools Reference.

Authentication

Authentication is the process of determining the identity of a user. In the context of the Java runtime environment, it is the process of identifying the user of an executing Java program. In certain cases, this process may rely on the services described in the section Java Cryptography.

The Java platform provides APIs that enable an application to perform user authentication via pluggable login modules. Applications call into the LoginContext class (in the javax.security.auth.login package), which in turn references a configuration. The configuration specifies which login module (an implementation of the javax.security.auth.spi.LoginModule interface) is to be used to perform the actual authentication.

Since applications solely talk to the standard LoginContext API, they can remain independent from the underlying plug-in modules. New or updated modules can be plugged in for an application without having to modify the application itself. The following figure illustrates the independence between applications and underlying login modules:



Application

Authentication Framework

Configuration

Smartcard

Kerberos

Username/
Password

Figure 1-3 Authentication Login Modules Plugging into the Authentication Framework

It is important to note that although login modules are pluggable components that can be configured into the Java platform, they are not plugged in via security providers. Therefore, they do not follow the provider searching model as described in the section Security Providers. Instead, as is shown in Figure 1-3, login modules are administered by their own unique configuration.

The Java platform provides the following built-in login modules, all in the com.sun.security.auth.module package:

- Krb5LoginModule for authentication using Kerberos protocols
- JndiLoginModule for username/password authentication using LDAP or NIS databases
- KeyStoreLoginModule for logging into any type of key store, including a PKCS#11 token key store

Authentication can also be achieved during the process of establishing a secure communication channel between two peers. The Java platform provides implementations of a number of standard communication protocols, which are discussed in the section Secure Communication.

Secure Communication

The data that travels across a network can be accessed by someone who is not the intended recipient. When the data includes private information, such as passwords and credit card numbers, steps must be taken to make the data unintelligible to unauthorized parties. It is also important to ensure that you are sending the data to the appropriate party, and that the data has not been modified, either intentionally or unintentionally, during transport.

Cryptography forms the basis required for secure communication; see the section Java Cryptography. The Java platform also provides API support and provider implementations for a number of standard secure communication protocols.



SSL, TLS, and DTLS Protocols

The JDK provides APIs and an implementation of the SSL, TLS, and DTLS protocols that includes functionality for data encryption, message integrity, and server and client authentication. Applications can use SSL/TLS/DTLS to provide for the secure passage of data between two peers over any application protocol, such as HTTP on top of TCP/IP.

The <code>javax.net.ssl.SSLSocket</code> class represents a network socket that encapsulates SSL/TLS support on top of a normal stream socket (<code>java.net.Socket</code>). Some applications might want to use alternate data transport abstractions (for example, New-I/O); the <code>javax.net.ssl.SSLEngine</code> class is available to produce and consume SSL/TLS/DTLS packets.

The JDK also includes APIs that support the notion of pluggable (provider-based) key managers and trust managers. *A key manager* is encapsulated by the <code>javax.net.ssl.KeyManager</code> class, and manages the keys used to perform authentication. A *trust manager* is encapsulated by the <code>TrustManager</code> class (in the same package), and makes decisions about who to trust based on certificates in the key store it manages.

The JDK includes a built-in provider that implements the SSL/TLS/DTLS protocols:

- SSLv3
- TLSv1
- TLSv1.1
- TLSv1.2
- DTLSv1.0
- DTLSv1.2

Simple Authentication and Security Layer (SASL)

Simple Authentication and Security Layer (SASL) is an Internet standard that specifies a protocol for authentication and optional establishment of a security layer between client and server applications. SASL defines how authentication data is to be exchanged, but does not itself specify the contents of that data. It is a framework into which specific authentication mechanisms that specify the contents and semantics of the authentication data can fit. There are a number of standard SASL mechanisms defined by the Internet community for various security levels and deployment scenarios.

The Java SASL API, which is in the <code>java.security.sasl</code> module, defines classes and interfaces for applications that use SASL mechanisms. It is defined to be mechanism-neutral; an application that uses the API need not be hardwired into using any particular SASL mechanism. Applications can select the mechanism to use based on desired security features. The API supports both client and server applications. The <code>javax.security.sasl.Sasl</code> class is used to create <code>SaslClient</code> and <code>SaslServer</code> objects.

SASL mechanism implementations are supplied in provider packages. Each provider may support one or more SASL mechanisms and is registered and invoked via the standard provider architecture.

The Java platform includes a built-in provider that implements the following SASL mechanisms:



- CRAM-MD5, DIGEST-MD5, EXTERNAL, GSSAPI, NTLM, and PLAIN client mechanisms
- CRAM-MD5, DIGEST-MD5, GSSAPI, and NTLM server mechanisms

Generic Security Service API and Kerberos

The Java platform contains an API with the Java language bindings for the Generic Security Service Application Programming Interface (GSS-API), which is in the java.security.jgss module. GSS-API offers application programmers uniform access to security services atop a variety of underlying security mechanisms. The Java GSS-API currently requires use of a Kerberos v5 mechanism, and the Java platform includes a built-in implementation of this mechanism. At this time, it is not possible to plug in additional mechanisms.



The Krb5LoginModule mentioned in the section Authentication can be used in conjunction with the GSS Kerberos mechanism.

The Java platform also includes a built-in implementation of the Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) GSS-API mechanism.

Before two applications can use GSS-API to securely exchange messages between them, they must establish a joint security context. The context encapsulates shared state information that might include, for example, cryptographic keys. Both applications create and use an org.ietf.jgss.GSSContext object to establish and maintain the shared information that makes up the security context. Once a security context has been established, it can be used to prepare secure messages for exchange.

The Java GSS APIs are in the org.ietf.jgss package. The Java platform also defines basic Kerberos classes, like <code>KerberosPrincipal</code>, <code>KerberosTicket</code>, <code>KerberosKey</code>, and <code>KeyTab</code>, which are located in the <code>javax.security.auth.kerberos</code> package.

Access Control

The access control architecture in the Java platform protects access to sensitive resources (for example, local files) or sensitive application code (for example, methods in a class). All access control decisions are mediated by a security manager, represented by the <code>java.lang.SecurityManager</code> class. A <code>SecurityManager</code> must be installed into the Java runtime in order to activate the access control checks.

Java applets and Java Web Start applications are automatically run with a securityManager installed. However, local applications executed via the java command are by default not run with a SecurityManager installed. In order to run local applications with a SecurityManager, either the application itself must programmatically set one via the setSecurityManager method (in the java.lang.System class), or java must be invoked with a -Djava.security.manager argument on the command line.

Permissions

A permission represents access to a system resource. In order for a resource access to be allowed for an applet (or an application running with a security manager), the



corresponding permission must be explicitly granted to the code attempting the access.

When Java code is loaded by a class loader into the Java runtime, the class loader automatically associates the following information with that code:

- Where the code was loaded from
- Who signed the code (if anyone)
- Default permissions granted to the code

This information is associated with the code regardless of whether the code is downloaded over an untrusted network (e.g., an applet) or loaded from the filesystem (e.g., a local application). The location from which the code was loaded is represented by a URL, the code signer is represented by the signer's certificate chain, and default permissions are represented by <code>java.security.Permission</code> objects.

The default permissions automatically granted to downloaded code include the ability to make network connections back to the host from which it originated. The default permissions automatically granted to code loaded from the local filesystem include the ability to read files from the directory it came from, and also from subdirectories of that directory.

Note that the identity of the user executing the code is not available at class loading time. It is the responsibility of application code to authenticate the end user if necessary (see the section Authentication). Once the user has been authenticated, the application can dynamically associate that user with executing code by invoking the doAs method in the javax.security.auth.Subject class.

Security Policy

A limited set of default permissions are granted to code by class loaders. Administrators have the ability to flexibly manage additional code permissions via a security policy.

Java SE encapsulates the notion of a security policy in the <code>java.security.Policy</code> class. There is only one <code>Policy</code> object installed into the Java runtime at any given time. The basic responsibility of the <code>Policy</code> object is to determine whether access to a protected resource is permitted to code (characterized by where it was loaded from, who signed it, and who is executing it). How a <code>Policy</code> object makes this determination is implementation-dependent. For example, it may consult a database containing authorization data, or it may contact another service.

Java SE includes a default Policy implementation that reads its authorization data from one or more ASCII (UTF-8) files configured in the security properties file. These policy files contain the exact sets of permissions granted to code: specifically, the exact sets of permissions granted to code loaded from particular locations, signed by particular entities, and executing as particular users. The policy entries in each file must conform to a documented proprietary syntax, and may be composed via a simple text editor or the graphical policytool utility.



The policytool is deprecated and marked for removal in the next major JDK release.



Access Control Enforcement

The Java runtime keeps track of the sequence of Java calls that are made as a program executes. When access to a protected resource is requested, the entire call stack, by default, is evaluated to determine whether the requested access is permitted.

As mentioned previously, resources are protected by the <code>SecurityManager</code>. Security-sensitive code in the JDK and in applications protects access to resources via code like the following:

```
SecurityManager sm = System.getSecurityManager();
if (sm != null) {
   sm.checkPermission(perm);
}
```

The Permission object perm corresponds to the requested access. For example, if an attempt is made to read the file / tmp/abc, the permission may be constructed as follows:

```
Permission perm = new java.io.FilePermission("/tmp/abc", "read");
```

The default implementation of SecurityManager delegates its decision to the java.security.AccessController implementation. The AccessController traverses the call stack, passing to the installed security Policy each code element in the stack, along with the requested permission (for example, the FilePermission in the previous example). The Policy determines whether the requested access is granted, based on the permissions configured by the administrator. If access is not granted, the AccessController throws a java.lang.SecurityException.

Figure 1-4 illustrates access control enforcement. In this particular example, there are initially two elements on the call stack, <code>ClassA</code> and <code>ClassB</code>. <code>ClassA</code> invokes a method in <code>ClassB</code>, which then attempts to access the file <code>/tmp/abc</code> by creating an instance of <code>java.io.FileInputStream</code>. The <code>FileInputStream</code> constructor creates a <code>FilePermission</code>, <code>perm</code>, as shown above, and then passes <code>perm</code> to the <code>SecurityManager</code> class's <code>checkPermission</code> method. In this particular case, only the <code>permissions</code> for <code>ClassA</code> and <code>ClassB</code> need to be <code>checked</code>, <code>because</code> all classes in the <code>java.base</code> module, including <code>FileInputStream</code>, <code>SecurityManager</code>, and <code>AccessController</code>, automatically receives all <code>permissions</code>.

In this example, <code>ClassA</code> and <code>ClassB</code> have different code characteristics — they come from different locations and have different signers. Each may have been granted a different set of permissions. The <code>AccessController</code> only grants access to the requested file if the <code>Policy</code> indicates that both classes have been granted the required <code>FilePermission</code>.



ClassB

ClassB

Location

Who

Signers

FileInputStream

sm.checkPermission(perm)

SecurityManager

AccessController

Policy

access granted or denied

abc

Figure 1-4 Controlling Access to Resources

XML Signature

The Java XML Digital Signature API is a standard Java API for generating and validating XML Signatures.

XML Signatures can be applied to data of any type, XML or binary (see XML Signature Syntax and Processing). The resulting signature is represented in XML. An XML Signature can be used to secure your data and provide data integrity, message authentication, and signer authentication.

The API is designed to support all of the required or recommended features of the W3C Recommendation for XML-Signature Syntax and Processing. The API is extensible and pluggable and is based on the Java Cryptography Service Provider Architecture.

The Java XML Digital Signature API, which is in the java.xml.crypto module, consists of six packages:

- javax.xml.crypto
- javax.xml.crypto.dsig
- javax.xml.crypto.dsig.keyinfo



- javax.xml.crypto.dsig.spec
- javax.xml.crypto.dom
- javax.xml.crypto.dsig.dom

Additional Information about Java Security

Find additional Java security documentation at Java SE Security.

Note:

Historically, as new types of security services were added to Java SE (sometimes initially as extensions), various acronyms were used to refer to them. Since these acronyms are still in use in the Java security documentation, here is an explanation of what they represent:

- JSSE (Java Secure Socket Extension) refers to the SSL-related services as described in the section SSL, TLS, and DTLS Protocols
- JCE (Java Cryptography Extension) refers to cryptographic services as described in the section Java Cryptography
- JAAS (Java Authentication and Authorization Service) refers to the authentication and user-based access control services as described in the sections Authentication and Access Control, respectively

Java Security Classes Summary

The following table describes some of the names, packages, and usage of the Java security classes and interfaces..

Table 1-3 Java security packages and classes

Package	Class/Interface Name	Usage	Module
java.lang	SecurityException	Indicates a security violation	java.base
java.lang	SecurityManager	Mediates all access control decisions	java.base
java.lang	System	Installs the SecurityManager	java.base
java.security	AccessController	Called by default implementation of SecurityManager to make access control decisions	java.base
java.security	DomainLoadStorePara meter	Stores parameters for the Domain keystore (DKS)	java.base
java.security	Key	Represents a cryptographic key	java.base



Table 1-3 (Cont.) Java security packages and classes

Package	Class/Interface Name	Usage	Module
java.security	KeyStore	Represents a repository of keys and trusted certificates	java.base
java.security	MessageDigest	Represents a message digest	java.base
java.security	Permission	Represents access to a particular resource	java.base
java.security	PKCS12Attribute	Supports attributes in PKCS12 keystores	java.base
java.security	Policy	Encapsulates the security policy	java.base
java.security	Provider	Encapsulates security service implementations	java.base
java.security	Security	Manages security providers and Security Properties	java.base
java.security	Signature	Creates and verifies digital signatures	java.base
java.security.cert	Certificate	Represents a public key certificate	java.base
java.security.cert	CertStore	Represents a repository of unrelated and typically untrusted certificates	java.base
java.security.cert	CRL	Represents a CRL	java.base
javax.crypto	Cipher	Performs encryption and decryption	java.base
javax.crypto	KeyAgreement	Performs a key exchange	java.base
javax.net.ssl	KeyManager	Manages keys used to perform SSL/TLS authentication	java.base
javax.net.ssl	SSLEngine	Produces/consumes SSL/TLS packets, allowing the application freedom to choose a transport mechanism	java.base
javax.net.ssl	SSLSocket	Represents a network socket that encapsulates SSL/TLS support on top of a normal stream socket	java.base



Table 1-3 (Cont.) Java security packages and classes

Package	Class/Interface Name	Usage	Module
javax.net.ssl	TrustManager	Makes decisions about who to trust in SSL/TLS interactions (for example, based on trusted certificates in key stores)	java.base
javax.security.auth	Subject	Represents a user	java.base
<pre>javax.security.auth .kerberos</pre>	KerberosPrincipal	Represents a Kerberos principal	java.base
javax.security.auth .kerberos	KerberosTicket	Represents a Kerberos ticket	java.base
<pre>javax.security.auth .kerberos</pre>	KerberosKey	Represents a Kerberos key	java.base
javax.security.auth .kerberos	KerberosTab	Represents a Kerberos keytab file	java.base
<pre>javax.security.auth .login</pre>	LoginContext	Supports pluggable authentication	java.base
javax.security.auth .spi	LoginModule	Implements a specific authentication mechanism	java.base
javax.security.sasl	Sasl	Creates SaslClient and SaslServer objects	java.security.sasl
javax.security.sasl	SaslClient	Performs SASL authentication as a client	java.security.sasl
javax.security.sasl	SaslServer	Performs SASL authentication as a server	java.security.sasl
org.ietf.jgss	GSSContext	Encapsulates a GSS- API security context and provides the security services available via the context	java.security.jgss
com.sun.security.au th.module	JndiLoginModule	Performs username/ password authentication using LDAP or NIS	jdk.security.auth
com.sun.security.au th.module	KeyStoreLoginModule	Performs authentication based on key store login	jdk.security.auth
com.sun.security.au th.module	Krb5LoginModule	Performs authentication using Kerberos protocols	jdk.security.auth

Deprecated Security APIs Marked for Removal

The following APIs are deprecated and eligible to be removed in a future release.



You can check the API dependencies using the jdeprscan tool. See jdeprscan in *Java Platform, Standard Edition Tools Reference*.

The following classes are deprecated and marked for removal:

- com.sun.security.auth.PolicyFile
- com.sun.security.auth.SolarisNumericGroupPrincipal
- com.sun.security.auth.SolarisNumericUserPrincipal
- com.sun.security.auth.SolarisPrincipal
- com.sun.security.auth.X500Principal
- com.sun.security.auth.module.SolarisLoginModule
- com.sun.security.auth.module.SolarisSystem

The following methods are deprecated and marked for removal:

- java.lang.SecurityManager.getInCheck
- java.lang.SecurityManager.checkMemberAccess
- java.lang.SecurityManager.classDepth
- java.lang.SecurityManager.currentClassLoader
- java.lang.SecurityManager.currentLoadedClass
- java.lang.SecurityManager.inClass
- java.lang.SecurityManager.inClassLoader
- java.lang.SecurityManager.checkAwtEventQueueAccess
- java.lang.SecurityManager.checkTopLevelWindow
- java.lang.SecurityManager.checkSystemClipboardAccess

The following field is deprecated and marked for removal:

java.lang.SecurityManager.incheck

Security Tools Summary

The following tables describe Java security and Kerberos-related tools.

See Security Tools and Commands in *Java Platform, Standard Edition Tools Reference*.

Table 1-4 Java Security Tools

Tool	Usage
jar	Creates Java Archive (JAR) files
jarsigner	Signs and verifies signatures on JAR files
keytool	Creates and manages key stores



Table 1-4 (Cont.) Java Security Tools

Usage
Creates and edits policy files for use with default Policy implementation
Note: policytool is deprecated and marked for removal.

There are also three Kerberos-related tools that are shipped with the JDK for Windows. Equivalent functionality is provided in tools of the same name that are automatically part of the Solaris and Linux operating environments.

Table 1-5 Kerberos-related Tools

Tool	Usage
kinit	Obtains and caches Kerberos ticket-granting tickets
klist	Lists entries in the local Kerberos credentials cache and key table
ktab	Manages the names and service keys stored in the local Kerberos key table

Built-In Providers

The Java SE implementation from Oracle includes a number of built-in provider packages. See JDK Providers Documentation.

Security Architecture

See Security Architecture in the JDK 8 documentation for an overview of the motivation of major security features, an introduction to security classes and their usage, a discussion of the impact of the security architecture on code, and thoughts on writing security-sensitive code.

Standard Algorithm Names

See Java Security Standard Algorithm Names Specification for information about the set of standard names for algorithms, certificate and keystore types that Java SE requires and uses.



Permissions in the Java Development Kit (JDK)

Information about the built-in JDK permission types and associated risks of granting each permission. Information about methods that require permissions to be in effect in order to be successful, and for each method lists the required permission.

A permission represents access to a system resource. In order for a resource access to be allowed for an applet (or an application running with a security manager), the corresponding permission must be explicitly granted to the code attempting the access.

A permission typically has a name (often referred to as a "target name") and, in some cases, a comma-separated list of one or more actions.

For example, the following code creates a FilePermission object representing read access to the file named abc in the /tmp directory:

```
perm = new java.io.FilePermission("/tmp/abc", "read");
```

Here, the target name is "/tmp/abc" and the action string is "read".



Important:

The above statement creates a permission object. A permission object represents, but does not grant access to, a system resource. Permission objects are constructed and assigned ("granted") to code based on the policy in effect. When a permission object is assigned to some code, that code is granted the permission to access the system resource specified in the permission object, in the specified manner, A permission object may also be constructed by the current security manager when making access decisions. In this case, the (target) permission object is created based on the requested access, and checked against the permission objects granted to and held by the code making the request.

The policy for a Java application environment is represented by a Policy object. In the "JavaPolicy" Policy implementation, the policy can be specified within one or more policy configuration files. The policy file(s) specify what permissions are allowed for code from specified code sources. A sample policy file entry that grants code from the /home/sysadmin directory read access to the file /tmp/abc is

```
grant codeBase "file:/home/sysadmin/" {
    permission java.io.FilePermission "/tmp/abc", "read";
};
```

To know more about policy file locations and granting permissions in policy files, see Default Policy Implementation and Policy File Syntax.

Using the policy tool saves typing and eliminates the need for you to know the required syntax of policy files. To know more about using the policy tool to specify the permissions, see Policy Tool. Using the policy tool saves typing and eliminates the need for you to know the required syntax of policy files.

Technically, whenever a resource access is attempted, all code traversed by the execution thread up to that point must have permission for that resource access,



unless some code on the thread has been marked as "privileged." See API for Privileged Blocks.

Permission Descriptions and Risks

List of built-in JDK permission types and the risks of granting each permission.

- AWTPermission
- FilePermission
- SerializablePermission
- ManagementPermission
- ReflectPermission
- RuntimePermission
 - NIO-Related Targets
- NetPermission
- SocketPermission
- LinkPermission
- URLPermission
- AllPermission
- SecurityPermission
- UnresolvedPermission
- SQLPermission
- LoggingPermission
- PropertyPermission
- MBeanPermission
- MBeanServerPermission
- MBeanTrustPermission
- SubjectDelegationPermission
- SSLPermission
- AuthPermission
- DelegationPermission
- ServicePermission
- PrivateCredentialPermission
- AudioPermission
- JAXBPermission
- WebServicePermission
- Methods and the Required Permissions
- java.lang.SecurityManager Method Permission Checks



NIO-Related Targets

NIO-related related target names.

Two NIO-related RuntimePermission targets were added in the 1.4 release of the JavaSE JDK:

selectorProvider charsetProvider

These RuntimePermissions are required to be granted to classes which subclass and implement java.nio.channel.spi.SelectorProvider or java.nio.charset.spi.CharsetProvider. The permission is checked during invocation of the abstract base class constructor. These permissions ensure trust in classes which implement these security-sensitive provider mechanisms. For more information, see

java.nio.channels.spi.SelectorProviderjava.nio.channels.spi.CharsetProvider

Methods and the Required Permissions

List of all the methods that require permissions, and for each method the corresponding SecurityManager method it calls.



The list of all the methods discussed in this document is not complete and does not include several methods that require permissions. See API Documentation for additional information on methods that throw SecurityException and the permissions that are required.

In the default <code>SecurityManager</code> method implementations, a call to a method in the <code>Method</code> column can only be successful if the permission specified in the corresponding entry in the <code>SecurityManager Method</code> column is allowed by the policy currently in effect.

Example 1-1 SecurityManager checkPermission Method

getSystemEventQueuejava.awt.ToolkitcheckPermission

java.awt.AWTPermission "accessEventQueue";

Method	SecurityManager Method	Permission
	checkPermission	java.awt.AWTPermission
<pre>java.awt.Toolkit getSystemEventQueue();</pre>		"accessEventQueue";

The following convention means the runtime value of foo replaces the string $\{foo\}$ in the permission name:



Method	SecurityManager Method	Permission
	checkXXX	SomePermission "{foo}";
<pre>some.package.class public static void someMethod(String foo);</pre>		

Example 1-2 SecurityManager checkRead Method

FileInputStreamjava.io.FileInputStream checkRead

Method	SecurityManager Method	Permission
<pre>java.io.FileInputStream FileInputStream(String name)</pre>	checkRead(String)	<pre>java.io.FilePermission "{name}", "read";</pre>

If the FileInputStream method (in this case, a constructor) is called with "/test/MyTestFile" as the name argument, as in

```
FileInputStream("/test/MyTestFile");
```

then in order for the call to succeed, the following permission must be set in the current policy, allowing read access to the file "/test/MyTestFile":

```
java.io.FilePermission "/test/MyTestFile", "read";
```

More specifically, the permission must either be explicitly set, as above, or implied by another permission, such as the following:

```
java.io.FilePermission "/test/*", "read";
```

which allows read access to any files in the "/test" directory.

Example 1-3 SecurityManager checkAccept Method

In some cases, a term in braces is not exactly the same as the name of a specific method argument but is meant to represent the relevant value:

Method	SecurityManager Method	Permission
java.net.DatagramSocket public synchronized void	<pre>checkAccept({host}, {port})</pre>	<pre>java.net.SocketPermission "{host}:{port}", "accept";</pre>
<pre>receive(DatagramPacket p);</pre>		

Here, the appropriate host and port values are calculated by the receive method and passed to <code>checkAccept</code>.

In most cases, just the name of the SecurityManager method called is listed. Where the method is one of multiple methods of the same name, the argument types are also



listed, for example for checkRead(String) and checkRead(FileDescriptor). In other cases where arguments may be relevant, they are also listed.

Methods and the Permissions

The following table is ordered by package name, the methods in classes in the <code>java.awt</code> package are listed first, followed by methods in classes in the <code>java.io</code> package, and so on:

Table 1-6 Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.awt.Graphics2d public abstract void setComposite(Composite comp)</pre>	checkPermission	java.awt.AWTPermission "readDisplayPixels" if this Graphics2D context is drawing to a Component on the display screen and the Composite is a custom object rather than an instance of the AlphaComposite class. Note: The setComposite method is actually abstract and thus can't invoke security checks. Each actual implementation of the method should call the java.lang.SecurityManager checkPermission method with a java.awt.AWTPermission("rea dDisplayPixels") permission under the conditions noted.
<pre>java.awt.Robot public Robot() public Robot(GraphicsDevice screen)</pre>	checkPermission	java.awt.AWTPermission "createRobot"
<pre>java.awt.Toolkit public void addAWTEventListener(</pre>	checkPermission	java.awt.AWTPermission "listenToAllAWTEvents"



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
	checkPrintJobAccess	java.lang.RuntimePermission "queuePrintJob"
<pre>java.awt.Toolkit public abstract PrintJob getPrintJob(</pre>		Note: The getPrintJob method is actually abstract and thus can't invoke security checks. Each actual implementation of the method should call the java.lang.SecurityManager checkPrintJobAccess method, which is successful only if the java.lang.RuntimePermission "queuePrintJob" permission is currently allowed.
java.awt.Toolkit	checkPermission	java.awt.AWTPermission "accessClipboard"
public abstract Clipboard		Note: The getSystemClipboard method is
<pre>getSystemClipboard()</pre>		actually abstract and thus can't invoke security checks. Each actual implementation of the method should call the checkPermission method, which is successful only if the java.awt.AWTPermission "accessClipboard" permission is currently allowed.
	checkPermission	java.awt.AWTPermission "accessEventQueue"
java.awt.Toolkit public final EventQueue		•
<pre>getSystemEventQueue()</pre>		
<pre>java.awt.Window Window()</pre>	checkPermission	If java.awt.AWTPermission "showWindowWithoutWarning Banner" is set, the window will be displayed without a banner warning that the window was created by an applet. It it's not set, such a banner will be displayed.



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
	checkPropertiesAccess	java.util.PropertyPermission
<pre>java.beans.Beans public static void setDesignTime(</pre>		"*", "read,write"
<pre>isDesignTime) public static void setGuiAvailable(</pre>		
isGuiAvailable)		
<pre>java.beans.Introspector public static synchronized void</pre>		
<pre>setBeanInfoSearchPath(Strin g path[])</pre>		
<pre>java.beans.PropertyEditorMa nager public static void registerEditor(</pre>		
<pre>java.io.File public boolean delete() public void deleteOnExit()</pre>	checkDelete(String)	java.io.FilePermission "{name}", "delete"
java.io.FileInputStream	checkRead(FileDescriptor)	java.lang.RuntimePermission "readFileDescriptor"
<pre>FileInputStream(FileDescrip tor fdObj)</pre>		



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
	checkRead(String)	java.io.FilePermission
java.io.FileInputStream		"{name}", "read"
FileInputStream(String		
name)		
FileInputStream(File		
file)		
java.io.File		
<pre>public boolean exists()</pre>		
<pre>public boolean canRead()</pre>		
<pre>public boolean isFile()</pre>		
public boolean		
<pre>isDirectory()</pre>		
<pre>public boolean isHidden()</pre>		
public long		
<pre>lastModified()</pre>		
<pre>public long length()</pre>		
<pre>public String[] list()</pre>		
<pre>public String[] list(</pre>		
FilenameFilter		
filter)		
<pre>public File[] listFiles()</pre>		
<pre>public File[] listFiles(</pre>		
FilenameFilter		
filter)		
<pre>public File[] listFiles(</pre>		
FileFilter		
filter)		
java.io.RandomAccessFile		
RandomAccessFile(String		
name, String mode)		
RandomAccessFile(File		
file, String mode)		
(where mode is "r"		
in both of these)		
	checkWrite(FileDescriptor)	java.lang.RuntimePermission
java.io.FileOutputStream	ssaktinto(i ilobosonptoi)	"writeFileDescriptor"
Java.10.F11eoutputstreall		
FileOutputStream(FileDescri		
ptor fdObj)		



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
java.io.FileOutputStream FileOutputStream(File file) FileOutputStream(String name) FileOutputStream(String name, boolean append) java.io.File public boolean canWrite() public boolean createNewFile() public static File createTempFile(String prefix, String suffix) public static File createTempFile(String prefix, String suffix, File directory) public boolean mkdir()	SecurityManager Method checkWrite(String)	Permission java.io.FilePermission "{name}", "write"
<pre>public boolean mkdirs() public boolean renameTo(File dest) public boolean setLastModified(long time) public boolean setReadOnly()</pre>		
	checkPermission	java.io.SerializablePermission
<pre>java.io.ObjectInputStream protected final boolean</pre>		"enableSubstitution"
<pre>enableResolveObject(boolean enable);</pre>		
<pre>java.io.ObjectOutputStream protected final boolean</pre>		
<pre>enableReplaceObject(boolean enable)</pre>		



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.io.ObjectInputStream protected ObjectInputStream() java.io.ObjectOutputStream protected</pre>	checkPermission	java.io.SerializablePermission "enableSubclassImplementati on"
<pre>java.io.RandomAccessFile RandomAccessFile(String name, String mode)</pre>	checkRead(String) and checkWrite(String)	java.io.FilePermission "{name}", "read,write"
<pre>java.lang.Class public static Class forName(String name, boolean initialize, ClassLoader loader)</pre>	checkPermission	If loader is null, and the caller's class loader is not null, then java.lang.RuntimePer mission("getClassLoader")
java.lang.Class public ClassLoader getClassLoader()	checkPermission	If the caller's class loader is null, or is the same as or an ancestor of the class loader for the class whose class loader is being requested, no permission is needed. Otherwise, java.lang.RuntimePermissio n "getClassLoader" is required.



Table 1-6 (Cont.) Methods and the Permissions

No. diameter	One of Manager Matheal	Bernitaria
Method	SecurityManager Method	Permission
<pre>java.lang.Class public Class[] getDeclaredClasses() public Field[] getDeclaredFields() public Method[] getDeclaredMethods() public Constructor[] getDeclaredConstructors() public Field getDeclaredField(</pre>	checkMemberAccess(this, Member.DECLARED) and, if this class is in a package, checkPackageAccess({pkgNa me})	Default checkMemberAccess does not require any permissions if "this" class's classloader is the same as that of the caller. Otherwise, it requires java.lang.RuntimePermission "accessDeclaredMembers". If this class is in a package, java.lang.RuntimePermission "accessClassInPackage. {pkgName}" is also required.
String name) public Method getDeclaredMethod() public Constructor getDeclaredConstructor()		
geobeetateaconstructor()		
<pre>java.lang.Class public Class[] getClasses() public Field[] getFields() public Method[] getMethods() public Constructor[] getConstructors() public Field getField(String name) public Method getMethod() public Constructor getConstructor</pre>	checkMemberAccess(this, Member.PUBLIC) and, if class is in a package, checkPackageAccess({pkgNa me})	Default checkMemberAccess does not require any permissions when the access type is Member.PUBLIC. If this class is in a package, java.lang.RuntimePermission "accessClassInPackage. {pkgName}" is required.
<pre>java.lang.Class public ProtectionDomain getProtectionDomain()</pre>	checkPermission	java.lang.RuntimePermission "getProtectionDomain"
<pre>java.lang.ClassLoader ClassLoader() ClassLoader(ClassLoader parent)</pre>	checkCreateClassLoader	java.lang.RuntimePermission "createClassLoader"



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.lang.ClassLoader public static ClassLoader getSystemClassLoader() public ClassLoader getParent()</pre>	checkPermission	If the caller's class loader is null, or is the same as or an ancestor of the class loader for the class whose class loader is being requested, no permission is needed. Otherwise, java.lang.RuntimePermission "getClassLoader" is required.
<pre>java.lang.Runtime public Process exec(String command) public Process exec(String command, String envp[]) public Process</pre>	checkExec	<pre>java.io.FilePermissi on "{command}", "execute"</pre>
<pre>exec(String cmdarray[]) public Process exec(String cmdarray[], String envp[])</pre>		
java.lang.Runtime public void exit(int status) public static void	checkExit(status) where status is 0 for runFinalizersOnExit	<pre>java.lang.RuntimePermissio n "exitVM.{status}"</pre>
runFinalizersOnExit(boolean value) java.lang.System public static void exit(int status) public static void		
runFinalizersOnExit(boolean value)		
<pre>java.lang.Runtime public void addShutdownHook(Thread hook) public boolean removeShutdownHook(Thread hook)</pre>	checkPermission	java.lang.RuntimePermission "shutdownHooks"



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.lang.Runtime public void load(String lib) public void loadLibrary(String lib) java.lang.System public static void load(String filename)</pre>	checkLink({libName}) where {libName} is the lib, filename or libname argument	java.lang.RuntimePermission "loadLibrary.{libName}"
public static void loadLibrary(
String libname)		
java.lang.SecurityManager methods	checkPermission	See java.lang.SecurityManager Method Permission Checks.
<pre>java.lang.System public static Properties getProperties() public static void</pre>	checkPropertiesAccess	<pre>java.util.PropertyPermissi on "*", "read,write"</pre>
setProperties(Properties props)		
<pre>java.lang.System public static String getProperty(String key) public static String getProperty(String key, String def)</pre>	checkPropertyAccess	java.util.PropertyPermissi on "{key}", "read"
<pre>java.lang.System public static void setIn(InputStream in) public static void setOut(PrintStream out) public static void setErr(PrintStream err)</pre>	checkPermission	java.lang.RuntimePermission "setIO"
<pre>java.lang.System public static String setProperty(String key, String value)</pre>	checkPermission	<pre>java.util.PropertyPermissi on "{key}", "write"</pre>



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
java.lang.System public static synchronized void	checkPermission	java.lang.RuntimePermission "setSecurityManager"
<pre>setSecurityManager(Security Manager s)</pre>		
<pre>java.lang.Thread public ClassLoader getContextClassLoader()</pre>	checkPermission	If the caller's class loader is null, or is the same as or an ancestor of the context class loader for the thread whose context class loader is being requested, no permission is needed. Otherwise, java.lang.RuntimePermission "getClassLoader" is required.
java.lang.Thread public void setContextClassLoader	checkPermission	java.lang.RuntimePermissio n "setContextClassLoader"
(ClassLoader cl)		
<pre>java.lang.Thread public final void checkAccess() public void interrupt() public final void suspend() public final void resume() public final void setPriority</pre>	checkAccess(this)	java.lang.RuntimePermissio n "modifyThread"
<pre>java.lang.Thread public static int enumerate(Thread tarray[])</pre>	checkAccess({threadGroup})	java.lang.RuntimePermission "modifyThreadGroup"



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.lang.Thread public final void stop()</pre>	checkAccess(this). Also checkPermission if the current thread is trying to stop a thread other than itself.	java.lang.RuntimePermissio n "modifyThread" Also java.lang.RuntimePermissio n "stopThread" if the current thread is trying to stop a thread other than itself.
<pre>java.lang.Thread public final synchronized void stop(Throwable obj)</pre>	checkAccess(this). Also checkPermission if the current thread is trying to stop a thread other than itself or obj is not an instance of ThreadDeath.	java.lang.RuntimePermissio n "modifyThread". Also java.lang.RuntimePermissio n "stopThread" if the current thread is trying to stop a thread other than itself or obj is not an instance of ThreadDeath.
<pre>java.lang.Thread Thread() Thread(Runnable target) Thread(String name) Thread(Runnable target, String name) java.lang.ThreadGroup</pre>	checkAccess({parentThreadGr oup})	java.lang.RuntimePermissio n "modifyThreadGroup"
ThreadGroup(String name) ThreadGroup(ThreadGroup parent, String name)		



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.lang.Thread Thread(ThreadGroup group,)</pre>	checkAccess(this) for ThreadGroup methods, or checkAccess(group) for Thread methods	java.lang.RuntimePermissio n "modifyThreadGroup"
<pre>java.lang.ThreadGroup public final void checkAccess() public int enumerate(Thread list[]) public int enumerate(Thread list[], boolean recurse) public int enumerate(ThreadGroup list[]) public int enumerate(ThreadGroup list[], boolean recurse) public final ThreadGroup getParent() public final void setDaemon(boolean daemon) public final void setMaxPriority(int pri) public final void suspend() public final void resume() public final void destroy()</pre>		
<pre>java.lang.ThreadGroup public final void interrupt()</pre>	checkAccess(this)	Requires java.lang.RuntimePermissio n "modifyThreadGroup". Also requires java.lang.RuntimePermissio n "modifyThread", since the java.lang.Thread interrupt() method is called for each thread in the thread group and in all of its subgroups. See the Thread interrupt() method.



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
java.lang.ThreadGroup public final void stop()	checkAccess(this)	Requires java.lang.RuntimePermissio n "modifyThreadGroup". Also requires java.lang.RuntimePermissio n "modifyThread" and possibly java.lang.RuntimePermissio n "stopThread", since the java.lang.Thread stop() method is called for each thread in the thread group and in all of its subgroups. See the Thread stop() method.
<pre>java.lang.reflect.Accessib eObject public static void setAccessible() public void setAccessible()</pre>	checkPermission	java.lang.reflect.ReflectP ermission "suppressAccessChecks"
java.net.Authenticator public static PasswordAuthentication	checkPermission	<pre>java.net.NetPermission "requestPasswordAuthentica tion"</pre>
<pre>requestPasswordAuthenticat on(</pre>		
String scheme java.net.Authenticator public static void		java.net.NetPermission "setDefaultAuthenticator"
setDefault(Authenticator a)	



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
java.net.MulticastSocket public void	checkMulticast(InetAddress)	<pre>java.net.SocketPermission(mcastaddr.getHostAddress() , "accept,connect")</pre>
<pre>joinGroup(InetAddress mcastaddr) public void</pre>		
<pre>leaveGroup(InetAddress mcastaddr)</pre>		
<pre>java.net.DatagramSocket public void send(DatagramPacket p)</pre>	checkMulticast(p.getAddress()) or checkConnect(p.getAddress() .getHostAddress(), p.getPort())	<pre>if (p.getAddress().isMulticast Address()) { java.net.SocketPermissio n((p.getAddress()).getHost Address(), "accept,connect") } else {port = p.getPort(); host = p.getAddress().getHostAddre ss(); if (port == -1) java.net.SocketPermission "{host}","resolve"; else java.net.SocketPermission "{host}:{port}","connect"</pre>
java.net.MulticastSocket public synchronized void send(DatagramPacket p, byte ttl)	checkMulticast(p.getAddress(), ttl) or checkConnect(p.getAddress().getHostAddress(), p.getPort())	<pre>if (p.getAddress().isMulticast Address()) { java.net.SocketPermissio n((p.getAddress()).getHost Address(), "accept,connect") } else { port = p.getPort(); host = p.getAddress().getHostAddre ss(); if (port == -1) java.net.SocketPermission "{host}","resolve"; else java.net.SocketPermission "{host}:{port}","connect" }</pre>



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.net.InetAddress public String getHostName() public static InetAddress[]</pre>	checkConnect({host}, -1)	<pre>java.net.SocketPermission "{host}", "resolve"</pre>
<pre>getAllByName(String host) public static InetAddress getLocalHost()</pre>		
<pre>java.net.DatagramSocket public InetAddress getLocalAddress()</pre>		
<pre>java.net.ServerSocket ServerSocket()</pre>	checkListen({port})	<pre>java.net.SocketPermission "localhost: {port}","listen";</pre>
<pre>java.net.DatagramSocket DatagramSocket()</pre>		
<pre>java.net.MulticastSocket MulticastSocket()</pre>		
<pre>java.net.ServerSocket public Socket accept() protected final void implAccept(Socket s)</pre>	checkAccept({host}, {port})	<pre>java.net.SocketPermission "{host}:{port}", "accept"</pre>



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
	checkSetFactory	java.lang.RuntimePermissio
java.net.ServerSocket		n "setFactory"
public static synchronized void		
setSocketFactory()		
java.net.Socket		
public static		
synchronized void		
<pre>setSocketImplFactory()</pre>		
java.net.URL		
public static		
synchronized void		
$\verb setURLStreamHandlerFactory \\ \dots)$		
java.net.URLConnection		
public static		
synchronized void		
setContentHandlerFactory(
.)		
public static void		
<pre>setFileNameMap(FileNameMap map)</pre>		
<pre>java.net.HttpURLConnection public static void</pre>		
<pre>setFollowRedirects(boolean set)</pre>		
java.rmi.activation.Activat		
ionGroup		
public static		
synchronized ActivationGroup		
createGroup()		
public static		
synchronized void		
<pre>setSystem(ActivationSystem system)</pre>		
java.rmi.server.RMISocketFa		
ctory public synchronized		
static void		
<pre>setSocketFactory()</pre>		



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.net.Socket Socket()</pre>	checkConnect({host}, {port})	<pre>java.net.SocketPermission "{host}:{port}", "connect"</pre>
java.net.DatagramSocket public synchronized void	checkAccept({host}, {port})	<pre>java.net.SocketPermission "{host}:{port}", "accept"</pre>
receive(DatagramPacket p)		
<pre>java.net.URL URL()</pre>	checkPermission	java.net.NetPermission "specifyStreamHandler"
<pre>java.net.URLClassLoader URLClassLoader()</pre>	checkCreateClassLoader	java.lang.RuntimePermissio n "createClassLoader"
<pre>java.security.AccessControl Context public AccessControlContext(Access ControlContext acc,</pre>	checkPermission	<pre>java.security.SecurityPerm ission "createAccessControlContex t"</pre>
DomainCombiner combiner) public DomainCombiner getDomainCombiner()		
<pre>java.security.Identity public void addCertificate()</pre>	checkSecurityAccess("addIde ntityCertificate")	<pre>java.security.SecurityPerm ission "addIdentityCertificate"</pre>
<pre>java.security.Identity public void removeCertificate()</pre>	checkSecurityAccess("remov eldentityCertificate")	<pre>java.security.SecurityPerm ission "removeIdentityCertificate "</pre>
<pre>java.security.Identity public void setInfo(String info)</pre>	checkSecurityAccess("setIde ntityInfo")	<pre>java.security.SecurityPerm ission "setIdentityInfo"</pre>



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.security.Identity public void setPublicKey(PublicKey key)</pre>	checkSecurityAccess("setIde ntityPublicKey")	java.security.SecurityPerm ission "setIdentityPublicKey"
<pre>java.security.Identity public String toString()</pre>	checkSecurityAccess("printld entity")	java.security.SecurityPerm ission "printIdentity"
<pre>java.security.IdentityScope protected static void setSystemScope()</pre>	checkSecurityAccess("setSys temScope")	<pre>java.security.SecurityPerm ission "setSystemScope"</pre>
<pre>java.security.Permission public void checkGuard(Object object)</pre>	checkPermission(this)	this Permission object is the permission checked
<pre>java.security.Policy public static Policy getPolicy()</pre>	checkPermission	java.security.SecurityPermission "getPolicy"
<pre>java.security.Policy public static void setPolicy(Policy policy)</pre>	checkPermission	java.security.SecurityPerm ission "setPolicy"
<pre>java.security.Policy public static Policy getInstance(String type, SpiParameter params) getInstance(String type, SpiParameter params, String provider) getInstance(String type, SpiParameter params, Provider provider)</pre>	checkPermission	<pre>java.security.SecurityPerm ission "createPolicy. {type}"</pre>
<pre>java.security.Provider public synchronized void clear()</pre>	checkSecurityAccess("clearPr oviderProperties."+{name})	java.security.SecurityPermission "clearProviderProperties. {name}" where name is the provider name.



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.security.Provider public synchronized Object put(Object key, Object value)</pre>	checkSecurityAccess("putPro viderProperty."+{name})	java.security.SecurityPerm ission "putProviderProperty. {name}" where name is the provider name.
<pre>java.security.Provider public synchronized Object remove(Object key)</pre>	checkSecurityAccess("remov eProviderProperty."+{name})	java.security.SecurityPerm ission "removeProviderProperty. {name}" where name is the provider name.
<pre>java.security.SecureClassLo ader SecureClassLoader()</pre>	checkCreateClassLoader	java.lang.RuntimePermissio n "createClassLoader"
<pre>java.security.Security public static void getProperty(String key)</pre>	checkPermission	<pre>java.security.SecurityPerm ission "getProperty.{key}"</pre>
<pre>java.security.Security public static int addProvider(Provider provider) public static int</pre>	checkSecurityAccess("insertP rovider."+provider.getName())	<pre>java.security.SecurityPerm ission "insertProvider. {name}"</pre>
<pre>insertProviderAt(Provider provider,</pre>		
<pre>java.security.Security public static void removeProvider(String name)</pre>	checkSecurityAccess("remov eProvider."+name)	<pre>java.security.SecurityPerm ission "removeProvider. {name}"</pre>
<pre>java.security.Security public static void setProperty(String key, String datum)</pre>	checkSecurityAccess("setPro perty."+key)	<pre>java.security.SecurityPerm ission "setProperty.{key}"</pre>



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>java.security.Signer public PrivateKey getPrivateKey()</pre>	checkSecurityAccess("getSig nerPrivateKey")	java.security.SecurityPerm ission "getSignerPrivateKey"
<pre>java.security.Signer public final void setKeyPair(KeyPair pair)</pre>	checkSecurityAccess("setSig nerKeypair")	java.security.SecurityPerm ission "setSignerKeypair"
	checkPermission	java.sql.SQLPermission
<pre>java.sql.DriverManager public static synchronized void</pre>		"setLog"
<pre>setLogWriter(PrintWriter out)</pre>		
	checkPermission	java.sql.SQLPermission
<pre>java.sql.DriverManager public static synchronized void</pre>		"setLog"
<pre>setLogStream(PrintWriter out)</pre>		
	checkPermission	java.util.PropertyPermissi
<pre>java.util.Locale public static synchronized void</pre>		on "user.language","write"
<pre>setDefault(Locale newLocale)</pre>		
	checkRead	<pre>java.io.FilePermission "{name}","read"</pre>
<pre>java.util.zip.ZipFile ZipFile(String name)</pre>		
	checkPermission	javax.security.auth.AuthPe
<pre>javax.security.auth.Subject public static Subject getSubject(final AccessControlContext acc)</pre>		rmission "getSubject"
<pre>javax.security.auth.Subject public void setReadOnly()</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "setReadOnly"</pre>



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>javax.security.auth.Subject public static Object doAs(final Subject subject, final PrivilegedAction action)</pre>	checkPermission	javax.security.auth.AuthPermission "doAs"
javax.security.auth.Subject public static Object doAs(final Subject subject, final PrivilegedExceptionAction action) throws java.security.PrivilegedAct ionException	checkPermission	javax.security.auth.AuthPermission "doAs"
<pre>javax.security.auth.Subject public static Object doAsPrivileged(final Subject subject, final PrivilegedAction action, final AccessControlContext acc)</pre>	checkPermission	javax.security.auth.AuthPermission "doAsPrivileged"
<pre>javax.security.auth.Subject public static Object doAsPrivileged(final Subject subject, final PrivilegedExceptionAction action, final AccessControlContext acc) throws java.security.PrivilegedAct ionException</pre>	checkPermission	javax.security.auth.AuthPermission "doAsPrivileged"



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>javax.security.auth.Subject DomainCombiner public Subject getSubject()</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "getSubjectFromDomainCombi ner"</pre>
<pre>javax.security.auth.Subject DomainCombiner public Subject getSubject()</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "getSubjectFromDomainCombi ner"</pre>
<pre>javax.security.auth.login.L oginContext public LoginContext(String name) throws LoginException</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "createLoginContext. {name}"</pre>
<pre>javax.security.auth.login.L oginContext public LoginContext(String name, Subject subject) throws LoginException</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "createLoginContext. {name}"</pre>
<pre>javax.security.auth.login.L oginContext public LoginContext(String name, CallbackHandler callbackHandler) throws LoginException</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "createLoginContext. {name}"</pre>



Table 1-6 (Cont.) Methods and the Permissions

Method	SecurityManager Method	Permission
<pre>javax.security.auth.login.L oginContext public LoginContext(String name, Subject subject, CallbackHandler</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "createLoginContext. {name}"</pre>
callbackHandler) throws LoginException		
<pre>javax.security.auth.login.C onfiguration public static Configuration getConfiguration()</pre>	checkPermission	javax.security.auth.AuthPermission "getLoginConfiguration"
<pre>javax.security.auth.login.C onfiguration public static void setConfiguration(Configurat ion configuration)</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "setLoginConfiguration"</pre>
<pre>javax.security.auth.login.C onfiguration public static void refresh()</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "refreshLoginConfiguration"</pre>
<pre>javax.security.auth.login.C onfiguration public static Configuration getInstance(String type, SpiParameter params) getInstance(String type, SpiParameter params, String provider) getInstance(String type, SpiParameter params, Provider provider)</pre>	checkPermission	<pre>javax.security.auth.AuthPe rmission "createLoginConfiguration. {type}"</pre>



java.lang.SecurityManager Method Permission Checks

List of permissions that are checked for by the default implementations of the java.lang.SecurityManager methods.

Each of the specified check methods calls the SecurityManager checkPermission method with the specified permission, except for the checkConnect and checkRead methods that take a context argument. Those methods expect the context to be an AccessControlContext and they call the context's checkPermission method with the specified permission.

Table 1-7 java.lang.SecurityManager Methods and Permissions

	Method	Permission
	public void checkAccept(String host, int port);	java.net.SocketPermission "{host}:{port}", "accept";
	<pre>public void checkAccess(Thread t);</pre>	java.lang.RuntimePermission "modifyThread";
	<pre>public void checkAccess(ThreadGroup g);</pre>	java.lang.RuntimePermission "modifyThreadGroup";
	<pre>public void checkAwtEventQueueAccess();</pre>	java.awt.AWTPermission "accessEventQueue";
	Note: This method is deprecated; use instead public void checkPermission(Permission perm);	
٠	public void checkConnect(String host, int port);	<pre>if (port == -1) java.net.SocketPermission "{host}","resolve"; else java.net.SocketPermission "{host}:</pre>
	<pre>public void checkConnect(String host, int port, Object context);</pre>	<pre>if (port == -1) java.net.SocketPermission "{host}","resolve"; else java.net.SocketPermission "{host}:</pre>
	public void checkCreateClassLoader();	java.lang.RuntimePermission "createClassLoader";
İ	public void checkDelete(String file);	java.io.FilePermission "{file}", "delete";
	public void checkExec(String cmd);	if cmd is an absolute path: java.io.FilePermission "{cmd}", "execute"; else java.io.FilePermission "< <all_files>>", "execute";</all_files>
	public void checkExit(int status);	java.lang.RuntimePermission "exitVM. {status}";
	public void checkLink(String lib);	java.lang.RuntimePermission "loadLibrary. {lib}";
	public void checkListen(int port);	java.net.SocketPermission "localhost: {port}","listen";



Table 1-7 (Cont.) java.lang.SecurityManager Methods and Permissions

	Method	Permission	
, de la	public void checkMemberAccess(Class clazz, int which); Note: This method is deprecated; use instead public void checkPermission(Permission perm);	<pre>if (which != Member.PUBLIC) { if (currentClassLoader() != clazz.getClassLoader()) { checkPermission(new java.lang.RuntimePermission("accessDeclaredMembers")); } }</pre>	
	public void checkMulticast(InetAddress maddr); public void checkMulticast(InetAddress maddr, byte ttl);	java.net.SocketPermission(maddr.getHostAdd ress(),"accept,connect"); java.net.SocketPermission(maddr.getHostAdd ress(),"accept,connect");	
A.M.	Note: This method is deprecated; use instead public void checkPermission(Permission perm);		
	public void checkPackageAccess(String pkg);	java.lang.RuntimePermission "accessClassInPackage.{pkg}";	
	public void checkPackageDefinition(String pkg);	java.lang.RuntimePermission "defineClassInPackage.{pkg}";	
	public void checkPrintJobAccess();	java.lang.RuntimePermission "queuePrintJob";	
	public void checkPropertiesAccess(); public void checkPropertyAccess(String key);	java.util.PropertyPermission "*", "read,write"; java.util.PropertyPermission "{key}", "read,write";	
	public void checkRead(FileDescriptor fd);	java.lang.RuntimePermission "readFileDescriptor";	
	public void checkRead(String file);	java.io.FilePermission "{file}", "read";	
	public void checkRead(String file, Object context);	java.io.FilePermission "{file}", "read";	
	public void checkSecurityAccess(String target);	java.security.SecurityPermission "{target}";	
	<pre>public void checkSetFactory();</pre>	java.lang.RuntimePermission "setFactory";	



Table 1-7 (Cont.) java.lang.SecurityManager Methods and Permissions

	Method	Permission
	public void checkSystemClipboardAccess();	java.awt.AWTPermission "accessClipboard";
	Note: This method is deprecated; use instead public void checkPermission(Permission perm);	
	public boolean checkTopLevelWindow(Object window);	java.awt.AWTPermission "showWindowWithoutWarningBanner";
***	Note: This method is deprecated; use instead public void checkPermission(Permission perm);	
	public void checkWrite(FileDescriptor fd);	java.lang.RuntimePermission "writeFileDescriptor";
	public void checkWrite(String file);	java.io.FilePermission "{file}", "write";
	public SecurityManager();	java.lang.RuntimePermission "createSecurityManager";

Default Policy Implementation and Policy File Syntax

The policy for a Java programming language application environment (specifying which permissions are available for code from various sources, and executing as various principals) is represented by a Policy object. More specifically, it is represented by a Policy subclass providing an implementation of the abstract methods in the Policy class (which is in the java.security package).

The source location for the policy information utilized by the Policy object is up to the Policy implementation. The Policy reference implementation obtains its information from static policy configuration files.

The rest of this document pertains to the Policy reference implementation and the syntax that must be used in policy files it reads:

- Default Policy Implementation
- Default Policy File Locations
- Modifying the Policy Implementation
- Policy File Syntax
- Policy File Examples



- · Property Expansion in Policy Files
- Windows Systems, File Paths, and Property Expansion
- General Expansion in Policy Files

Default Policy Implementation

Compose a policy file with any text editor.

In the Policy reference implementation, the policy can be specified within one or more policy configuration files. The configuration file(s) specify what permissions are allowed for code from a specified code source, and executed by a specified principal. Each configuration file must be encoded in UTF-8.

There is by default a single system-wide policy file, and a single (optional) user policy file. By default, permissions required by JDK modules that are loaded by the platform class loader or its ancestors are always granted.

The Policy reference implementation is initialized the first time its <code>getPermissions</code> method is called, or whenever its <code>refresh</code> method is called. Initialization involves parsing the policy configuration file(s) (see Policy File Syntax), and then populating the <code>Policy</code> object.

Default Policy File Locations

There is by default a single system-wide policy file, and a single (optional) user policy file. When the Policy is initialized, the system policy is loaded in first, and then the user policy is added to it. If neither policy is present, a built-in policy is used. This built-in policy is the same as the <code>java.policy</code> file installed with the JRE.

System Policy File Locations

By default, the system policy file is < java-home > / conf/security/java.policy.

The system policy file is meant to grant system-wide code permissions. The <code>java.policy</code> file installed with the JDK allows anyone to listen on dynamic ports, and allows any code to read certain "standard" properties that are not security-sensitive, such as the <code>os.name</code> and <code>file.separator</code> properties.

User Policy File Location

By default, the user policy file is <user-home>/.java.policy.

Configure Policy File Location and Format

Policy file locations are specified in the security properties file < java-home > / conf / security/java.security.

The policy file locations are specified as the values of properties whose names are of the form

policy.url.n

Here, n is a number. You specify each such property value in a line of the following form:

policy.url.n=URL



Here, URL is a URL specification. For example, the default system and user policy files are defined in the security properties file as:

```
policy.url.1=file:${java.home}/conf/security/java.policy
policy.url.2=file:${user.home}/.java.policy
```

(See Property Expansion for information about specifying property values via a special syntax, such as specifying the <code>java.home</code> property value via \${java.home}.)

You can actually specify a number of URLs (including ones of the form "http://"), and all the designated policy files will get loaded. You can also comment out or change the second one to disable reading the default user policy file.

The algorithm starts at policy.url.1, and keeps incrementing until it does not find a URL. Thus if you have policy.url.1 and policy.url.3, and policy.url.3 will never be read.

Example 1-4 Specifying an Additional Policy File at Runtime

It is also possible to specify an additional or a different policy file when invoking execution of an application. This can be done via the <code>-Djava.security.policy</code> command line argument, which sets the value of the <code>java.security.policy</code> property. For example, if you use following command, where <code>someURL</code> is a URL specifying the location of a policy file, then the specified policy file will be loaded in addition to all the policy files that are specified in the security properties file.

```
java -Djava.security.manager -Djava.security.policy=someURL SomeApp
```

The URL can be any regular URL or simply the name of a policy file in the current directory, as in:

```
java -Djava.security.manager -Djava.security.policy=mypolicy SomeApp
```

The -Djava.security.manager option ensures that the default security manager is installed, and thus the application is subject to policy checks. It is not required if the application SomeApp installs a security manager.

If you use the following command (note the double equals) then *just* the specified policy file will be used; all the ones indicated in the security properties file will be ignored.

java -Djava.security.manager -Djava.security.policy==someURL SomeApp



The policy file value of the -Djava.security.policy option is ignored if the policy.allowSystemProperty property in the security properties file is set to false. The default is true.

Modifying the Policy Implementation

The Policy reference implementation can be modified by editing the security properties file, which is the <code>java.security</code> file in the <code>conf/security</code> directory of the JDK.

An alternative policy class can be given to replace the Policy reference implementation class, as long as the former is a subclass of the abstract Policy class and implements the getPermissions method (and other methods as necessary).



One of the types of properties you can set in java.security is of the following form:

```
policy.provider=PolicyClassName
```

PolicyClassName must specify the fully qualified name of the desired Policy implementation class.

The default security properties file entry for this property is the following:

```
policy.provider=sun.security.provider.PolicyFile
```

To customize, you can change the property value to specify another class, as in

```
policy.provider=com.mycom.MyPolicy
```

Policy File Syntax

The policy configuration file(s) for a JDK installation specifies what permissions (which types of system resource accesses) are granted to code from a specified code source, and executed as a specified principal.

For an applet (or an application running under a security manager) to be allowed to perform secured actions (such as reading or writing a file), the applet (or application) must be granted permission for that particular action. In the Policy reference implementation, that permission must be granted by a grant entry in a policy configuration file. See below and the Java Security Architecture Specification for more information. (The only exception is that code always automatically has permission to read files from its same (URL) location, and subdirectories of that location; it does not need explicit permission to do so.)

A policy configuration file essentially contains a list of entries. It may contain a "keystore" entry, and contains zero or more "grant" entries.

Keystore Entry

The keytool utility is used to create and administer keystores.

A *keystore* is a database of private keys and their associated digital certificates such as X.509 certificate chains authenticating the corresponding public keys. The keytool utility is used to create and administer keystores. The keystore specified in a policy configuration file is used to look up the public keys of the signers specified in the grant entries of the file. A keystore entry must appear in a policy configuration file if any grant entries specify signer aliases, or if any grant entries specify principal aliases.

At this time, there can be only one keystore/keystorePasswordURL entry in the policy file (other entries following the first one are ignored). This entry can appear anywhere outside the file's grant entries. It has the following syntax:

```
keystore "some_keystore_url", "keystore_type", "keystore_provider";
keystorePasswordURL "some password url";
```

Here,



some_keystore_url

Specify the URL location of the keystore.

some_password_url

Specify the URL location of the keystore password.

keystore type

Specify the keystore type.

keystore provider

Specify the keystore provider.

Note:

- The input stream from some_keystore_url is passed to the KeyStore.load method.
- If NONE is specified as the URL, then a null stream is passed to the KeyStore.load method. NONE should be specified in the URL if the KeyStore is not file-based. For example, if it resides on a hardware token device.
- The URL is relative to the policy file location. If the policy file is specified in the security properties file as:

```
policy.url.1=http://foo.example.com/fum/some.policy
```

and that policy file has an entry:

```
keystore ".keystore";
```

then the keystore will be loaded from:

```
http://foo.example.com/fum/.keystore
```

The URL can also be absolute.

A **keystore type** defines the storage and data format of the keystore information, and the algorithms used to protect private keys in the keystore and the integrity of the keystore itself. The default type is a proprietary keystore type named "PKCS12". Thus, if the keystore type is "PKCS12", it does not need to be specified in the keystore entry.

Grant Entries

Grant entry is used to specify which code you want to grant permissions.

Code being executed is always considered to come from a particular "code source" (represented by an object of type <code>codeSource</code>). The code source includes not only the location (URL) where the code originated from, but also a reference to the certificate(s) containing the public key(s) corresponding to the private key(s) used to sign the code. Certificates in a code source are referenced by symbolic alias names from the user's keystore. Code is also considered to be executed as a particular principal (represented by an object of type <code>Principal</code>), or group of principals.



Each **grant entry** includes one or more "permission entries" preceded by optional codeBase, signedBy, and principal name/value pairs that specify which code you want to grant the permissions. The basic format of a grant entry is the following:

```
grant signedBy "signer_names", codeBase "URL",
    principal principal_class_name "principal_name",
    principal principal_class_name "principal_name",
    ... {

    permission permission_class_name "target_name", "action",
        signedBy "signer_names";
    permission permission_class_name "target_name", "action",
        signedBy "signer_names";
    ...
};
```

All non-italicized items above must appear as is (although case doesn't matter and some are optional, as noted below). Italicized items represent variable values.

A grant entry must begin with the word grant.

The SignedBy, Principal, and CodeBase Fields

The signedBy, codeBase, and principal values are optional, and the order of these fields does not matter.

A signedBy value indicates the alias for a certificate stored in the keystore. The public key within that certificate is used to verify the digital signature on the code; you grant the permission(s) to code signed by the private key corresponding to the public key in the keystore entry specified by the alias.

The signedBy value can be a comma-separated list of multiple aliases. An example is "Adam,Eve,Charles", which means "signed by Adam and Eve and Charles"; the relationship is AND, not OR. To be more exact, a statement like "Code signed by Adam" means "Code in a class file contained in a JAR which is signed using the private key corresponding to the public key certificate in the keystore whose entry is aliased by Adam".

The signedBy field is optional in that, if it is omitted, it signifies "any signer". It doesn't matter whether the code is signed or not or by whom.

A principal value specifies a class_name/principal_name pair which must be present within the executing thread's principal set. The principal set is associated with the executing code by way of a Subject.

The principal_class_name may be set to the wildcard value, *, which allows it to match any Principal class. In addition, the principal_name may also be set to the wildcard value, *, allowing it to match any Principal name. When setting the principal_class_name or principal_name to *, do not surround the * with quotes. Also, if you specify a wildcard principal class, you must also specify a wildcard principal name.

The principal field is optional in that, if it is omitted, it signifies "any principals".

KeyStore Alias Replacement

The principal <code>class_name/principal_name pair</code> is specified as a single quoted string, it is treated as a keystore alias.



The keystore is consulted and queried (via the alias) for an X509 Certificate. If one is found, the principal class_name is automatically treated as

javax.security.auth.x500.X500Principal, and the principal_name is automatically treated as the subject distinguished name from the certificate. If an X509 Certificate mapping is not found, the entire grant entry is ignored.

A codeBase value indicates the code source location; you grant the permission(s) to code from that location. An empty codeBase entry signifies "any code"; it doesn't matter where the code originates from.



AcodeBase value is a URL and thus should always utilize slashes (never backslashes) as the directory separator, even when the code source is actually on a Windows system. Thus, if the source location for code on a Windows system is actually $C:\somepath\api\$, then the policy codeBase entry should look like:

```
grant codeBase "file:/C:/somepath/api/" {
    ...
};
```

The exact meaning of a codeBase value depends on the characters at the end. A codeBase with a trailing "/" matches all class files (not JAR files) in the specified directory. A codeBase with a trailing "/*" matches all files (both class and JAR files) contained in that directory. A codeBase with a trailing "/-" matches all files (both class and JAR files) in the directory and recursively all files in subdirectories contained in that directory. The following table illustrates the different cases:

Table 1-8 KeyStore Alias

Codebase URL of Downloaded Code	Codebase URL in Policy	Match?
www.example.com/people/ gong/	www.example.com/people/ gong	Yes
www.example.com/people/ gong/	www.example.com/people/ gong/	Yes
www.example.com/people/ gong/	www.example.com/people/ gong/*	Yes
www.example.com/people/ gong/	www.example.com/people/ gong/-	Yes
www.example.com/people/ gong/appl.jar	www.example.com/people/ gong/	No
www.example.com/people/ gong/appl.jar	www.example.com/people/ gong/-	Yes
www.example.com/people/ gong/appl.jar	www.example.com/people/ gong/*	Yes
www.example.com/people/ gong/appl.jar	www.example.com/people/-	Yes



Nο

Codebase URL of Downloaded Code	Codebase URL in Policy	Match?
www.example.com/people/ gong/appl.jar	www.example.com/people/*	No
www.example.com/people/	www.example.com/people/-	Yes

www.example.com/people/*

Table 1-8 (Cont.) KeyStore Alias

www.example.com/people/

The Permission Entries

gong/

gong/

A permission entry is specified in the order (permission, permission_class_name, "target_name", "action", and signedBy "signer_names

A **permission entry** must begin with the word permission. The word permission_class_name in the template above would actually be a specific permission type, Such as java.io.FilePermission Or java.lang.RuntimePermission.

The "action" is required for many permission types, such as <code>java.io.FilePermission</code> (where it specifies what type of file access is permitted). It is not required for categories such as <code>java.lang.RuntimePermission</code> where it is not necessary, you either have the permission specified by the "<code>target_name</code>" value following the <code>permission_class_name</code> or you don't.

The signedBy name/value pair for a permission entry is optional. If present, it indicates a signed permission. That is, the permission class itself must be signed by the given alias(es) in order for the permission to be granted. For example, suppose you have the following grant entry:

```
grant {
    permission Foo "foobar", signedBy "FooSoft";
};
```

Then this permission of type Foo is granted if the Foo.class permission was placed in a JAR file and the JAR file was signed by the private key corresponding to the public key in the certificate specified by the "FooSoft" alias, or if Foo.class is a system class, since system classes are not subject to policy restrictions.

Items that appear in a permission entry must appear in the specified order (permission, permission_class_name, "target_name", "action", and signedBy "signer_names"). An entry is terminated with a semicolon.

Case is unimportant for the identifiers (permission, signedBy, codeBase, etc.) but is significant for the *permission_class_name* or for any string that is passed in as a value.

File Path Specifications on Windows Systems

The file path specifications on Windows systems should include two backslashes for each actual single backslash.



Note:

When you are specifying a <code>java.io.FilePermission</code>, the "target_name" is a file path. On Windows systems, whenever you directly specify a file path in a string (but not in a codeBase URL), you need to include two backslashes for each actual single backslash in the path, as in

```
grant {
    permission java.io.FilePermission "C:\\users\\cathy\\foo.bat", "read";
};
```

The reason this is necessary is because the strings are processed by a tokenizer (java.io.StreamTokenizer), which allows "\" to be used as an escape string (for example, "\n" to indicate a new line) and which thus requires two backslashes to indicate a single backslash. After the tokenizer has processed the above file path string, converting double backslashes to single backslashes, the end result is

```
"C:\users\cathy\foo.bat"
```

Policy File Examples

Examples of policy configuration files, with different configuration of the codeBase and signedBy values. Examples of grant statements with different principal based entry and KeyStore values.

Example 1-5 Sample Policy Configuration File

An example of two entries in a policy configuration file is as follows:

```
// If the code is signed by "Duke", grant it read/write access to all
// files in /tmp:
grant signedBy "Duke" {
    permission java.io.FilePermission "/tmp/*", "read,write";
};

// Grant everyone the following permission:
grant {
    permission java.util.PropertyPermission "java.vendor", "read";
};
```

Example 1-6 Sample Policy Configuration File

The following code specifies that *only* code that satisfies the following conditions can call methods in the Security class to add or remove providers or to set Security Properties:

- The code was loaded from a signed JAR file that is in the "/home/sysadmin/"
 directory on the local file system.
- The signature can be verified using the public key referenced by the alias name "sysadmin" in the keystore.



```
grant signedBy "sysadmin", codeBase "file:/home/sysadmin/*" {
    permission java.security.SecurityPermission "Security.insertProvider.*";
    permission java.security.SecurityPermission "Security.removeProvider.*";
    permission java.security.SecurityPermission "Security.setProperty.*";
};
```

Example 1-7 Sample Where codeBase is Missing

Either component of the code source (or both) may be missing. An example where codeBase is missing:

```
grant signedBy "sysadmin" {
    permission java.security.SecurityPermission "Security.insertProvider.*";
    permission java.security.SecurityPermission "Security.removeProvider.*";
};
```

If this policy is in effect, code that comes in a JAR File signed by "sysadmin" can add/remove providers, regardless of where the JAR File originated from.

Example 1-8 Sample Without signedBy

```
grant codeBase "file:/home/sysadmin/-" {
   permission java.security.SecurityPermission "Security.insertProvider.*";
   permission java.security.SecurityPermission "Security.removeProvider.*";
};
```

In this case, code that comes from anywhere beneath the "/home/sysadmin/" directory on the local filesystem can add/remove providers. The code does not need to be signed.

Example 1-9 Sample Without codeBase or signedBy

```
grant {
    permission java.security.SecurityPermission "Security.insertProvider.*";
    permission java.security.SecurityPermission "Security.removeProvider.*";
};
```

Here, with both code source components missing, any code (regardless of where it originated from, or whether or not it is signed, or who signed it) can add/remove providers.

Example 1-10 Sample Executing As X500Principal

```
grant principal javax.security.auth.x500.X500Principal "cn=Alice" {
    permission java.io.FilePermission "/home/Alice", "read, write";
};
```

This permits any code executing as the X500Principal, "cn=Alice", permission to read and write to "home/Alice".

Example 1-11 Sample Executing As X500Principal Without a Distinguished Name

```
grant principal javax.security.auth.x500.X500Principal * {
    permission java.io.FilePermission "/tmp", "read, write";
};
```



This permits any code executing as an X500Principal (regardless of the distinguished name), permission to read and write to "/tmp".

Example 1-12 Grant Statement With CodeBase and X500Principal Information

This allows code downloaded from "www.games.example.com", signed by "Duke", and executed by "cn=Alice", permission to read and write into the "/tmp/games" directory.

Example 1-13 Grant Statement With KeyStore Alias

```
keystore "http://foo.example.com/blah/.keystore";
grant principal "alice" {
    permission java.io.FilePermission "/tmp/games", "read, write";
};
"alice" will be replaced by
    javax.security.auth.x500.X500Principal "cn=Alice"
```

assuming the X.509 certificate associated with the keystore alias, alice, has a subject distinguished name of "cn=Alice". This allows code executed by the X500Principal "cn=Alice" permission to read and write into the "/tmp/games" directory.

Property Expansion in Policy Files

Property expansion is possible in policy files and in the security properties file.

Property expansion is similar to expanding variables in a shell. That is, when a string like

```
${some.property}
```

appears in a policy file, or in the security properties file, it will be expanded to the value of the system property. For example,

```
permission java.io.FilePermission "${user.home}", "read";
```

will expand " $\{user.home\}$ " to use the value of the "user.home" system property. If that property's value is "/home/cathy", then the above is equivalent to

```
permission java.io.FilePermission "/home/cathy", "read";
```

In order to assist in platform-independent policy files, you can also use the special notation of " $\{f|\}$ ", which is a shortcut for $\{file.separator\}$ ". This allows things like

```
permission java.io.FilePermission "${user.home}${/}*", "read";
```



If the value of the "user.home" property is /home/cathy, and you are on Solaris, Linux, or macOS, the above gets converted to:

```
permission java.io.FilePermission "/home/cathy/*", "read";
```

If on the other hand the "user.home" value is $C:\users\cathy$ and you are on a Windows system, the above gets converted to:

```
permission java.io.FilePermission "C:\users\cathy\*", "read";
```

Also, as a special case, if you expand a property in a codebase, such as

```
grant codeBase "file:${java.home}/lib/ext/"
```

then any file.separator characters will be automatically converted to / 's. Thus on a Windows system, the above would get converted to

```
grant codeBase "file:C:/jdk1.4/lib/ext/"
```

even if "java.home" is set to c:\jdk1.4\ Thus you don't need to use \${/} in codebase strings (and you shouldn't). Property expansion takes place anywhere a double quoted string is allowed in the policy file. This includes the "signer_names", "URL", "target_name", and "action" fields. Whether or not property expansion is allowed is controlled by the value of the "policy.expandProperties" property in the security properties file. If the value of this property is true (the default), expansion is allowed.

Note:

You can't use nested properties; they will not work. For example,

```
"${user.${foo}}"
```

doesn't work, even if the "foo" property is set to "home". The reason is the property parser doesn't recognize nested properties; it simply looks for the first "\${", and then keeps looking until it finds the first "}" and tries to interpret the result (in this case, "\${user.\$foo}") as a property, but fails if there is no such property.



Note:

If a property can't be expanded in a grant entry, permission entry, or keystore entry, that entry is ignored. For example, if the system property "foo" is not defined and you have:

```
grant codeBase "${foo}" {
    permission ...;
    permission ...;
};
```

then all the permissions in this grant entry are ignored. If you have

```
grant {
    permission Foo "${foo}";
    permission Bar "barTarget";
};
```

then only the "permission Foo..." entry is ignored. And finally, if you have

```
keystore "${foo}";
```

then the keystore entry is ignored.

Windows Systems, File Paths, and Property Expansion

The file path specifications on Windows systems should include two backslashes for each actual single backslash.

In Windows systems, when you directly specify a file path in a string (but not in a codeBase URL), you need to include two backslashes for each actual single backslash in the path, as in

```
grant {
    permission java.io.FilePermission "C:\\users\\cathy\\foo.bat", "read";
};
```

This is because the strings are processed by a tokenizer (java.io.StreamTokenizer), which allows "\" to be used as an escape string (e.g., "\n" to indicate a new line) and which thus requires two backslashes to indicate a single backslash. After the tokenizer has processed the above file path string, converting double backslashes to single backslashes, the end result is

```
"C:\users\cathy\foo.bat"
```

Expansion of a property in a string takes place after the tokenizer has processed the string. Thus if you have the string

```
"${user.home}\\foo.bat"
```



then first the tokenizer processes the string, converting the double backslashes to a single backslash, and the result is

```
"${user.home}\foo.bat"
```

Then the \${user.home} property is expanded and the end result is

```
"C:\users\cathy\foo.bat"
```

assuming the "user.home" value is c:\users\cathy. Of course, for platform independence, it would be better if the string was initially specified without any explicit slashes, i.e., using the \${/} property instead, as in

```
"${user.home}${/}foo.bat"
```

Path-Name Canonicalization

A canonical path is a path that doesn't contain any links or shortcuts. Performing pathname canonicalization in FilePermission object can negatively affect performance.

Before JDK 9, path names were canonicalized when two FilePermission objects were compared. This allowed a program to access a file using a different name than the name that was granted to a FilePermission object in a policy file, as long as the object pointed to the same file. Because the canonicalization had to access the underlying file system, it could be quite slow.

In JDK 9, path-name canonicalization is disabled by default. This means two FilePermission objects aren't equal to each other if one uses an absolute path and the other uses a relative path, or one uses a symbolic link and the other uses a target, or one uses a Windows long name and the other uses a DOS-style 8.3 name. This is true even if they all point to the same file in the file system.

Therefore, if a pathname is granted to a FilePermission object in a policy file, then the program should also access that file using the same path-name style. For example, if the path name in the policy file is using a symbolic link, then the program should also use that symbolic link. Accessing the file with the target path name will fail the permission check.

Compatibility Layer

A compatibility layer has been added to ensure that granting a FilePermission object for a relative path will permit applications to access the file with an absolute path (and conversly). This works for the default Policy provider and the Limited doPrivileged calls.

For example, a FilePermission object on a file with a relative path name of "a" no longer implies a FilePermission object on the same file with an absolute path name as "/pwd/a" ("pwd" is the current working directory). Granting code a FilePermission object to read "a" allows that code to also read "/pwd/a" when a Security Manager is enabled.

The compatibility layer doesn't cover translations between symbolic links and targets, or Windows long names and DOS-style 8.3 names, or any other different name forms that can be canonicalized to the same name.



Customizing Path-Name Canonicalization

The system properties in Table 1-9 can be used to customize the FilePermission pathname canonicalization. See How to Specify a java.lang.System Property.

 Table 1-9
 System Properties to Customize Pathname Canonicalization

System Property	Default Value	Description
jdk.io.permissionsUseCanon icalPath	false	The system property can be used to enable or disable pathname canonicalization in the FilePermission object.
		 To disable FilePermission path- name canonicalization, set jdk.io.permissionsUseC anonicalPath=false. To enable FilePermission path- name canonicalization, set jdk.io.permissionsUseC anonicalPath=true.
jdk.security.filePermCompa t	false	The system property can be used to extend the compatibility layer to support third-party Policy implementations.
		 To disable the system property, set jdk.security.filePermC ompat=false.
		The FilePermission for a relative path will permit applications to access the file with an absolute path for the default Policy provider and the Limited doPrivileged method.
		• To extend the compatibility layer to support third-party Policy implementations, set jdk.security.filePermC ompat=true.
		The FilePermission for a relative path will permit applications to access the file with an absolute path for the default Policy provider, the Limited doPrivileged method, and for third-party Policy implementations.



General Expansion in Policy Files

The policy files can be expanded using two protocols self and alias forms of expansion in the policy files.

```
${{protocol:protocol_data}}
```

If such a string occurs in a permission name, then the value in *protocol* determines the exact type of expansion that should occur, and *protocol_data* may be empty, in which case the above string should simply take the form:

```
${{protocol}}
```

There are two protocols supported in the default policy file implementation:

1. \${{self}}

The protocol, self, denotes a replacement of the entire string, \${{self}}}, with one or more principal class/name pairs. The exact replacement performed depends upon the contents of the grant clause to which the permission belongs.

If the grant clause does not contain any principal information, the permission will be ignored (permissions containing $\{\{self\}\}\}$ in their target names are only valid in the context of a principal-based grant clause). For example, BarPermission will always be ignored in the following grant clause:

```
grant codebase "www.example.com", signedby "duke" {
    permission BarPermission "... ${{self}} ...";
};
```

If the grant clause contains principal information, \${{self}} will be replaced with that same principal information. For example, \${{self}} in BarPermission will be replaced with javax.security.auth.x500.X500Principal "cn=Duke" in the following grant clause:

```
grant principal javax.security.auth.x500.X500Principal "cn=Duke" {
    permission BarPermission "... ${{self}} ...";
};
```

If there is a comma-separated list of principals in the grant clause, then $\{\{self\}\}\}$ will be replaced by the same comma-separated list or principals. In the case where both the principal class and name are wildcarded in the grant clause, $\{\{self\}\}\}$ is replaced with all the principals associated with the Subject in the current AccessControlContext.

The following example describes a scenario involving both self and KeyStore Alias Replacement together:

```
keystore "http://foo.example.com/blah/.keystore";
grant principal "duke" {
    permission BarPermission "... ${{self}} ...";
};
```



In the above example, "duke" will first be expanded into javax.security.auth.x500.X500Principal "cn=Duke" assuming the X.509 certificate associated with the KeyStore alias, "duke", has a subject distinguished name of "cn=Duke". Next, \${{self}} will be replaced with the same principal information that was just expanded in the grant clause: javax.security.auth.x500.X500Principal "cn=Duke".

2. \${{alias:alias_name}}

The protocol, alias, denotes a java.security.KeyStore alias substitution. The KeyStore used is the one specified in the Keystore Entry. alias_name represents an alias into the KeyStore. \${{alias:alias_name}}} is replaced with javax.security.auth.x500.X500Principal "DN", where DN represents the subject distinguished name of the certificate belonging to alias name. For example:

```
keystore "http://foo.example.com/blah/.keystore";
grant codebase "www.example.com" {
    permission BarPermission "... ${{alias:duke}} ...";
};
```

In the above example the X.509 certificate associated with the alias, *duke*, is retrieved from the <code>KeyStore</code>, *foo.example.com/blah/.keystore*. Assuming duke's certificate specifies "o=dukeOrg, cn=duke" as the subject distinguished name, then \$ {alias:duke}} is replaced with <code>javax.security.auth.x500.X500Principal</code> "o=dukeOrg, cn=duke".

The permission entry is ignored under the following error conditions:

- · The keystore entry is unspecified
- The alias_name is not provided
- The certificate for alias_name can not be retrieved
- The certificate retrieved is not an X.509 certificate

API for Privileged Blocks

Background information about what *privileged* code is and what it is used for, followed by illustrations of the use of the API. It covers the following topics:

Using the doPrivileged API

What It Means to Have Privileged Code

Reflection

Using the doPrivileged API

Description of the doPrivileged API and the use of the *privileged* feature.

No Return Value, No Exception Thrown

Accessing Local Variables

Handling Exceptions

Asserting a Subset of Privileges

Least Privilege



More Privilege

No Return Value, No Exception Thrown

If you do not need to return a value from within the *privileged* block, your call to doPrivileged can look like Example 1-14.

Note that the invocation of doPrivileged with a lambda expression explicitly casts the lambda expression as of type PrivilegedAction<Void>. Another version of the method doPrivileged exists that takes an object of type PrivilegedExceptionAction; see Handling Exceptions.

PrivilegedAction is a functional interface with a single abstract method, named run, that returns a value of type specified by its type parameter.

Note that this example ignores the return value of the run method. Also, depending on what *privileged code* actually consists of, you might have to make some changes due to the way inner classes work. For example, if *privileged code* throws an exception or attempts to access local variables, then you will have to make some changes, which is described later.

Be very careful in your use of the privileged construct, and always remember to make the privileged code section as small as possible. That is, try to limit the code within the run method to only what needs to be run with privileges, and do more general things outside the run method. Also note that the call to doPrivileged should be made in the code that wants to enable its privileges. Do not be tempted to write a utility class that itself calls doPrivileged as that could lead to security holes. You can write utility classes for PrivilegedAction classes though, as shown in the preceding example. See Guideline 9-3: Safely invoke java.security.AccessController.doPrivileged in Secure Coding Guidelines for the Java Programming Language.

Example 1-14 Sample Code for Privileged Block

- In a class that implements the interface PrivilegedAction.
- In an anonymous class.
- In a lambda expression.



Returning Values

Sample code to return a value.

If you need to return a value, then you can do something like the following:

Accessing Local Variables

If you are using a lambda expression or anonymous inner class, then any local variables you access must be final or effectively final.

For example:

The variable 11b is effectively final because its value has not been modified. For example, suppose you add the following assignment statement after the declaration of the variable 11b:

```
lib = "swing";
```

The compiler generates the following errors when it encounters the invocation System.loadLibrary both in the lambda expression and the anonymous class:

- error: local variables referenced from a lambda expression must be final or effectively final
- error: local variables referenced from an inner class must be final or effectively final

See Accessing Members of an Enclosing Class in Local Classes for more information.

If there are cases where you cannot make an existing variable effectively final (because it gets set multiple times), then you can create a new final variable right before invoking the doPrivileged method, and set that variable equal to the other variable. For example:

Handling Exceptions

If the action performed in your run method could throw a *checked* exception (one that must be listed in the throws clause of a method), then you need to use the PrivilegedExceptionAction interface instead of the PrivilegedAction interface.

Example 1-15 Sample for Handling Exceptions

If a checked exception is thrown during execution of the run method, then it is placed in a PrivilegedActionException *wrapper* exception that is then thrown and should be caught by your code, as illustrated in the following example:



```
}
```

Asserting a Subset of Privileges

Variant of the doPrivileged has three parameters, one of which you use to specify the subset of privileges.

As of JDK 8, a variant of doPrivileged is available that enables code to assert a subset of its privileges, without preventing the full traversal of the stack to check for other permissions. This variant of the doPrivileged variant has three parameters, one of which you use to specify this subset of privileges. For example, the following excerpt asserts a privilege to retrieve system properties:

The first parameter of this version of doPrivileged is of type java.security.PrivilegedAction. In this example, the first parameter is a lambda expression that implements the functional interface PrivilegedAction whose run method returns the value of the system property specified by the parameter prop.

The second parameter of this version of doPrivileged is of type AccessControlContext. Sometimes you need to perform an additional security check within a different context, such as a worker thread. You can obtain an AccessControlContext instance from a particular calling context with the method AccessControlContext. If you specify null for this parameter (as in this example), then the invocation of doPrivileged does not perform any additional security checks.

The third parameter of this version of doPrivileged is of type Permission..., which is a varargs parameter. This means that you can specify one or more Permission parameters or an array of Permission objects, as in Permission[]. In this example, the invocation of doPrivileged can retrieve the properties app.version and app.vendor.

You can use this three parameter variant of <code>doPrivileged</code> in a mode of least privilege or a mode of more privilege.

Least Privilege

The typical use case of the <code>doPrivileged</code> method is to enable the method that invokes it to perform one or more actions that require permission checks without requiring the callers of the current method to have all the necessary permissions.

For example, the current method might need to open a file or make a network request for its own internal implementation purposes.

Before JDK 8, calls to <code>doPrivileged</code> methods had only two parameters. They worked by granting temporary privileges to the calling method and stopping the normal full traversal of the stack for access checking when it reached that class, rather than

continuing up the call stack where it might reach a method whose defining class does not have the required permission. Typically, the class that is calling <code>doPrivileged</code> might have additional permissions that are not required in that code path and which might also be missing from some caller classes.

Normally, these extra permissions are not exercised at runtime. Not elevating them through use of <code>doPrivileged</code> helps to block exploitation of any incorrect code that could perform unintended actions. This is especially true when the <code>PrivilegedAction</code> is more complex than usual, or when it calls code outside the class or package boundary that might evolve independently over time.

The three-parameter variant of <code>doPrivileged</code> is generally safer to use because it avoids unnecessarily elevating permissions that are not intended to be required. However, it executes less efficiently so simple or performance-critical code paths might choose not to use it.

More Privilege

When coding the current method, you want to temporarily extend the permission of the calling method to perform an action.

For example, a framework I/O API might have a general purpose method for opening files of a particular data format. This API would take a normal file path parameter and use it to open an underlying FileInputStream using the calling code's permissions. However, this might also allow any caller to open the data files in a special directory that contains some standard demonstration samples.

The callers of this API could be directly granted a FilePermission for *read* access. However, it might not be convenient or possible for the security policy of the calling code to be updated. For example, the calling code could be a sandboxed applet.

One way to implement this is for the code to check the incoming path and determine if it refers to a file in the special directory. If it does, then it would call <code>doPrivileged</code>, enabling all permissions, then open the file inside the <code>PrivilegedAction</code>. If the file was not in the special directory, the code would open the file without using <code>doPrivileged</code>.

This technique requires the implementation to carefully handle the requested file path to determine if it refers to the special shared directory. The file path must be canonicalized before calling <code>doPrivileged</code> so that any relative path will be processed (and permission to read the <code>user.dir</code> system property will be checked) prior to determining if the path refers to a file in the special directory. It must also prevent malicious "../" path elements meant to escape out of the special directory.

A simpler and better implementation would use the variant of <code>doPrivileged</code> with the third parameter. It would pass a <code>FilePermission</code> with <code>read</code> access to the special directory as the third parameter. Then any manipulation of the file would be inside the <code>PrivilegedAction</code>. This implementation is simpler and much less prone to contain a security flaw.

What It Means to Have Privileged Code

Marking code as *privileged* enables a piece of trusted code to temporarily enable access to more resources than are available directly to the code that called it.

The policy for a JDK installation specifies what permissions which types of system resource accesses — are allowed for code from specified code sources. A *code source* (of type <code>codeSource</code>) essentially consists of the code location (URL) and a



reference to the certificates containing the public keys corresponding to the private keys used to sign the code (if it was signed).

The policy is represented by a Policy object. More specifically, it is represented by a Policy subclass providing an implementation of the abstract methods in the Policy class (which is in the java.security package).

The source location for the policy information used by the Policy object depends on the Policy implementation. The Policy reference implementation obtains its information from policy configuration files. See Default Policy Implementation and Policy File Syntax for information about the Policy reference implementation and the syntax that must be used in policy files it reads. For information about using the **Policy Tool** to create a policy file (without needing to know the required syntax), see Policy Tool .

A *protection domain* encompasses a codeSource instance and the permissions granted to code from that CodeSource, as determined by the security policy currently in effect. Thus, classes signed by the same keys and from the same URL are typically placed in the same domain, and a class belongs to one and only one protection domain. (However, classes signed by the same keys and from the same URL but loaded by separate class loader instances are typically placed in separate domains.) Classes that have the same permissions but are from different code sources belong to different domains.

Each applet or application runs in its appropriate domain, determined by its code source. For an applet (or an application running under a security manager) to be allowed to perform a secured action (such as reading or writing a file), the applet or application must be granted permission for that particular action.

More specifically, whenever a resource access is attempted, *all* code traversed by the execution thread up to that point must have permission for that resource access, *unless some code on the thread has been marked as privileged*. That is, suppose that access control checking occurs in a thread of execution that has a chain of multiple callers. (Think of this as multiple method calls that potentially cross the protection domain boundaries.) When the AccessController.checkPermission method is invoked by the most recent caller, the basic algorithm for deciding whether to allow or deny the requested access is as follows: If the code for any caller in the call chain does not have the requested permission, then an AccessControlException is thrown, *unless* the following is true: a caller whose code is granted the said permission has been marked as *privileged*, and all parties subsequently called by this caller (directly or indirectly) have the said permission.



Note:

The method AccessController.checkPermission is normally invoked indirectly through invocations of specific SecurityManager methods that begin with the word check such as checkConnect or through the method SecurityManager.checkPermission. Normally, these checks only occur if a SecurityManager has been installed; code checked by the AccessController.checkPermission method first checks if the method System.getSecurityManager returns null.

Marking code as *privileged* enables a piece of trusted code to temporarily enable access to more resources than are available directly to the code that called it. This is necessary in some situations. For example, an application might not be allowed direct access to files that contain fonts, but the system utility to display a document must obtain those fonts, on behalf of the user. The system utility must become privileged in order to obtain the fonts.

Reflection

doPrivileged method can be invoked reflectively using java.lang.reflect.Method.invoke.

One subtlety that must be considered is the interaction of this API with reflection. The doPrivileged method can be invoked reflectively using java.lang.reflect.Method.invoke. In this case, the privileges granted in privileged mode are not those of Method.invoke but of the non-reflective code that invoked it. Otherwise, system privileges could erroneously (or maliciously) be conferred on user code. Note that similar requirements exist when using reflection in the existing API.

Troubleshooting Security

To monitor security access, you can set the <code>java.security.debug</code> system property, which determines what trace messages are printed during execution.

To see a list of all debugging options, use the help setting:

java -Djava.security.debug=help



To use more than one option, separate options with a comma.

JSSE also provides dynamic debug tracing support for SSL/TLS/DTLS troubleshooting. See Debugging Utilities.

The following table lists <code>java.security.debug</code> options and links to further information about each option:



Table 1-10 java.security.debug Options

Option	Description	Further Information
all	Turn on all the debugging options	None
access	Print all results from the AccessController.checkPermissio n method. You can use the following options with the access option:	Permissions in the Java Development Kit (JDK)
	1. stack: Include stack trace	
	2. domain: Dump all domains in context	
	3. failure: Before throwing exception, dump stack and domain that do not have permission	
	You can use the following options with the stack and domain options:	
	 permission=<classname>: Only dump output if specified permission is being checked</classname> 	
	codebase=<url>: Only dump output if specified codebase is being checked</url>	
certpath	Turns on debugging for the PKIX CertPathValidator and CertPathBuilder implementations. Use the ocsp option with the certpath option for OCSP protocol tracing. A hexadecimal dump of the OCSP request and response bytes is displayed.	PKI Programmers Guide Overview
combiner	SubjectDomainCombiner debugging	Permissions in the Java Development Kit (JDK)
configfile	JAAS (Java Authentication and Authorization Service) configuration file loading	Java Authentication and Authorization Service (JAAS) Reference Guide
		Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges
configparser	JAAS configuration file parsing	Java Authentication and Authorization Service (JAAS) Reference Guide
		Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges



Table 1-10 (Cont.) java.security.debug Options

Option	Description	Further Information
gssloginconfig	Java GSS (Generic Security Services) login configuration file debugging	Java Generic Security Services: (Java GSS) and Kerberos
		JAAS and Java GSS-API Tutorial
		javax.security.auth.login.Configuration: A Configuration object is responsible for specifying which javax.net.ssl.SSLEngine should be used for a particular application, and in what order the LoginModules should be invoked.
		JAAS Login Configuration File
		Advanced Security Programming in Java SE Authentication, Secure Communication and Single Sign-On
jar	JAR file verification	Verifying Signed JAR Files from The Java Tutorials
jca	JCA engine class debugging	Engine Classes and Algorithms
keystore	Keystore debugging	Keystores
		KeyStore
logincontext	LoginContext results	Java Authentication and Authorization Service (JAAS) Reference Guide
		Use of JAAS Login Utility and Java GSS-API for Secure Message Exchanges
pkcs11	PKCS11 session manager debugging	PKCS#11 Reference Guide
pkcs11keystore	PKCS11 KeyStore debugging	PKCS#11 Reference Guide
pkcs12	PKCS12 KeyStore debugging	None
policy	Loading and granting permissions with policy file	Set up the Policy File to Grant the Required Permissions (Controlling Applications) from The Java Tutorials Default Policy Implementation and Policy File Syntax



Table 1-10 (Cont.) java.security.debug Options

Option	Description	Further Information
provider	Security provider debugging The following options can be used with the provider option:	Java Cryptography Architecture (JCA) Reference Guide
	engine= <engines>: The output is displayed only for a specified list of JCA engines.</engines>	
	The supported values for <engines> are:</engines>	
	 Cipher 	
	 KeyAgreement 	
	 KeyGenerator 	
	 KeyPairGenerator 	
	 KeyStore 	
	• Mac	
	 MessageDigest 	
	 SecureRandom 	
	 Signature 	
scl	Permissions that SecureClassLoader assigns	Permissions in the Java Development Kit (JDK)
securerandom	SecureRandom debugging	The SecureRandom Class
sunpkcs11	SunPKCS11 provider debugging	PKCS#11 Reference Guide
ts	Timestamping debugging	None



Java Cryptography Architecture (JCA) Reference Guide

The Java Cryptography Architecture (JCA) is a major piece of the platform, and contains a "provider" architecture and a set of APIs for digital signatures, message digests (hashes), certificates and certificate validation, encryption (symmetric/asymmetric block/stream ciphers), key generation and management, and secure random number generation, to name a few.

Introduction to Java Cryptography Architecture

The Java platform strongly emphasizes security, including language safety, cryptography, public key infrastructure, authentication, secure communication, and access control.

The JCA is a major piece of the platform, and contains a "provider" architecture and a set of APIs for digital signatures, message digests (hashes), certificates and certificate validation, encryption (symmetric/asymmetric block/stream ciphers), key generation and management, and secure random number generation, to name a few. These APIs allow developers to easily integrate security into their application code. The architecture was designed around the following principles:

- Implementation independence: Applications do not need to implement security
 algorithms. Rather, they can request security services from the Java platform.
 Security services are implemented in providers (see Cryptographic Service
 Providers), which are plugged into the Java platform via a standard interface. An
 application may rely on multiple independent providers for security functionality.
- **Implementation interoperability**: Providers are interoperable across applications. Specifically, an application is not bound to a specific provider, and a provider is not bound to a specific application.
- Algorithm extensibility: The Java platform includes a number of built-in providers
 that implement a basic set of security services that are widely used today.
 However, some applications may rely on emerging standards not yet
 implemented, or on proprietary services. The Java platform supports the
 installation of custom providers that implement such services.

Other cryptographic communication libraries available in the JDK use the JCA provider architecture, but are described elsewhere. The JSSE components provides access to Secure Socket Layer (SSL), Transport Layer Security (TLS), and Datagram Transport Layer Security (DTLS) implementations; see Java Secure Socket Extension (JSSE) Reference Guide. You can use Java Generic Security Services (JGSS) (via Kerberos) APIs, and Simple Authentication and Security Layer (SASL) to securely exchange messages between communicating applications; see Java GSS-API and JAAS Tutorials for Use with Kerberos and Java SASL API Programming and Deployment Guide.



Notes on Terminology

- Prior to JDK 1.4, the JCE was an unbundled product, and as such, the JCA and JCE were regularly referred to as separate, distinct components. As JCE is now bundled in the JDK, the distinction is becoming less apparent. Since the JCE uses the same architecture as the JCA, the JCE should be more properly thought of as a part of the JCA.
- The JCA within the JDK includes two software components:
 - The framework that defines and supports cryptographic services for which providers supply implementations. This framework includes packages such as java.security, javax.crypto, javax.crypto.spec, and javax.crypto.interfaces.
 - The actual providers such as Sun, SunRsaSign, SunJCE, which contain the actual cryptographic implementations.

Whenever a specific JCA provider is mentioned, it will be referred to explicitly by the provider's name.

WARNING:

The JCA makes it easy to incorporate security features into your application. However, this document does not cover the theory of security/cryptography beyond an elementary introduction to concepts necessary to discuss the APIs. This document also does not cover the strengths/weaknesses of specific algorithms, not does it cover protocol design. Cryptography is an advanced topic and one should consult a solid, preferably recent, reference in order to make best use of these tools.

You should always understand what you are doing and why: DO NOT simply copy random code and expect it to fully solve your usage scenario. Many applications have been deployed that contain significant security or performance problems because the wrong tool or algorithm was selected.

JCA Design Principles

The JCA was designed around these principles:

- Implementation independence and interoperability
- Algorithm independence and extensibility

Implementation independence and algorithm independence are complementary; you can use cryptographic services, such as digital signatures and message digests, without worrying about the implementation details or even the algorithms that form the basis for these concepts. While complete algorithm-independence is not possible, the JCA provides standardized, algorithm-specific APIs. When implementationindependence is not desirable, the JCA lets developers indicate a specific implementation.

Algorithm independence is achieved by defining types of cryptographic "engines" (services), and defining classes that provide the functionality of these cryptographic engines. These classes are called *engine classes*, and examples are the MessageDigest, Signature, KeyFactory, KeyPairGenerator, and Cipher classes.



Implementation independence is achieved using a "provider"-based architecture. The term Cryptographic Service Provider (CSP), which is used interchangeably with the term "provider," (see Cryptographic Service Providers) refers to a package or set of packages that implement one or more cryptographic services, such as digital signature algorithms, message digest algorithms, and key conversion services. A program may simply request a particular type of object (such as a signature object) implementing a particular service (such as the DSA signature algorithm) and get an implementation from one of the installed providers. If desired, a program may instead request an implementation from a specific provider. Providers may be updated transparently to the application, for example when faster or more secure versions are available.

Implementation interoperability means that various implementations can work with each other, use each other's keys, or verify each other's signatures. This would mean, for example, that for the same algorithms, a key generated by one provider would be usable by another, and a signature generated by one provider would be verifiable by another.

Algorithm extensibility means that new algorithms that fit in one of the supported engine classes can be added easily.

Provider Architecture

Providers contain a package (or a set of packages) that supply concrete implementations for the advertised cryptographic algorithms.

Cryptographic Service Providers

<code>java.security.Provider</code> is the base class for all security providers. Each CSP contains an instance of this class which contains the provider's name and lists all of the security services/algorithms it implements. When an instance of a particular algorithm is needed, the JCA framework consults the provider's database, and if a suitable match is found, the instance is created.

Providers contain a package (or a set of packages) that supply concrete implementations for the advertised cryptographic algorithms. Each JDK installation has one or more providers installed and configured by default. Additional providers may be added statically or dynamically. Clients may configure their runtime environment to specify the provider *preference order*. The preference order is the order in which providers are searched for requested services when no specific provider is requested.

To use the JCA, an application simply requests a particular type of object (such as a MessageDigest) and a particular algorithm or service (such as the "SHA-256" algorithm), and gets an implementation from one of the installed providers. For example, the following statement requests a SHA-256 message digest from an installed provider:

```
md = MessageDigest.getInstance("SHA-256");
```

Alternatively, the program can request the objects from a specific provider. Each provider has a name used to refer to it. For example, the following statement requests a SHA-256 message digest from the provider named ProviderC:

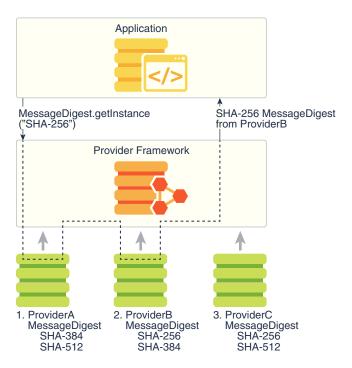
```
md = MessageDigest.getInstance("SHA-256", "ProviderC");
```

The following figures illustrates requesting an SHA-256 message digest implementation. They show three different providers that implement various message digest algorithms (SHA-256, SHA-384, and SHA-512). The providers are ordered by preference from left to right (1-3). In Figure 2-1, an application requests a SHA-256



algorithm implementation **without** specifying a provider name. The providers are searched in preference order and the implementation from the first provider supplying that particular algorithm, ProviderB, is returned. In Figure 2-2, the application requests the SHA-256 algorithm implementation **from a specific provider**, ProviderC. This time, the implementation from ProviderC is returned, even though a provider with a higher preference order, ProviderB, also supplies an MD5 implementation.

Figure 2-1 Request SHA-256 Message Digest Implementation Without Specifying Provider





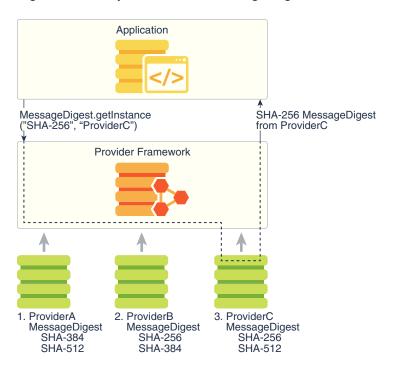


Figure 2-2 Request SHA-256 Message Digest with ProviderC

Cryptographic implementations in the JDK are distributed via several different providers (Sun, SunJSSE, SunJCE, SunRsaSign) primarily for historical reasons, but to a lesser extent by the type of functionality and algorithms they provide. Other Java runtime environments may not necessarily contain these providers, so applications should not request a provider-specific implementation unless it is known that a particular provider will be available.

The JCA offers a set of APIs that allow users to query which providers are installed and what services they support.

This architecture also makes it easy for end-users to add additional providers. Many third party provider implementations are already available. See The Provider Class for more information on how providers are written, installed, and registered.

How Providers Are Actually Implemented

Algorithm independence is achieved by defining a generic high-level Application Programming Interface (API) that all applications use to access a service type. Implementation independence is achieved by having all provider implementations conform to well-defined interfaces. Instances of engine classes are thus "backed" by implementation classes which have the same method signatures. Application calls are routed through the engine class and are delivered to the underlying backing implementation. The implementation handles the request and return the proper results.

The application API methods in each engine class are routed to the provider's implementations through classes that implement the corresponding Service Provider Interface (SPI). That is, for each engine class, there is a corresponding abstract SPI class which defines the methods that each cryptographic service provider's algorithm must implement. The name of each SPI class is the same as that of the corresponding engine class, followed by Spi. For example, the Signature engine class provides



access to the functionality of a digital signature algorithm. The actual provider implementation is supplied in a subclass of SignatureSpi. Applications call the engine class' API methods, which in turn call the SPI methods in the actual implementation.

Each SPI class is abstract. To supply the implementation of a particular type of service for a specific algorithm, a provider must subclass the corresponding SPI class and provide implementations for all the abstract methods.

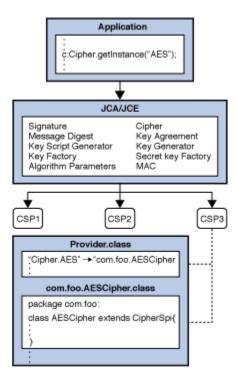
For each engine class in the API, implementation instances are requested and instantiated by calling the <code>getInstance()</code> factory method in the engine class. A factory method is a static method that returns an instance of a class. The engine classes use the framework provider selection mechanism described above to obtain the actual backing implementation (SPI), and then creates the actual engine object. Each instance of the engine class encapsulates (as a private field) the instance of the corresponding SPI class, known as the SPI object. All API methods of an API object are declared final and their implementations invoke the corresponding SPI methods of the encapsulated SPI object.

To make this clearer, review Example 2-1 and Figure 2-3:

Example 2-1 Sample Code for Getting an Instance of an Engine Class

```
import javax.crypto.*;
Cipher c = Cipher.getInstance("AES");
c.init(ENCRYPT_MODE, key);
```

Figure 2-3 Application Retrieves "AES" Cipher Instance



Here an application wants an "AES" <code>javax.crypto.Cipher</code> instance, and doesn't care which provider is used. The application calls the <code>getInstance()</code> factory methods of the <code>cipher</code> engine class, which in turn asks the JCA framework to find the first provider

instance that supports "AES". The framework consults each installed provider, and obtains the provider's instance of the Provider class. (Recall that the Provider class is a database of available algorithms.) The framework searches each provider, finally finding a suitable entry in CSP3. This database entry points to the implementation class com.foo.AESCipher which extends CipherSpi, and is thus suitable for use by the Cipher engine class. An instance of com.foo.AESCipher is created, and is encapsulated in a newly-created instance of javax.crypto.Cipher, which is returned to the application. When the application now does the init() operation on the Cipher instance, the Cipher engine class routes the request into the corresponding engineInit() backing method in the com.foo.AESCipher class.

Java Security Standard Algorithm Names Specification lists the Standard Names defined for the Java environment. Other third-party providers may define their own implementations of these services, or even additional services.

Keystores

A database called a "keystore" can be used to manage a repository of keys and certificates. Keystores are available to applications that need data for authentication, encryption, or signing purposes.

Applications can access a keystore via an implementation of the <code>KeyStore</code> class, which is in the <code>java.security</code> package. As of JDK 9, the default and recommended keystore type (format) is "pkcs12", which is based on the RSA PKCS12 Personal Information Exchange Syntax Standard. Previously, the default keystore type was "jks", which is a proprietary format. Other keystore formats are available, such as "jceks", which is an alternate proprietary keystore format, and "pkcs11", which is based on the RSA PKCS11 Standard and supports access to cryptographic tokens such as hardware security modules and smartcards.

Applications can choose different keystore implementations from different providers, using the same provider mechanism described previously. See Key Management.

Engine Classes and Algorithms

An engine class provides the interface to a specific type of cryptographic service, independent of a particular cryptographic algorithm or provider.

The engines provides one of the following:

- cryptographic operations (encryption, digital signatures, message digests, etc.),
- generators or converters of cryptographic material (keys and algorithm parameters), or
- objects (keystores or certificates) that encapsulate the cryptographic data and can be used at higher layers of abstraction.

The following engine classes are available:

- SecureRandom: used to generate random or pseudo-random numbers.
- MessageDigest: used to calculate the message digest (hash) of specified data.
- Signature: initialized with keys, these are used to sign data and verify digital signatures.



- Cipher: initialized with keys, these used for encrypting/decrypting data. There are various types of algorithms: symmetric bulk encryption (e.g. AES), asymmetric encryption (e.g. RSA), and password-based encryption (e.g. PBE).
- Message Authentication Codes (MAC): like MessageDigests, these also generate
 hash values, but are first initialized with keys to protect the integrity of messages.
- KeyFactory: used to convert existing opaque cryptographic keys of type Key into key specifications (transparent representations of the underlying key material), and vice versa.
- SecretKeyFactory: used to convert existing opaque cryptographic keys of type SecretKey into key specifications (transparent representations of the underlying key material), and vice versa. SecretKeyFactorys are specialized KeyFactorys that create secret (symmetric) keys only.
- KeyPairGenerator: used to generate a new pair of public and private keys suitable for use with a specified algorithm.
- KeyGenerator: used to generate new secret keys for use with a specified algorithm.
- KeyAgreement: used by two or more parties to agree upon and establish a specific key to use for a particular cryptographic operation.
- AlgorithmParameters: used to store the parameters for a particular algorithm, including parameter encoding and decoding.
- AlgorithmParameterGenerator: used to generate a set of AlgorithmParameters suitable for a specified algorithm.
- KeyStore: used to create and manage a keystore. A keystore is a database of keys.
 Private keys in a keystore have a certificate chain associated with them, which authenticates the corresponding public key. A keystore also contains certificates from trusted entities.
- CertificateFactory: used to create public key certificates and Certificate Revocation Lists (CRLs).
- CertPathBuilder: used to build certificate chains (also known as certification paths).
- CertPathValidator: used to validate certificate chains.
- CertStore: used to retrieve Certificates and CRLs from a repository.



A *generator* creates objects with brand-new contents, whereas a *factory* creates objects from existing material (for example, an encoding).

Core Classes and Interfaces

The following are the core classes and interfaces provided in the JCA.

- Provider and Security
- SecureRandom, MessageDigest, Signature, Cipher, Mac, KeyFactory, SecretKeyFactory, KeyPairGenerator, KeyGenerator, KeyAgreement, AlgorithmParameter, AlgorithmParameterGenerator, KeyStore, CertificateFactory, and engine
- Key Interface, KeyPair



- AlgorithmParameterSpec Interface, AlgorithmParameters,
 AlgorithmParameterGenerator, and algorithm parameter specification interfaces and
 classes in the java.security.spec and javax.crypto.spec packages.
- KeySpec Interface, EncodedKeySpec, PKCS8EncodedKeySpec, and X509EncodedKeySpec.
- SecretKeyFactory, KeyFactory, KeyPairGenerator, KeyGenerator, KeyAgreement, and KeyStore.



See CertPathBuilder, CertPathValidator, and CertStoreengine classes in the Java PKI Programmers Guide.

The guide will cover the most useful high-level classes first (Provider, Security, SecureRandom, MessageDigest, Signature, Cipher, and Mac), then delve into the various support classes. For now, it is sufficient to simply say that Keys (public, private, and secret) are generated and represented by the various JCA classes, and are used by the high-level classes as part of their operation.

This section shows the signatures of the main methods in each class and interface. Examples for some of these classes (MessageDigest, Signature, KeyPairGenerator, SecureRandom, KeyFactory, and key specification classes) are supplied in the corresponding Code Examples sections.

The complete reference documentation for the relevant Security API packages can be found in the package summaries:

- java.security
- javax.crypto
- java.security.cert
- java.security.spec
- javax.crypto.spec
- java.security.interfaces
- javax.crypto.interfaces

The Provider Class

The term "Cryptographic Service Provider" (used interchangeably with "provider" in this document) refers to a package or set of packages that supply a concrete implementation of a subset of the JDK Security API cryptography features. The Provider class is the interface to such a package or set of packages. It has methods for accessing the provider name, version number, and other information. Please note that in addition to registering implementations of cryptographic services, the Provider class can also be used to register implementations of other security services that might get defined as part of the JDK Security API or one of its extensions.

To supply implementations of cryptographic services, an entity (e.g., a development group) writes the implementation code and creates a subclass of the Provider class. The constructor of the Provider subclass sets the values of various properties; the JDK Security API uses these values to look up the services that the provider implements. In



other words, the subclass specifies the names of the classes implementing the services.

Figure 2-4 Provider Class

```
ProviderC

provider.java

public class fooJCA extends Provider {
    .
    .
    put("MessageDigest.SHA-256"."com.foo.SHA256");
    .
}

com.foo.SHA256.java

package com.foo;
public class SHA256 extends MessageDigestSpi {
    .
    .
    .
}
```

There are several types of services that can be implemented by provider packages; See Engine Classes and Algorithms.

The different implementations may have different characteristics. Some may be software-based, while others may be hardware-based. Some may be platform-independent, while others may be platform-specific. Some provider source code may be available for review and evaluation, while some may not. The JCA lets both endusers and developers decide what their needs are.

You can find information about how end-users install the cryptography implementations that fit their needs, and how developers request the implementations that fit theirs.



To implement a provider, see Steps to Implement and Integrate a Provider.

How Provider Implementations Are Requested and Supplied

For each engine class (see Engine Classes and Algorithms) in the API, a implementation instance is requested and instantiated by calling one of the <code>getInstance</code> methods on the engine class, specifying the name of the desired algorithm and, optionally, the name of the provider (or the <code>Provider</code> class) whose implementation is desired.

```
static EngineClassName getInstance(String algorithm)
    throws NoSuchAlgorithmException

static EngineClassName getInstance(String algorithm, String provider)
    throws NoSuchAlgorithmException, NoSuchProviderException
```

static EngineClassName getInstance(String algorithm, Provider provider)
throws NoSuchAlgorithmException

where

EngineClassName

is the desired engine type (MessageDigest/Cipher/etc). For example:

```
MessageDigest md = MessageDigest.getInstance("SHA-256");
KeyAgreement ka = KeyAgreement.getInstance("DH", "SunJCE");
```

return an instance of the "SHA-256" MessageDigest and "DH" KeyAgreement objects, respectively.

Java Security Standard Algorithm Names contains the list of names that have been standardized for use with the Java environment. Some providers may choose to also include alias names that also refer to the same algorithm. For example, the "SHA256" algorithm might be referred to as "SHA-256". Applications should use standard names instead of an alias, as not all providers may alias algorithm names in the same way.

Note:

The algorithm name is not case-sensitive. For example, all the following calls are equivalent:

```
MessageDigest.getInstance("SHA256")
MessageDigest.getInstance("sha256")
MessageDigest.getInstance("sHa256")
```

If no provider is specified, <code>getInstance</code> searches the registered providers for an implementation of the requested cryptographic service associated with the named algorithm. In any given Java Virtual Machine (JVM), providers are installed in a given preference order, the order in which the provider list is searched if a specific provider is not requested. (See Installing Providers.) For example, suppose there are two providers installed in a JVM, <code>PROVIDER_1</code> and <code>PROVIDER_2</code>. Assume that:

- PROVIDER_1 implements SHA-256 and DESede. PROVIDER_1 has preference order 1
 (the highest priority).
- PROVIDER_2 implements SHA256withDSA, SHA-256, RC5, and RSA. PROVIDER_2
 has preference order 2.

Now let's look at three scenarios:

- 1. If we are looking for an SHA-256 implementation. Both providers supply such an implementation. The PROVIDER_1 implementation is returned since PROVIDER_1 has the highest priority and is searched first.
- 2. If we are looking for an SHA256withDSA signature algorithm, PROVIDER_1 is first searched for it. No implementation is found, so PROVIDER_2 is searched. Since an implementation is found, it is returned.
- 3. Suppose we are looking for a SHA256withRSA signature algorithm. Since no installed provider implements it, a NoSuchAlgorithmException is thrown.



The <code>getInstance</code> methods that include a provider argument are for developers who want to specify which provider they want an algorithm from. A federal agency, for example, will want to use a provider implementation that has received federal certification. Let's assume that the SHA256withDSA implementation from <code>PROVIDER_1</code> has not received such certification, while the DSA implementation of <code>PROVIDER_2</code> has received it.

A federal agency program would then have the following call, specifying PROVIDER_2 since it has the certified implementation:

```
Signature dsa = Signature.getInstance("SHA256withDSA", "PROVIDER_2");
```

In this case, if PROVIDER_2 was not installed, a NoSuchProviderException would be thrown, even if another installed provider implements the algorithm requested.

A program also has the option of getting a list of all the installed providers (using the getProviders method in The Security Class class) and choosing one from the list.



General purpose applications **SHOULD NOT** request cryptographic services from specific providers. Otherwise, applications are tied to specific providers which may not be available on other Java implementations. They also might not be able to take advantage of available optimized providers (for example hardware accelerators via PKCS11 or native OS implementations such as Microsoft's MSCAPI) that have a higher preference order than the specific requested provider.

Installing Providers

In order to be used, a cryptographic provider must first be installed, then registered either statically or dynamically. There are a variety of Sun providers shipped with this release (SUN, SunJCE, SunJSSE, SunRsaSign, etc.) that are already installed and registered. The following sections describe how to install and register additional providers.

All JDK providers are already installed and registered. However, if you require any third-party providers, see Step 8: Prepare for Testing from Steps to Implement and Integrate a Provider for information about how to add providers to the class or module path, register providers (statically or dynamically), and add any required permissions.

Provider Class Methods

Each Provider class instance has a (currently case-sensitive) name, a version number, and a string description of the provider and its services.

You can query the Provider instance for this information by calling the following methods:

```
public String getName()
public double getVersion()
public String getInfo()
```



The Security Class

The Security class manages installed providers and security-wide properties. It only contains static methods and is never instantiated. The methods for adding or removing providers, and for setting Security properties, can only be executed by a trusted program. Currently, a "trusted program" is either

- A local application not running under a security manager, or
- An applet or application with permission to execute the specified method (see below).

The determination that code is considered trusted to perform an attempted action (such as adding a provider) requires that the applet is granted the proper permission(s) for that particular action. The policy configuration file(s) for a JDK installation specify what permissions (which types of system resource accesses) are allowed by code from specified code sources. (See below and the Default Policy Implementation and Policy File Syntax and Java Security Architecture Specification files.)

Code being executed is always considered to come from a particular "code source". The code source includes not only the location (URL) where the code originated from, but also a reference to any public key(s) corresponding to the private key(s) that may have been used to sign the code. Public keys in a code source are referenced by (symbolic) alias names from the user's .

In a policy configuration file, a code source is represented by two components: a code base (URL), and an alias name (preceded by signedBy), where the alias name identifies the keystore entry containing the public key that must be used to verify the code's signature.

Each "grant" statement in such a file grants a specified code source a set of permissions, specifying which actions are allowed.

Here is a sample policy configuration file:

```
grant codeBase "file:/home/sysadmin/", signedBy "sysadmin" {
   permission java.security.SecurityPermission "insertProvider";
   permission java.security.SecurityPermission "removeProvider";
   permission java.security.SecurityPermission "putProviderProperty.*";
};
```

This configuration file specifies that code loaded from a signed JAR file in the <code>/home/sysadmin/</code> directory on the local file system can add or remove providers or set provider properties. (Note that the signature of the JAR file can be verified using the public key referenced by the alias name <code>sysadmin</code> in the user's keystore.).

Either component of the code source (or both) may be missing. Here's an example of a configuration file where the codeBase is omitted:

```
grant signedBy "sysadmin" {
    permission java.security.SecurityPermission "insertProvider.*";
    permission java.security.SecurityPermission "removeProvider.*";
};
```



If this policy is in effect, code that comes in a JAR File signed by $\protect\ensuremath{\text{home/sysadmin/}}\protect\ensuremath{\text{directory}}\protect\ensuremath{\text{on}}$ the local filesystem can add or remove providers. The code does not need to be signed.

An example where neither codeBase nor signedBy is included is:

```
grant {
   permission java.security.SecurityPermission "insertProvider.*";
   permission java.security.SecurityPermission "removeProvider.*";
};
```

Here, with both code source components missing, any code (regardless of where it originates, or whether or not it is signed, or who signed it) can add/remove providers. Obviously, this is definitely not recommended, as this grant could open a security hole. Untrusted code could install a Provider, thus affecting later code that is depending on a properly functioning implementation. (For example, a rogue cipher object might capture and store the sensitive information it receives.)

Managing Providers

The following tables summarize the methods in the <code>Security</code> class you can use to query which <code>Providers</code> are installed, as well as to install or remove providers at runtime.

Querying Providers

Method	Description
<pre>static Provider[] getProviders()</pre>	Returns an array containing all the installed providers (technically, the Provider subclass for each package provider). The order of the Providers in the array is their preference order.
<pre>static Provider getProvider (String providerName)</pre>	Returns the Provider named providerName. It returns null if the Provider is not found.

Adding Providers

Method	Description
static int addProvider(Provider provider)	Adds a Provider to the end of the list of installed Providers. It returns the preference position in which the Provider was added, or -1 if the Provider was not added because it was already installed.
static int insertProviderAt (Provider provider, int position)	Adds a new Provider at a specified position. If the given provider is installed at the requested position, the provider formerly at that position and all providers with a position greater than position are shifted up one position (towards the end of the list). This method returns the preference position in which the Provider was added, or -1 if the Provider was not added because it was already installed.



Removing Providers

Method	Description
static void removeProvider(String name)	Removes the Provider with the specified name. It returns silently if the provider is not installed. When the specified provider is removed, all providers located at a position greater than where the specified provider was are shifted down one position (towards the head of the list of installed providers).



If you want to change the preference position of a provider, you must first remove it, and then insert it back in at the new preference position.

Security Properties

The security class maintains a list of system-wide Security Properties. These properties are similar to the system properties, but are security-related. These properties can be set statically (through the <java-home>/conf/security/java.security file) or dynamically (using an API). See Step 8.1: Configure the Provider from Steps to Implement and Integrate a Provider. for an example of registering a provider statically with the security.provider.n Security Property. If you want to set properties dynamically, trusted programs can use the following methods:

static String getProperty(String key)
static void setProperty(String key, String datum)



The list of security providers is established during VM startup; therefore, the methods described above must be used to alter the provider list.

The SecureRandom Class

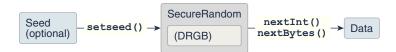
The SecureRandom class is an engine class (see Engine Classes and Algorithms) that provides cryptographically strong random numbers, either by accessing a pseudorandom number generator (PRNG), a deterministic algorithm that produces a pseudorandom sequence from an initial seed value, or by reading a native source of randomness (for example, /dev/random or a true random number generator). One example of a PRNG is the Deterministic Random Bits Generator (DRBG) as specified in NIST SP 800-90Ar1. Other implementations may produce true random numbers, and yet others may use a combination of both techniques. A cryptographically strong random number minimally complies with the statistical random number generator tests specified in FIPS 140-2, Security Requirements for Cryptographic Modules, section 4.9.1.

All Java SE implementations must indicate the strongest (most random) implementation of SecureRandom that they provide in the

securerandom.strongAlgorithms property of the java.security.Security class. This implementation can be used when a particularly strong random value is required.

The securerandom.drbg.config property is used to specify the DRBG SecureRandom configuration and implementations in the SUN provider. The securerandom.drbg.config is a property of the java.security.Security class. Other DRBG implementations can also use the securerandom.drbg.config property.

Figure 2-5 SecureRandom class



Creating a SecureRandom Object

There are several ways to obtain an instance of SecureRandom:

- All Java SE implementations provide a default SecureRandom using the no-argument constructor: new SecureRandom(). This constructor traverses the list of registered security providers, starting with the most preferred provider, then returns a new SecureRandom object from the first provider that supports a SecureRandom random number generator (RNG) algorithm. If none of the providers support a RNG algorithm, then it returns a SecureRandom object that uses SHA1PRNG from the SUN provider.
- To get a specific implementation of SecureRandom, use one of the How Provider Implementations Are Requested and Supplied.
- Use the <code>getInstanceStrong()</code> method to obtain a strong <code>SecureRandom</code> implementation as defined by the <code>securerandom.strongAlgorithms</code> property of the <code>java.security.Security</code> class. This property lists platform implementations that are suitable for generating important values.

Seeding or Re-Seeding the SecureRandom Object

The SecureRandom object is initialized with a random seed unless the call to getInstance() is followed by a call to one of the following setSeed methods.

```
void setSeed(byte[] seed)
void setSeed(long seed)
```

You must call setSeed before the first nextBytes call to prevent any environmental randomness.

The randomness of the bits produced by the SecureRandom object depends on the randomness of the seed bits

At any time a SecureRandom object may be re-seeded using one of the setSeed or reseed methods. The given seed for setSeed supplements, rather than replaces, the existing seed; therefore, repeated calls are guaranteed never to reduce randomness.



Using a SecureRandom Object

To get random bytes, a caller simply passes an array of any length, which is then filled with random bytes:

void nextBytes(byte[] bytes)

Generating Seed Bytes

If desired, it is possible to invoke the generateSeed method to generate a given number of seed bytes (to seed other random number generators, for example):

byte[] generateSeed(int numBytes)

The MessageDigest Class

The MessageDigest class is an engine class (see Engine Classes and Algorithms) designed to provide the functionality of cryptographically secure message digests such as SHA-256 or SHA-512. A cryptographically secure message digest takes arbitrary-sized input (a byte array), and generates a fixed-size output, called a *digest* or hash.

Figure 2-6 MessageDigest Class



For example, the SHA-256 algorithm produces a 32-byte digest, and SHA-512's is 64 bytes.

A digest has two properties:

- It should be computationally infeasible to find two messages that hash to the same value.
- The digest should not reveal anything about the input that was used to generate it.

Message digests are used to produce unique and reliable identifiers of data. They are sometimes called "checksums" or the "digital fingerprints" of the data. Changes to just one bit of the message should produce a different digest value.

Message digests have many uses and can determine when data has been modified, intentionally or not. Recently, there has been considerable effort to determine if there are any weaknesses in popular algorithms, with mixed results. When selecting a digest algorithm, one should always consult a recent reference to determine its status and appropriateness for the task at hand.

Creating a MessageDigest Object

Procedure to create a MessageDigest object.



• To compute a digest, create a message digest instance. The MessageDigest objects are obtained by using one of the getInstance() methods in the MessageDigest class. See How Provider Implementations Are Requested and Supplied.

The factory method returns an initialized message digest object. It thus does not need further initialization.

Updating a Message Digest Object

Procedure to update the Message Digest object.

 To calculate the digest of some data, you have to supply the data to the initialized message digest object. It can be provided all at once, or in chunks. Pieces can be fed to the message digest by calling one of the update methods:

```
void update(byte input)
void update(byte[] input)
void update(byte[] input, int offset, int len)
```

Computing the Digest

Procedure to compute the digest using different types of digest() methods.

The data chunks have to be supplied by calls to update method. See Updating a Message Digest Object.

The digest is computed using a call to one of the digest methods:

```
byte[] digest()
byte[] digest(byte[] input)
int digest(byte[] buf, int offset, int len)
```

- 1. The byte[] digest() method return the computed digest.
- 2. The byte[] digest(byte[] input) method does a final update(input) with the input byte array before calling digest(), which returns the digest byte array.
- 3. The int digest(byte[] buf, int offset, int len) method stores the computed digest in the provided buffer buf, starting at offset. len is the number of bytes in buf allotted for the digest, the method returns the number of bytes actually stored in buf. If there is not enough room in the buffer, the method will throw an exception.

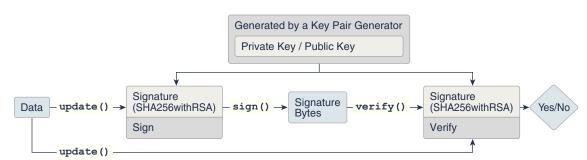
See Computing a MessageDigest Object.

The Signature Class

The Signature class is an engine class (see Engine Classes and Algorithms (designed to provide the functionality of a cryptographic digital signature algorithm such as SHA256withDSA or SHA512withRSA. A cryptographically secure signature algorithm takes arbitrary-sized input and a private key and generates a relatively short (often fixed-size) string of bytes, called the *signature*, with the following properties:

 Only the owner of a private/public key pair is able to create a signature. It should be computationally infeasible for anyone having only the public key and a number of signatures to recover the private key. • Given the public key corresponding to the private key used to generate the signature, it should be possible to verify the authenticity and integrity of the input.

Figure 2-7 Signature Class



A Signature object is initialized for signing with a Private Key and is given the data to be signed. The resulting signature bytes are typically kept with the signed data. When verification is needed, another Signature object is created and initialized for verification and given the corresponding Public Key. The data and the signature bytes are fed to the signature object, and if the data and signature match, the Signature object reports success.

Even though a signature seems similar to a message digest, they have very different purposes in the type of protection they provide. In fact, algorithms such as "SHA256WithRSA" use the message digest "SHA256" to initially "compress" the large data sets into a more manageable form, then sign the resulting 32 byte message digest with the "RSA" algorithm.

For an example for signing and verifying data, see Generating and Verifying a Signature Using Generated Keys.

Signature Object States

signature objects are modal objects. This means that a signature object is always in a given state, where it may only do one type of operation.

States are represented as final integer constants defined in their respective classes.

The three states a Signature object may have are:

- UNINITIALIZED
- SIGN
- VERIFY

When it is first created, a signature object is in the UNINITIALIZED state. The signature class defines two initialization methods, initsign and initVerify, which change the state to SIGN and VERIFY, respectively.

Creating a Signature Object

The first step for signing or verifying a signature is to create a Signature instance.

Signature objects are obtained by using one of the Signature getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.



Initializing a Signature Object

A Signature object must be initialized before it is used. The initialization method depends on whether the object is going to be used for signing or for verification.

If it is going to be used for signing, the object must first be initialized with the private key of the entity whose signature is going to be generated. This initialization is done by calling the method:

```
final void initSign(PrivateKey privateKey)
```

This method puts the <code>signature</code> object in the <code>signature</code> object is going to be used for verification, it must first be initialized with the public key of the entity whose signature is going to be verified. This initialization is done by calling either of these methods:

```
final void initVerify(PublicKey publicKey)
final void initVerify(Certificate certificate)
```

This method puts the Signature object in the VERIFY state.

Signing with a Signature Object

If the Signature object has been initialized for signing (if it is in the SIGN state), the data to be signed can then be supplied to the object. This is done by making one or more calls to one of the update methods:

```
final void update(byte b)
final void update(byte[] data)
final void update(byte[] data, int off, int len)
```

Calls to the update method(s) should be made until all the data to be signed has been supplied to the Signature Object.

To generate the signature, simply call one of the sign methods:

```
final byte[] sign()
final int sign(byte[] outbuf, int offset, int len)
```

The first method returns the signature result in a byte array. The second stores the signature result in the provided buffer *outbuf*, starting at *offset*. *len* is the number of bytes in *outbuf* allotted for the signature. The method returns the number of bytes actually stored.

Signature encoding is algorithm specific. See Java Security Standard Algorithm Names Specification to know more about the use of ASN.1 encoding in the Java Cryptography Architecture.

A call to a sign method resets the signature object to the state it was in when previously initialized for signing via a call to initsign. That is, the object is reset and available to generate another signature with the same private key, if desired, via new calls to update and sign.



Alternatively, a new call can be made to initSign specifying a different private key, or to initVerify (to initialize the Signature object to verify a signature).

Verifying with a Signature Object

If the Signature object has been initialized for verification (if it is in the VERIFY state), it can then verify if an alleged signature is in fact the authentic signature of the data associated with it. To start the process, the data to be verified (as opposed to the signature itself) is supplied to the object. The data is passed to the object by calling one of the update methods:

```
final void update(byte b)
final void update(byte[] data)
final void update(byte[] data, int off, int len)
```

Calls to the update method(s) should be made until all the data to be verified has been supplied to the Signature object. The signature can now be verified by calling one of the verify methods:

```
final boolean verify(byte[] signature)
final boolean verify(byte[] signature, int offset, int length)
```

The argument must be a byte array containing the signature. This byte array would hold the signature bytes which were returned by a previous call to one of the sign methods.

The verify method returns a boolean indicating whether or not the encoded signature is the authentic signature of the data supplied to the update method(s).

A call to the verify method resets the signature object to its state when it was initialized for verification via a call to initVerify. That is, the object is reset and available to verify another signature from the identity whose public key was specified in the call to initVerify.

Alternatively, a new call can be made to <code>initVerify</code> specifying a different public key (to initialize the <code>signature</code> object for verifying a signature from a different entity), or to <code>initSign</code> (to initialize the <code>signature</code> object for generating a signature).

The Cipher Class

The cipher class provides the functionality of a cryptographic cipher used for encryption and decryption. Encryption is the process of taking data (called *cleartext*) and a *key*, and producing data (*ciphertext*) meaningless to a third-party who does not know the key. Decryption is the inverse process: that of taking ciphertext and a key and producing cleartext.



Plaintext

Decrypt

Cipher Algorithm Parameters Cipher update() (AES)

Figure 2-8 The Cipher Class

Plaintext

Symmetric vs. Asymmetric Cryptography

Encrypt

There are two major types of encryption: *symmetric* (also known as *secret key*), and asymmetric (or *public key cryptography*). In symmetric cryptography, the same secret key to both encrypt and decrypt the data. Keeping the key private is critical to keeping the data confidential. On the other hand, asymmetric cryptography uses a public/ private key pair to encrypt data. Data encrypted with one key is decrypted with the other. A user first generates a public/private key pair, and then publishes the public key in a trusted database that anyone can access. A user who wishes to communicate securely with that user encrypts the data using the retrieved public key. Only the holder of the private key will be able to decrypt. Keeping the private key confidential is critical to this scheme.

Asymmetric algorithms (such as RSA) are generally much slower than symmetric ones. These algorithms are not designed for efficiently protecting large amounts of data. In practice, asymmetric algorithms are used to exchange smaller secret keys which are used to initialize symmetric algorithms.

Stream vs. Block Ciphers

There are two major types of ciphers: *block* and *stream*. Block ciphers process entire blocks at a time, usually many bytes in length. If there is not enough data to make a complete input block, the data must be *padded*: that is, before encryption, dummy bytes must be added to make a multiple of the cipher's block size. These bytes are then stripped off during the decryption phase. The padding can either be done by the application, or by initializing a cipher to use a padding type such as "PKCS5PADDING". In contrast, stream ciphers process incoming data one small unit (typically a byte or even a bit) at a time. This allows for ciphers to process an arbitrary amount of data without padding.

Modes Of Operation

When encrypting using a simple block cipher, two identical blocks of plaintext will always produce an identical block of cipher text. Cryptanalysts trying to break the ciphertext will have an easier job if they note blocks of repeating text. A cipher mode of operation makes the ciphertext less predictable with output block alterations based on block position or the values of other ciphertext blocks. The first block will need an initial value, and this value is called the *initialization vector (IV)*. Since the IV simply alters the data before any encryption, the IV should be random but does not necessarily need to be kept secret. There are a variety of modes, such as CBC (Cipher Block Chaining), CFB (Cipher Feedback Mode), and OFB (Output Feedback Mode). ECB (Electronic Codebook Mode) is a mode in which there is no influence from block position or other ciphertext blocks. Because ECB ciphertexts are the same if they use



the same plaintext/key, this mode is not typically suitable for cryptographic applications and should not be used.

Some algorithms such as AES and RSA allow for keys of different lengths, but others are fixed, such as 3DES. Encryption using a longer key generally implies a stronger resistance to message recovery. As usual, there is a trade off between security and time, so choose the key length appropriately.

Most algorithms use binary keys. Most humans do not have the ability to remember long sequences of binary numbers, even when represented in hexadecimal. Character passwords are much easier to recall. Because character passwords are generally chosen from a small number of characters (for example, [a-zA-Z0-9]), protocols such as "Password-Based Encryption" (PBE) have been defined which take character passwords and generate strong binary keys. In order to make the task of getting from password to key very time-consuming for an attacker (via so-called "rainbow table attacks" or "precomputed dictionary attacks" where common dictionary word->value mappings are precomputed), most PBE implementations will mix in a random number, known as a *salt*, to reduce the usefulness of precomputed tables.

Newer cipher modes such as Authenticated Encryption with Associated Data (AEAD) (for example, Galois/Counter Mode (GCM)) encrypt data and authenticate the resulting message simultaneously. Additional Associated Data (AAD) can be used during the calculation of the resulting AEAD tag (MAC), but this AAD data is not output as ciphertext. (For example, some data might not need to be kept confidential, but should figure into the tag calculation to detect modifications.) The Cipher.updateAAD() methods can be used to include AAD in the tag calculations.

Using an AES Cipher with GCM Mode

AES Cipher with GCM is an AEAD Cipher which has different usage patterns than the non-AEAD ciphers. Apart from the regular data, it also takes AAD which is optional for encryption/decryption but AAD must be supplied before data for encryption/decryption. In addition, in order to use GCM securely, callers should not re-use key and IV combinations for encryption. This means that the cipher object should be explicitly re-initialized with a different set of parameters every time for each encryption operation.

Example 2-2 Sample Code for Using an AES Cipher with GCM Mode

```
SecretKey myKey = ...
   byte[] myAAD = ...
    byte[] plainText = ...
        int myTLen = ...
       byte[] myIv = ...
    GCMParameterSpec myParams = new GCMParameterSpec(myTLen, myIv);
    Cipher c = Cipher.getInstance("AES/GCM/NoPadding");
    c.init(Cipher.ENCRYPT_MODE, myKey, myParams);
    // AAD is optional, if present, it must be supplied before any update/doFinal
calls.
    c.updateAAD(myAAD); // if AAD is non-null
    byte[] cipherText = new byte[c.getOutputSize(plainText.length)];
    // conclusion of encryption operation
    int actualOutputLen = c.doFinal(plainText, 0, plainText.length, cipherText);
    // To decrypt, same AAD and GCM parameters must be supplied
    c.init(Cipher.DECRYPT_MODE, myKey, myParams);
    c.updateAAD(myAAD);
```



```
byte[] recoveredText = c.doFinal(cipherText, 0, actualOutputLen);

// MUST CHANGE IV VALUE if the same key were to be used again for encryption
    byte[] newIv = ...;

myParams = new GCMParameterSpec(myTLen, newIv);
```

Creating a Cipher Object

cipher objects are obtained by using one of the cipher getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied. Here, the algorithm name is slightly different than with other engine classes, in that it specifies not just an algorithm name, but a "transformation". A transformation is a string that describes the operation (or set of operations) to be performed on the given input to produce some output. A transformation always includes the name of a cryptographic algorithm (e.g., AES), and may be followed by a mode and padding scheme.

A transformation is of the form:

- "algorithm/mode/padding" or
- "algorithm"

For example, the following are valid transformations:

```
"AES"
```

If just a transformation name is specified, the system will determine if there is an implementation of the requested transformation available in the environment, and if there is more than one, returns there is a preferred one.

If both a transformation name and a package provider are specified, the system will determine if there is an implementation of the requested transformation in the package requested, and throw an exception if there is not.

It is recommended to use a transformation that fully specifies the algorithm, mode, and padding. By not doing so, the provider will use a default. For example, the SunJCE and SunPKCS11 providers use ECB as the default mode, and PKCS5Padding as the default padding for many symmetric ciphers.

This means that in the case of the <code>SunJCE</code> provider:

```
Cipher c1 = Cipher.getInstance("AES/ECB/PKCS5Padding");
and
Cipher c1 = Cipher.getInstance("AES");
```

Note:

are equivalent statements.

ECB mode is the easiest block cipher mode to use and is the default in the JDK and JRE. ECB works for single blocks of data when using different keys, but it absolutely should not be used for multiple data blocks. Other cipher modes such as Cipher Block Chaining (CBC) or Galois/Counter Mode (GCM) are more appropriate.



Using modes such as CFB and OFB, block ciphers can encrypt data in units smaller than the cipher's actual block size. When requesting such a mode, you may optionally specify the number of bits to be processed at a time by appending this number to the mode name as shown in the "AES/CFB8/NoPadding" and "AES/OFB32/PKCS5Padding" transformations. If no such number is specified, a provider-specific default is used. (For example, the SunJCE provider uses a default of 128 bits for AES.) Thus, block ciphers can be turned into byte-oriented stream ciphers by using an 8 bit mode such as CFB8 or OFB8.

Java Security Standard Algorithm Names contains a list of standard names that can be used to specify the algorithm name, mode, and padding scheme components of a transformation.

The objects returned by factory methods are uninitialized, and must be initialized before they become usable.

Initializing a Cipher Object

A Cipher object obtained via <code>getInstance</code> must be initialized for one of four modes, which are defined as final integer constants in the <code>cipher</code> class. The modes can be referenced by their symbolic names, which are shown below along with a description of the purpose of each mode:

ENCRYPT_MODE

Encryption of data.

DECRYPT MODE

Decryption of data.

WRAP_MODE

Wrapping a java.security.Key into bytes so that the key can be securely transported.

UNWRAP_MODE

Unwrapping of a previously wrapped key into a java.security.Key object.

Each of the Cipher initialization methods takes an operational mode parameter (opmode), and initializes the Cipher object for that mode. Other parameters include the key (key) or certificate containing the key (certificate), algorithm parameters (params), and a source of randomness (random).

To initialize a Cipher object, call one of the following init methods:



```
AlgorithmParameters params);

public void init(int opmode, Key key,

AlgorithmParameters params, SecureRandom random);
```

If a Cipher object that requires parameters (e.g., an initialization vector) is initialized for encryption, and no parameters are supplied to the <code>init</code> method, the underlying cipher implementation is supposed to supply the required parameters itself, either by generating random parameters or by using a default, provider-specific set of parameters.

However, if a Cipher object that requires parameters is initialized for decryption, and no parameters are supplied to the init method, an InvalidKeyException or InvalidAlgorithmParameterException exception will be raised, depending on the init method that has been used.

See Managing Algorithm Parameters.

The same parameters that were used for encryption must be used for decryption.

Note that when a Cipher object is initialized, it loses all previously-acquired state. In other words, initializing a Cipher is equivalent to creating a new instance of that Cipher, and initializing it. For example, if a Cipher is first initialized for decryption with a given key, and then initialized for encryption, it will lose any state acquired while in decryption mode.

Encrypting and Decrypting Data

Data can be encrypted or decrypted in one step (*single-part operation*) or in multiple steps (*multiple-part operation*). A multiple-part operation is useful if you do not know in advance how long the data is going to be, or if the data is too long to be stored in memory all at once.

To encrypt or decrypt data in a single step, call one of the doFinal methods:

To encrypt or decrypt data in multiple steps, call one of the update methods:



A multiple-part operation must be terminated by one of the above doFinal methods (if there is still some input data left for the last step), or by one of the following doFinal methods (if there is no input data left for the last step):

```
public byte[] doFinal();
public int doFinal(byte[] output, int outputOffset);
```

All the doFinal methods take care of any necessary padding (or unpadding), if padding (or unpadding) has been requested as part of the specified transformation.

A call to doFinal resets the Cipher object to the state it was in when initialized via a call to init. That is, the Cipher object is reset and available to encrypt or decrypt (depending on the operation mode that was specified in the call to init) more data.

Wrapping and Unwrapping Keys

Wrapping a key enables secure transfer of the key from one place to another.

The wrap/unwrap API makes it more convenient to write code since it works with key objects directly. These methods also enable the possibility of secure transfer of hardware-based keys.

To **wrap** a Key, first initialize the Cipher object for WRAP_MODE, and then call the following:

```
public final byte[] wrap(Key key);
```

If you are supplying the wrapped key bytes (the result of calling wrap) to someone else who will unwrap them, be sure to also send additional information the recipient will need in order to do the unwrap:

- The name of the key algorithm.
- The type of the wrapped key (one of Cipher.SECRET_KEY, Cipher.PRIVATE_KEY, or Cipher.PUBLIC_KEY).

The key algorithm name can be determined by calling the <code>getAlgorithm</code> method from the Key interface:

```
public String getAlgorithm();
```

To **unwrap** the bytes returned by a previous call to wrap, first initialize a Cipher object for UNWRAP_MODE, then call the following:

Here, wrappedKey is the bytes returned from the previous call to wrap, wrappedKeyAlgorithm is the algorithm associated with the wrapped key, and wrappedKeyType is the type of the wrapped key. This must be one of Cipher.SECRET_KEY, Cipher.PRIVATE_KEY, OF Cipher.PUBLIC_KEY.



Managing Algorithm Parameters

The parameters being used by the underlying Cipher implementation, which were either explicitly passed to the init method by the application or generated by the underlying implementation itself, can be retrieved from the Cipher object by calling its getParameters method, which returns the parameters as a java.security.AlgorithmParameters object (or null if no parameters are being used). If the parameter is an initialization vector (IV), it can also be retrieved by calling the getIV method.

In the following example, a Cipher object implementing password-based encryption (PBE) is initialized with just a key and no parameters. However, the selected algorithm for password-based encryption requires two parameters - a *salt* and an *iteration count*. Those will be generated by the underlying algorithm implementation itself. The application can retrieve the generated parameters from the Cipher object, see Example 2-3.

The same parameters that were used for encryption must be used for decryption. They can be instantiated from their encoding and used to initialize the corresponding Cipher object for decryption, see Example 2-4.

If you did not specify any parameters when you initialized a Cipher object, and you are not sure whether or not the underlying implementation uses any parameters, you can find out by simply calling the <code>getParameters</code> method of your Cipher object and checking the value returned. A return value of <code>null</code> indicates that no parameters were used.

The following cipher algorithms implemented by the SunJCE provider use parameters:

- AES, DES-EDE, and Blowfish, when used in feedback (i.e., CBC, CFB, OFB, or PCBC) mode, use an initialization vector (IV). The <code>javax.crypto.spec.IvParameterSpec</code> class can be used to initialize a Cipher object with a given IV. In addition, CTR and GCM modes require an IV.
- PBE Cipher algorithms use a set of parameters, comprising a salt and an iteration count. The <code>javax.crypto.spec.PBEParameterSpec</code> class can be used to initialize a Cipher object implementing a PBE algorithm (for example: PBEWithHmacSHA256AndAES 256) with a given salt and iteration count.

Note that you do not have to worry about storing or transferring any algorithm parameters for use by the decryption operation if you use the The SealedObject Class class. This class attaches the parameters used for sealing (encryption) to the encrypted object contents, and uses the same parameters for unsealing (decryption).

Example 2-3 Sample Code for Retrieving Parameters from the Cipher Object

The application can retrieve the generated parameters for encryption from the Cipher object as follows:

```
import javax.crypto.*;
import java.security.AlgorithmParameters;

// get cipher object for password-based encryption
Cipher c = Cipher.getInstance("PBEWithHmacSHA256AndAES_256");

// initialize cipher for encryption, without supplying
// any parameters. Here, "myKey" is assumed to refer
// to an already-generated key.
c.init(Cipher.ENCRYPT MODE, myKey);
```



```
// encrypt some data and store away ciphertext
// for later decryption
byte[] cipherText = c.doFinal("This is just an example".getBytes());

// retrieve parameters generated by underlying cipher
// implementation
AlgorithmParameters algParams = c.getParameters();

// get parameter encoding and store it away
byte[] encodedAlgParams = algParams.getEncoded();
```

Example 2-4 Sample Code for Initializing the Cipher Object for Decryption

The same parameters that were used for encryption must be used for decryption. They can be instantiated from their encoding and used to initialize the corresponding Cipher object for decryption as follows:

```
import javax.crypto.*;
import java.security.AlgorithmParameters;

// get parameter object for password-based encryption
AlgorithmParameters algParams;
algParams = AlgorithmParameters.getInstance("PBEWithHmacSHA256AndAES_256");

// initialize with parameter encoding from above
algParams.init(encodedAlgParams);

// get cipher object for password-based encryption
Cipher c = Cipher.getInstance("PBEWithHmacSHA256AndAES_256");

// initialize cipher for decryption, using one of the
// init() methods that takes an AlgorithmParameters
// object, and pass it the algParams object from above
c.init(Cipher.DECRYPT_MODE, myKey, algParams);
```

Cipher Output Considerations

Some of the update and doFinal methods of Cipher allow the caller to specify the output buffer into which to encrypt or decrypt the data. In these cases, it is important to pass a buffer that is large enough to hold the result of the encryption or decryption operation.

The following method in Cipher can be used to determine how big the output buffer should be:

```
public int getOutputSize(int inputLen)
```

Other Cipher-based Classes

There are some helper classes which internally use ciphers to provide easy access to common cipher uses.

Topics

The Cipher Stream Classes

The SealedObject Class



The Cipher Stream Classes

The CipherInputStream and CipherOutputStream classes are Cipher stream classes.

The CipherInputStream Class

This class is a FilterInputStream that encrypts or decrypts the data passing through it. It is composed of an InputStream. CipherInputStream represents a secure input stream into which a Cipher object has been interposed. The read methods of CipherInputStream return data that are read from the underlying InputStream but have additionally been processed by the embedded Cipher object. The Cipher object must be fully initialized before being used by a CipherInputStream.

For example, if the embedded Cipher has been initialized for decryption, the CipherInputStream will attempt to decrypt the data it reads from the underlying InputStream before returning them to the application.

This class adheres strictly to the semantics, especially the failure semantics, of its ancestor classes <code>java.io.FilterInputStream</code> and <code>java.io.InputStream</code>. This class has exactly those methods specified in its ancestor classes, and overrides them all, so that the data are additionally processed by the embedded cipher. Moreover, this class catches all exceptions that are not thrown by its ancestor classes. In particular, the <code>skip(long)</code> method skips only data that has been processed by the <code>Cipher</code>.

It is crucial for a programmer using this class not to use methods that are not defined or overridden in this class (such as a new method or constructor that is later added to one of the super classes), because the design and implementation of those methods are unlikely to have considered security impact with regard to CipherInputStream. See Example 2-5 for its usage, suppose cipher1 has been initialized for encryption. The program reads and encrypts the content from the file /tmp/a.txt and then stores the result (the encrypted bytes) in /tmp/b.txt.

Example 2-6 demonstrates how to easily connect several instances of CipherInputStream and FileInputStream. In this example, assume that cipher1 and cipher2 have been initialized for encryption and decryption (with corresponding keys), respectively. The program copies the content from file /tmp/a.txt to /tmp/b.txt, except that the content is first encrypted and then decrypted back when it is read from /tmp/a.txt. Of course since this program simply encrypts text and decrypts it back right away, it's actually not very useful except as a simple way of illustrating chaining of CipherInputStreams.

Note that the read methods of the CipherInputStream will block until data is returned from the underlying cipher. If a block cipher is used, a full block of cipher text will have to be obtained from the underlying InputStream.

Example 2-5 Sample Code for Using CipherInputStream and FileInputStream

The code below demonstrates how to use a CipherInputStream containing that cipher and a FileInputStream in order to encrypt input stream data:

```
try (FileInputStream fis = new FileInputStream("/tmp/a.txt");
CipherInputStream cis = new CipherInputStream(fis, cipher1);
FileOutputStream fos = new FileOutputStream("/tmp/b.txt")) {
    byte[] b = new byte[8];
    int i = cis.read(b);
    while (i != -1) {
        fos.write(b, 0, i);
    }
}
```



```
i = cis.read(b);
}
```

Example 2-6 Sample Code for Connecting CipherInputStream and FileInputStream

The following example demonstrates how to easily connect several instances of CipherInputStream and FileInputStream. In this example, assume that cipher1 and cipher2 have been initialized for encryption and decryption (with corresponding keys), respectively:

The CipherOutputStream Class

This class is a FilterOutputStream that encrypts or decrypts the data passing through it. It is composed of an OutputStream, or one of its subclasses, and a Cipher. CipherOutputStream represents a secure output stream into which a Cipher object has been interposed. The write methods of CipherOutputStream first process the data with the embedded Cipher object before writing them out to the underlying OutputStream. The Cipher object must be fully initialized before being used by a CipherOutputStream.

For example, if the embedded Cipher has been initialized for encryption, the CipherOutputStream will encrypt its data, before writing them out to the underlying output stream.

This class adheres strictly to the semantics, especially the failure semantics, of its ancestor classes <code>java.io.OutputStream</code> and <code>java.io.FilterOutputStream</code>. This class has exactly those methods specified in its ancestor classes, and overrides them all, so that all data are additionally processed by the embedded cipher. Moreover, this class catches all exceptions that are not thrown by its ancestor classes.

It is crucial for a programmer using this class not to use methods that are not defined or overridden in this class (such as a new method or constructor that is later added to one of the super classes), because the design and implementation of those methods are unlikely to have considered security impact with regard to CipherOutputStream.

See Example 2-7, for its usage, suppose cipher1 has been initialized for encryption. The program reads the content from the file /tmp/a.txt, then encrypts and stores the result (the encrypted bytes) in /tmp/b.txt.

Example 2-7 demonstrates how to easily connect several instances of CipherOutputStream and FileOutputStream. In this example, assume that cipher1 and cipher2 have been initialized for decryption and encryption (with corresponding keys), respectively. The program copies the content from file /tmp/a.txt to /tmp/b.txt, except that the content is first encrypted and then decrypted back before it is written to /tmp/b.txt.



One thing to keep in mind when using *block* cipher algorithms is that a full block of plaintext data must be given to the CipherOutputStream before the data will be encrypted and sent to the underlying output stream.

There is one other important difference between the flush and close methods of this class, which becomes even more relevant if the encapsulated Cipher object implements a block cipher algorithm with padding turned on:

- flush flushes the underlying OutputStream by forcing any buffered output bytes that have already been processed by the encapsulated Cipher object to be written out. Any bytes buffered by the encapsulated Cipher object and waiting to be processed by it will **not** be written out.
- close closes the underlying OutputStream and releases any system resources associated with it. It invokes the doFinal method of the encapsulated Cipher object, causing any bytes buffered by it to be processed and written out to the underlying stream by calling its flush method.

Example 2-7 Sample Code for Using CipherOutputStream and FileOutputStream

```
CipherOutputStreamFileOutputStream

try (FileInputStream fis = new FileInputStream("/tmp/a.txt");
        FileOutputStream fos = new FileOutputStream("/tmp/b.txt");
        CipherOutputStream cos = new CipherOutputStream(fos, cipher1)) {
        byte[] b = new byte[8];
        int i = fis.read(b);
        while (i != -1) {
            cos.write(b, 0, i);
            i = fis.read(b);
        }
        cos.flush();
}
```

Example 2-8 Sample Code for Connecting CipherOutputStream and FileOutputStream

```
CipherOutputStreamFileOutputStreamcipher1cipher2

try (FileInputStream fis = new FileInputStream("/tmp/a.txt");
        FileOutputStream fos = new FileOutputStream("/tmp/b.txt");
        CipherOutputStream cos1 = new CipherOutputStream(fos, cipher1);
        CipherOutputStream cos2 = new CipherOutputStream(cos1, cipher2)) {
        byte[] b = new byte[8];
        int i = fis.read(b);
        while (i != -1) {
            cos2.write(b, 0, i);
            i = fis.read(b);
        }
        cos2.flush();
}
```

The SealedObject Class

This class enables a programmer to create an object and protect its confidentiality with a cryptographic algorithm.

Given any object that implements the <code>java.io.Serializable</code> interface, one can create a <code>SealedObject</code> that encapsulates the original object, in serialized format (i.e., a "deep

copy"), and seals (encrypts) its serialized contents, using a cryptographic algorithm such as AES, to protect its confidentiality. The encrypted content can later be decrypted (with the corresponding algorithm using the correct decryption key) and deserialized, yielding the original object.

A typical usage is illustrated in the following code segment: In order to seal an object, you create a <code>SealedObject</code> from the object to be sealed and a fully initialized <code>Cipher</code> object that will encrypt the serialized object contents. In this example, the String "This is a secret" is sealed using the AES algorithm. Note that any algorithm parameters that may be used in the sealing operation are stored inside of <code>SealedObject</code>:

```
// create Cipher object
// NOTE: sKey is assumed to refer to an already-generated
// secret AES key.
Cipher c = Cipher.getInstance("AES");
c.init(Cipher.ENCRYPT_MODE, sKey);

// do the sealing
SealedObject so = new SealedObject("This is a secret", c);
```

The original object that was sealed can be recovered in two different ways:

• by using a Cipher object that has been initialized with the exact same algorithm, key, padding scheme, etc., that were used to seal the object:

```
c.init(Cipher.DECRYPT_MODE, sKey);
try {
    String s = (String)so.getObject(c);
} catch (Exception e) {
    // do something
};
```

This approach has the advantage that the party who unseals the sealed object does not require knowledge of the decryption key. For example, after one party has initialized the cipher object with the required decryption key, it could hand over the cipher object to another party who then unseals the sealed object.

 by using the appropriate decryption key (since AES is a symmetric encryption algorithm, we use the same key for sealing and unsealing):

```
try {
    String s = (String)so.getObject(sKey);
} catch (Exception e) {
    // do something
};
```

In this approach, the <code>getObject</code> method creates a cipher object for the appropriate decryption algorithm and initializes it with the given decryption key and the algorithm parameters (if any) that were stored in the sealed object. This approach has the advantage that the party who unseals the object does not need to keep track of the parameters (e.g., the IV) that were used to seal the object.

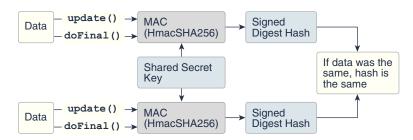


The Mac Class

Similar to a MessageDigest, a Message Authentication Code (MAC) provides a way to check the integrity of information transmitted over or stored in an unreliable medium, but includes a secret key in the calculation.

Only someone with the proper key will be able to verify the received message. Typically, message authentication codes are used between two parties that share a secret key in order to validate information transmitted between these parties.

Figure 2-9 The Mac Class



A MAC mechanism that is based on cryptographic hash functions is referred to as HMAC. HMAC can be used with any cryptographic hash function, e.g., SHA-256, in combination with a secret shared key.

The Mac class provides the functionality of a Message Authentication Code (MAC). See HMAC-SHA256 Example.

Creating a Mac Object

Mac objects are obtained by using one of the Mac getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing a Mac Object

A Mac object is always initialized with a (secret) key and may optionally be initialized with a set of parameters, depending on the underlying MAC algorithm.

To initialize a Mac object, call one of its init methods:

```
public void init(Key key);
public void init(Key key, AlgorithmParameterSpec params);
```

You can initialize your Mac object with any (secret-)key object that implements the javax.crypto.SecretKey interface. This could be an object returned by javax.crypto.KeyGenerator.generateKey(), or one that is the result of a key agreement protocol, as returned by

javax.crypto.KeyAgreement.generateSecret(), or an instance of javax.crypto.spec.SecretKeySpec.

With some MAC algorithms, the (secret-)key algorithm associated with the (secret-)key object used to initialize the Mac object does not matter (this is the case with the HMAC-MD5 and HMAC-SHA1 implementations of the <code>SunJCE</code> provider). With others,

however, the (secret-)key algorithm does matter, and an InvalidKeyException is thrown if a (secret-)key object with an inappropriate (secret-)key algorithm is used.

Computing a MAC

A MAC can be computed in one step (*single-part operation*) or in multiple steps (*multiple-part operation*). A multiple-part operation is useful if you do not know in advance how long the data is going to be, or if the data is too long to be stored in memory all at once.

To compute the MAC of some data in a single step, call the following doFinal method:

```
public byte[] doFinal(byte[] input);
```

To compute the MAC of some data in multiple steps, call one of the update methods:

```
public void update(byte input);
public void update(byte[] input);
public void update(byte[] input, int inputOffset, int inputLen);
```

A multiple-part operation must be terminated by the above doFinal method (if there is still some input data left for the last step), or by one of the following doFinal methods (if there is no input data left for the last step):

```
public byte[] doFinal();
public void doFinal(byte[] output, int outOffset);
```

Key Interfaces

The <code>java.security.Key</code> interface is the top-level interface for all opaque keys. It defines the functionality shared by all opaque key objects.

To this point, we have focused the high-level uses of the JCA without getting lost in the details of what keys are and how they are generated/represented. It is now time to turn our attention to keys.

An *opaque* key representation is one in which you have no direct access to the key material that constitutes a key. In other words: "opaque" gives you limited access to the key--just the three methods defined by the Key interface (see below): getAlgorithm, getFormat, and getEncoded.

This is in contrast to a *transparent* representation, in which you can access each key material value individually, through one of the get methods defined in the corresponding KeySpec interface (see The KeySpec Interface).

All opaque keys have three characteristics:

An Algorithm

The key algorithm for that key. The key algorithm is usually an encryption or asymmetric operation algorithm (such as AES, DSA or RSA), which will work with those



algorithms and with related algorithms (such as SHA256withRSA). The name of the algorithm of a key is obtained using this method:

String getAlgorithm()

An Encoded Form

The external encoded form for the key used when a standard representation of the key is needed outside the Java Virtual Machine, as when transmitting the key to some other party. The key is encoded according to a standard format (such as X.509 or PKCS8), and is returned using the method:

byte[] getEncoded()

A Format

The name of the format of the encoded key. It is returned by the method:

String getFormat()

Keys are generally obtained through key generators such as the KeyGenerator class and the KeyPairGenerator class, certificates, key specifications (see the The KeySpec Interface) using a KeyFactory, or a Keystore implementation accessing a keystore database used to manage keys. It is possible to parse encoded keys, in an algorithm-dependent manner, using a KeyFactory.

It is also possible to parse certificates, using a CertificateFactory.

Here is a list of interfaces which extend the Key interface in the java.security.interfaces and javax.crypto.interfaces packages:

- SecretKey
 - PBEKey
- PrivateKey
 - DHPrivateKey
 - DSAPrivateKey
 - ECPrivateKey
 - RSAMultiPrimePrivateCrtKey
 - RSAPrivateCrtKey
 - RSAPrivateKey
- PublicKey
 - DHPublicKey
 - DSAPublicKey
 - ECPublicKey
 - RSAPublicKey

The PublicKey and PrivateKey Interfaces

The Publickey and Privatekey interfaces (which both extend the key interface) are methodless interfaces, used for type-safety and type-identification.



The KeyPair Class

The KeyPair class is a simple holder for a key pair (a public key and a private key).

It has two public methods, one for returning the private key, and the other for returning the public key:

PrivateKey getPrivate()
PublicKey getPublic()

Key Specification Interfaces and Classes

 κ ey objects and key specifications (κ eyspecs) are two different representations of key data. Ciphers use κ ey objects to initialize their encryption algorithms, but keys may need to be converted into a more portable format for transmission or storage.

A *transparent* representation of keys means that you can access each key material value individually, through one of the <code>get</code> methods defined in the corresponding specification class. For example, <code>DSAPrivateKeySpec</code> defines <code>getX</code>, <code>getP</code>, <code>getQ</code>, and <code>getG</code> methods, to access the private key <code>x</code>, and the DSA algorithm parameters used to calculate the key: the prime <code>p</code>, the sub-prime <code>q</code>, and the base <code>g</code>. If the key is stored on a hardware device, its specification may contain information that helps identify the key on the device.

This representation is contrasted with an *opaque* representation, as defined by the Key Interfaces interface, in which you have no direct access to the key material fields. In other words, an "opaque" representation gives you limited access to the key--just the three methods defined by the Key interface: getAlgorithm, getFormat, and getEncoded.

A key may be specified in an algorithm-specific way, or in an algorithm-independent encoding format (such as ASN.1). For example, a DSA private key may be specified by its components x, p, q, and g (see DSAPrivateKeySpec), or it may be specified using its DER encoding (see PKCS8EncodedKeySpec).

The The KeyFactory Class and The SecretKeyFactory Class classes can be used to convert between opaque and transparent key representations (that is, between $\kappa_{\rm EyS}$ and $\kappa_{\rm EySpec}$ s, assuming that the operation is possible. (For example, private keys on smart cards might not be able leave the card. Such $\kappa_{\rm EyS}$ are not convertible.)

In the following sections, we discuss the key specification interfaces and classes in the java.security.spec package.

The KeySpec Interface

This interface contains no methods or constants. Its only purpose is to group and provide type safety for all key specifications. All key specifications must implement this interface.

The KeySpec Subinterfaces

Like the Key interface, there are a similar set of KeySpec interfaces.

• SecretKeySpec



- EncodedKeySpec
 - PKCS8EncodedKeySpec
 - X509EncodedKeySpec
- DESKeySpec
- DESedeKeySpec
- PBEKeySpec
- DHPrivateKeySpec
- DSAPrivateKeySpec
- ECPrivateKeySpec
- RSAPrivateKeySpec
 - RSAMultiPrimePrivateCrtKeySpec
 - RSAPrivateCrtKeySpec
- DHPublicKeySpec
- DSAPublicKeySpec
- ECPublicKeySpec
- RSAPublicKeySpec

The EncodedKeySpec Class

This abstract class (which implements the The KeySpec Interface interface) represents a public or private key in encoded format. Its getEncoded method returns the encoded key:

```
abstract byte[] getEncoded();
```

and its getFormat method returns the name of the encoding format:

```
abstract String getFormat();
```

See the next sections for the concrete implementations $\tt PKCS8EncodedKeySpec$ and $\tt X509EncodedKeySpec$.

The PKCS8EncodedKeySpec Class

This class, which is a subclass of <code>EncodedKeySpec</code>, represents the DER encoding of a private key, according to the format specified in the PKCS8 standard.

Its getEncoded method returns the key bytes, encoded according to the PKCS8 standard. Its getFormat method returns the string "PKCS#8".

The X509EncodedKeySpec Class

This class, which is a subclass of <code>EncodedKeySpec</code>, represents the DER encoding of a public key, according to the format specified in the X.509 standard.

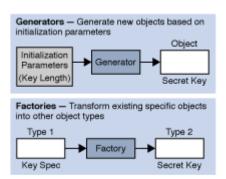
Its getEncoded method returns the key bytes, encoded according to the X.509 standard. Its getFormat method returns the string "X.509".



Generators and Factories

Newcomers to Java and the JCA APIs in particular sometimes do not grasp the distinction between generators and factories.

Figure 2-10 Generators and Factories



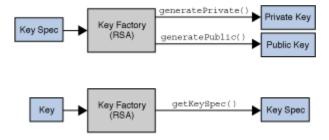
Generators are used to **generate brand new objects**. Generators can be initialized in either an algorithm-dependent or algorithm-independent way. For example, to create a Diffie-Hellman (DH) keypair, an application could specify the necessary P and G values, or the generator could simply be initialized with the appropriate key length, and the generator will select appropriate P and G values. In both cases, the generator will produce brand new keys based on the parameters.

On the other hand, factories are used to **convert data from one existing object type to another**. For example, an application might have available the components of a DH private key and can package them as a The KeySpec Interface, but needs to convert them into a PrivateKey object that can be used by a KeyAgreement object, or vice-versa. Or they might have the byte array of a certificate, but need to use a CertificateFactory to convert it into a X509Certificate object. Applications use factory objects to do the conversion.

The KeyFactory Class

The KeyFactory class is an Engine Classes and Algorithms designed to perform conversions between opaque cryptographic Key Interfaces and Key Specification Interfaces and Classes (transparent representations of the underlying key material).

Figure 2-11 KeyFactory Class





Key factories are bi-directional. They allow you to build an opaque key object from a given key specification (key material), or to retrieve the underlying key material of a key object in a suitable format.

Multiple compatible key specifications can exist for the same key. For example, a DSA public key may be specified by its components y, p, q, and g (see java.security.spec.DSAPublicKeySpec), or it may be specified using its DER encoding according to the X.509 standard (see The X509EncodedKeySpec Class).

A key factory can be used to translate between compatible key specifications. Key parsing can be achieved through translation between compatible key specifications, e.g., when you translate from x509EncodedKeySpec to DSAPublicKeySpec, you basically parse the encoded key into its components. For an example, see the end of the Generating/Verifying Signatures Using Key Specifications and KeyFactory section.

Creating a KeyFactory Object

KeyFactory objects are obtained by using one of the KeyFactorygetInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.

Converting Between a Key Specification and a Key Object

If you have a key specification for a public key, you can obtain an opaque PublicKey object from the specification by using the generatePublic method:

PublicKey generatePublic(KeySpec keySpec)

Similarly, if you have a key specification for a private key, you can obtain an opaque PrivateKey object from the specification by using the generatePrivate method:

PrivateKey generatePrivate(KeySpec keySpec)

Converting Between a Key Object and a Key Specification

If you have a Key object, you can get a corresponding key specification object by calling the getKeySpec method:

KeySpec getKeySpec(Key key, Class keySpec)

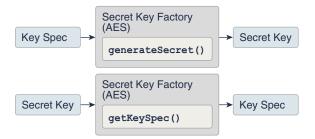
keySpec identifies the specification class in which the key material should be returned. It could, for example, be DSAPublicKeySpec.class, to indicate that the key material should be returned in an instance of the DSAPublicKeySpec class. See Generating/ Verifying Signatures Using Key Specifications and KeyFactory.

The SecretKeyFactory Class

The SecretKeyFactory class represents a factory for secret keys. Unlike the KeyFactory class (see The KeyFactory Class), a javax.crypto.SecretKeyFactory object operates only on secret (symmetric) keys, whereas a java.security.KeyFactory object processes the public and private key components of a key pair.



Figure 2-12 SecretKeyFactory Class



Key factories are used to convert Key Interfaces (opaque cryptographic keys of type <code>java.security.Key</code>) into Key Specification Interfaces and Classes (transparent representations of the underlying key material in a suitable format), and vice versa.

Objects of type <code>java.security.Key</code>, of which <code>java.security.PublicKey</code>, <code>java.security.PrivateKey</code>, and <code>javax.crypto.SecretKey</code> are subclasses, are opaque key objects, because you cannot tell how they are implemented. The underlying implementation is provider-dependent, and may be software or hardware based. Key factories allow providers to supply their own implementations of cryptographic keys.

For example, if you have a key specification for a Diffie-Hellman public key, consisting of the public value y, the prime modulus p, and the base g, and you feed the same specification to Diffie-Hellman key factories from different providers, the resulting PublicKey objects will most likely have different underlying implementations.

A provider should document the key specifications supported by its secret key factory. For example, the <code>SecretKeyFactory</code> for DES keys supplied by the <code>SumJCE</code> provider supports <code>DESKeySpec</code> as a transparent representation of DES keys, the <code>SecretKeyFactory</code> for DES-EDE keys supports <code>DESedeKeySpec</code> as a transparent representation of DES-EDE keys, and the <code>SecretKeyFactory</code> for PBE supports <code>PBEKeySpec</code> as a transparent representation of the underlying password.

The following is an example of how to use a SecretKeyFactory to convert secret key data into a SecretKey object, which can be used for a subsequent Cipher operation:

```
// Note the following bytes are not realistic secret key data
// bytes but are simply supplied as an illustration of using data
// bytes (key material) you already have to build a DESedeKeySpec.

byte[] desEdeKeyData = getKeyData();

DESedeKeySpec desEdeKeySpec = new DESedeKeySpec(desEdeKeyData);
SecretKeyFactory keyFactory = SecretKeyFactory.getInstance("DESede");
SecretKey secretKey = keyFactory.generateSecret(desEdeKeySpec);
```

In this case, the underlying implementation of <code>SecretKey</code> is based on the provider of <code>KeyFactory</code>.

An alternative, provider-independent way of creating a functionally equivalent SecretKey object from the same key material is to use the javax.crypto.spec.SecretKeySpec class, which implements the javax.crypto.SecretKey interface:

```
byte[] aesKeyData = getKeyData();
SecretKeySpec secretKey = new SecretKeySpec(aesKeyData, "AES");
```



Creating a SecretKeyFactory Object

SecretKeyFactory objects are obtained by using one of the SecretKeyFactory getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.

Converting Between a Key Specification and a Secret Key Object

If you have a key specification for a secret key, you can obtain an opaque SecretKey object from the specification by using the generateSecret method:

SecretKey generateSecret(KeySpec keySpec)

Converting Between a Secret Key Object and a Key Specification

If you have a SecretKey object, you can get a corresponding key specification object by calling the getKeySpec method:

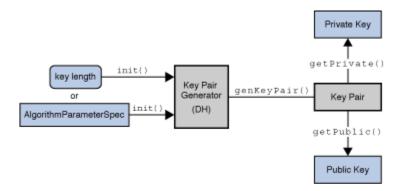
KeySpec getKeySpec(Key key, Class keySpec)

keySpec identifies the specification class in which the key material should be returned. It could, for example, be <code>DESKeySpec.class</code>, to indicate that the key material should be returned in an instance of the <code>DESKeySpec.class</code>.

The KeyPairGenerator Class

The KeyPairGenerator class is an engine class (see Engine Classes and Algorithms) used to generate pairs of public and private keys.

Figure 2-13 KeyPairGenerator Class



There are two ways to generate a key pair: in an algorithm-independent manner, and in an algorithm-specific manner. The only difference between the two is the initialization of the object.

See Generating a Pair of Keys for examples of calls to the methods documented below.

Creating a KeyPairGenerator

All key pair generation starts with a KeyPairGenerator. KeyPairGenerator Objects are obtained by using one of the KeyPairGenerator getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.



Initializing a KeyPairGenerator

A key pair generator for a particular algorithm creates a public/private key pair that can be used with this algorithm. It also associates algorithm-specific parameters with each of the generated keys.

A key pair generator needs to be initialized before it can generate keys. In most cases, algorithm-independent initialization is sufficient. But in other cases, algorithm-specific initialization can be used.

Algorithm-Independent Initialization

All key pair generators share the concepts of a keysize and a source of randomness. The keysize is interpreted differently for different algorithms. For example, in the case of the DSA algorithm, the keysize corresponds to the length of the modulus. (See Java Security Standard Algorithm Names Specification for information about the keysizes for specific algorithms.)

An initialize method takes two universally shared types of arguments:

```
void initialize(int keysize, SecureRandom random)
```

Another initialize method takes only a keysize argument; it uses a system-provided source of randomness:

```
void initialize(int keysize)
```

Since no other parameters are specified when you call the above algorithm-independent <code>initialize</code> methods, it is up to the provider what to do about the algorithm-specific parameters (if any) to be associated with each of the keys.

If the algorithm is a "DSA" algorithm, and the modulus size (keysize) is 512, 768, 1024, 2048, or 3072, then the $_{\rm SUN}$ provider uses a set of precomputed values for the $_{\rm P},~_{\rm Q},$ and $_{\rm g}$ parameters. If the modulus size is not one of the above values, the $_{\rm SUN}$ provider creates a new set of parameters. Other providers might have precomputed parameter sets for more than just the three modulus sizes mentioned above. Still others might not have a list of precomputed parameters at all and instead always create new parameter sets.

Algorithm-Specific Initialization

For situations where a set of algorithm-specific parameters already exists (such as "community parameters" in DSA), there are two initialize methods that have an The AlgorithmParameterSpec Interface argument. One also has a SecureRandom argument, while the source of randomness is system-provided for the other:

See Generating a Pair of Keys.

Generating a Key Pair

The procedure for generating a key pair is always the same, regardless of initialization (and of the algorithm). You always call the following method from <code>KeyPairGenerator</code>:

```
KeyPair generateKeyPair()
```

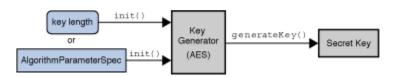


Multiple calls to generate KeyPair will yield different key pairs.

The KeyGenerator Class

A key generator is used to generate secret keys for symmetric algorithms.

Figure 2-14 The KeyGenerator Class



Creating a KeyGenerator

KeyGenerator objects are obtained by using one of the KeyGenerator getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing a KeyGenerator Object

A key generator for a particular symmetric-key algorithm creates a symmetric key that can be used with that algorithm. It also associates algorithm-specific parameters (if any) with the generated key.

There are two ways to generate a key: in an algorithm-independent manner, and in an algorithm-specific manner. The only difference between the two is the initialization of the object:

Algorithm-Independent Initialization

All key generators share the concepts of a *keysize* and a *source of randomness*. There is an init method that takes these two universally shared types of arguments. There is also one that takes just a keysize argument, and uses a system-provided source of randomness, and one that takes just a source of randomness:

```
public void init(SecureRandom random);
public void init(int keysize);
public void init(int keysize, SecureRandom random);
```

Since no other parameters are specified when you call the above algorithm-independent init methods, it is up to the provider what to do about the algorithm-specific parameters (if any) to be associated with the generated key.

Algorithm-Specific Initialization

For situations where a set of algorithm-specific parameters already exists, there are two init methods that have an AlgorithmParameterSpec argument. One also has a SecureRandom argument, while the source of randomness is system-provided for the other:

public void init(AlgorithmParameterSpec params);



public void init(AlgorithmParameterSpec params, SecureRandom random);

In case the client does not explicitly initialize the KeyGenerator (via a call to an init method), each provider must supply (and document) a default initialization.

Creating a Key

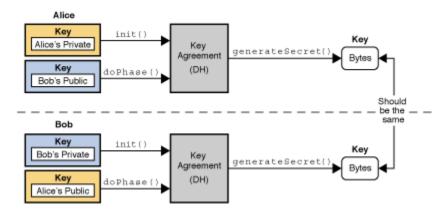
The following method generates a secret key:

public SecretKey generateKey();

The KeyAgreement Class

Key agreement is a protocol by which 2 or more parties can establish the same cryptographic keys, without having to exchange any secret information.

Figure 2-15 The KeyAgreement Class



Each party initializes their key agreement object with their private key, and then enters the public keys for each party that will participate in the communication. In most cases, there are just two parties, but algorithms such as Diffie-Hellman allow for multiple parties (3 or more) to participate. When all the public keys have been entered, each <code>keyAgreement</code> object will generate (agree upon) the same key.

The KeyAgreement class provides the functionality of a key agreement protocol. The keys involved in establishing a shared secret are created by one of the key generators (KeyPairGenerator Or KeyGenerator), a KeyFactory, or as a result from an intermediate phase of the key agreement protocol.

Creating a KeyAgreement Object

Each party involved in the key agreement has to create a KeyAgreement object.

KeyAgreement objects are obtained by using one of the KeyAgreement getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing a KeyAgreement Object

You initialize a KeyAgreement object with your private information. In the case of Diffie-Hellman, you initialize it with your Diffie-Hellman private key. Additional initialization information may contain a source of randomness and/or a set of algorithm



parameters. Note that if the requested key agreement algorithm requires the specification of algorithm parameters, and only a key, but no parameters are provided to initialize the KeyAgreement object, the key must contain the required algorithm parameters. (For example, the Diffie-Hellman algorithm uses a prime modulus ${\tt p}$ and a base generator ${\tt g}$ as its parameters.)

To initialize a KeyAgreement object, call one of its init methods:

Executing a KeyAgreement Phase

Every key agreement protocol consists of a number of phases that need to be executed by each party involved in the key agreement.

To execute the next phase in the key agreement, call the doPhase method:

```
public Key doPhase(Key key, boolean lastPhase);
```

The key parameter contains the key to be processed by that phase. In most cases, this is the public key of one of the other parties involved in the key agreement, or an intermediate key that was generated by a previous phase. doPhase may return an intermediate key that you may have to send to the other parties of this key agreement, so they can process it in a subsequent phase.

The lastPhase parameter specifies whether or not the phase to be executed is the last one in the key agreement: A value of FALSE indicates that this is not the last phase of the key agreement (there are more phases to follow), and a value of TRUE indicates that this is the last phase of the key agreement and the key agreement is completed, i.e., generateSecret can be called next.

In the example of Diffie-Hellman Key Exchange between 2 Parties , you call doPhase once, with lastPhase set to TRUE. In the example of Diffie-Hellman between three parties, you call doPhase twice: the first time with lastPhase set to FALSE, the 2nd time with lastPhase set to TRUE.

Generating the Shared Secret

After each party has executed all the required key agreement phases, it can compute the shared secret by calling one of the <code>generateSecret</code> methods:

```
public byte[] generateSecret();
public int generateSecret(byte[] sharedSecret, int offset);
public SecretKey generateSecret(String algorithm);
```



Key Management

A database called a "keystore" can be used to manage a repository of keys and certificates. (A *certificate* is a digitally signed statement from one entity, saying that the public key of some other entity has a particular value.)

Keystore Location

The user keystore is by default stored in a file named .keystore in the user's home directory, as determined by the user.home system property whose default value depends on the operating system:

- Solaris, Linux, and MacOS: /home/username/
- Windows: C:\Users\username\

Of course, keystore files can be located as desired. In some environments, it may make sense for multiple keystores to exist. For example, one keystore might hold a user's private keys, and another might hold certificates used to establish trust relationships.

In addition to the user's keystore, the JDK also maintains a system-wide keystore which is used to store trusted certificates from a variety of Certificate Authorities (CA's). These CA certificates can be used to help make trust decisions. For example, in SSL/TLS/DTLS when the SunJSSE provider is presented with certificates from a remote peer, the default trustmanager will consult one of the following files to determine if the connection is to be trusted:

- Solaris, Linux, and MacOS: < java-home > /lib/security/cacerts
- Windows: < java-home > \lib\security\cacerts

Instead of using the system-wide cacerts keystore, applications can set up and use their own keystores, or even use the user keystore described above.

Keystore Implementation

The KeyStore class supplies well-defined interfaces to access and modify the information in a keystore. It is possible for there to be multiple different concrete implementations, where each implementation is that for a particular *type* of keystore.

Currently, there are two command-line tools that make use of <code>KeyStore</code>: <code>keytool</code> and <code>jarsigner</code>, and also a GUI-based tool named <code>policytool</code>. It is also used by the <code>Policy</code> reference implementation when it processes policy files specifying the permissions (allowed accesses to system resources) to be granted to code from various sources. Since <code>KeyStore</code> is publicly available, JDK users can write additional security applications that use it.

Applications can choose different *types* of keystore implementations from different providers, using the <code>getInstance</code> factory method in the <code>KeyStore</code> class. A keystore type defines the storage and data format of the keystore information, and the algorithms used to protect private keys in the keystore and the integrity of the keystore itself. Keystore implementations of different types are not compatible.

The default keystore implementation is "pkcs12". This is a cross-platform keystore based on the RSA PKCS12 Personal Information Exchange Syntax Standard. This standard is primarily meant for storing or transporting a user's private keys,



certificates, and miscellaneous secrets. Arbitrary attributes can be associated with individual entries in a PKCS12 keystore.

keystore.type=pkcs12

To have tools and other applications use a different default keystore implementation, you can change that line to specify a different type.

Some applications, such as keytool, also let you override the default keystore type (via the -storetype command-line parameter).



Keystore type designations are case-insensitive. For example, "jks" would be considered the same as "JKS".

PKCS12 is the default and recommened keystore type. However, there are three other types of keystores that come with the JDK implementation.

1. "jceks" is an alternate proprietary keystore format to "jks" that uses Password-Based Encryption with Triple-DES.

The "jceks" implementation can parse and convert a "jks" keystore file to the "jceks" format. You may upgrade your keystore of type "jks" to a keystore of type "jceks" by changing the password of a private-key entry in your keystore and specifying "-storetype jceks" as the keystore type. To apply the cryptographically strong(er) key protection supplied to a private key named "signkey" in your default keystore, use the following command, which will prompt you for the old and new key passwords:

keytool -keypass -alias signkey -storetype jceks

See keytool in Java Platform, Standard Edition Tools Reference.

- 2. "jks" is another option. It implements the keystore as a file, utilizing a proprietary keystore type (format). It protects each private key with its own individual password, and also protects the integrity of the entire keystore with a (possibly different) password.
- 3. "dks" is a domain keystore. It is a collection of keystores presented as a single logical keystore. The keystores that comprise a given domain are specified by configuration data whose syntax is described in DomainLoadStoreParameter.

Keystore implementations are provider-based. If you want to write your own KeyStore implementations, see How to Implement a Provider in the Java Cryptography Architecture.

The KeyStore Class

The KeyStore class supplies well-defined interfaces to access and modify the information in a keystore.

The KeyStore class is an Engine Classes and Algorithms.



Alias Type Data

Brad Private Key/Certificate ...

Deb Secret Key ...

Milton Trusted Certificate ...

Duke Trusted Certificate ...

store() load()

Figure 2-16 KeyStore Class

This class represents an in-memory collection of keys and certificates. KeyStore manages two types of entries:

- Key Entry: This type of keystore entry holds very sensitive cryptographic key
 information, which must be protected from unauthorized access. Typically, a key
 stored in this type of entry is a secret key, or a private key accompanied by the
 certificate chain authenticating the corresponding public key.
 - Private keys and certificate chains are used by a given entity for selfauthentication using digital signatures. For example, software distribution organizations digitally sign JAR files as part of releasing and/or licensing software.
- **Trusted Certificate Entry**: This type of entry contains a single public key certificate belonging to another party. It is called a *trusted certificate* because the keystore owner trusts that the public key in the certificate indeed belongs to the identity identified by the *subject* (owner) of the certificate.

This type of entry can be used to authenticate other parties.

Each entry in a keystore is identified by an "alias" string. In the case of private keys and their associated certificate chains, these strings distinguish among the different ways in which the entity may authenticate itself. For example, the entity may authenticate itself using different certificate authorities, or using different public key algorithms.

Whether keystores are persistent, and the mechanisms used by the keystore if it is persistent, are not specified here. This convention allows use of a variety of techniques for protecting sensitive (e.g., private or secret) keys. Smart cards or other integrated cryptographic engines (SafeKeyper) are one option, and simpler mechanisms such as files may also be used (in a variety of formats).

The main KeyStore methods are described below.

Creating a KeyStore Object

KeyStore objects are obtained by using one of the KeyStore getInstance() method. See How Provider Implementations Are Requested and Supplied.



Loading a Particular Keystore into Memory

Before a KeyStore object can be used, the actual keystore data must be loaded into memory via the load method:

```
final void load(InputStream stream, char[] password)
```

The optional password is used to check the integrity of the keystore data. If no password is supplied, no integrity check is performed.

To create an empty keystore, you pass null as the InputStream argument to the load method.

A DKS keystore is loaded by passing a <code>DomainLoadStoreParameter</code> to the alternative load method:

```
final void load(KeyStore.LoadStoreParameter param)
```

Getting a List of the Keystore Aliases

All keystore entries are accessed via unique *aliases*. The aliases method returns an enumeration of the alias names in the keystore:

```
final Enumeration aliases()
```

Determining Keystore Entry Types

As stated in the KeyStore class, there are two different types of entries in a keystore. The following methods determine whether the entry specified by the given alias is a key/certificate or a trusted certificate entry, respectively:

```
final boolean isKeyEntry(String alias)
final boolean isCertificateEntry(String alias)
```

Adding/Setting/Deleting Keystore Entries

The setCertificateEntry method assigns a certificate to a specified alias:

```
final void setCertificateEntry(String alias, Certificate cert)
```

If alias doesn't exist, a trusted certificate entry with that alias is created. If alias exists and identifies a trusted certificate entry, the certificate associated with it is replaced by cert.

The setKeyEntry methods add (if alias doesn't yet exist) or set key entries:



In the method with key as a byte array, it is the bytes for a key in protected format. For example, in the keystore implementation supplied by the SUN provider, the key byte array is expected to contain a protected private key, encoded as an <code>EncryptedPrivateKeyInfo</code> as defined in the PKCS8 standard. In the other method, the password is the password used to protect the key.

The deleteEntry method deletes an entry:

```
final void deleteEntry(String alias)
```

PKCS #12 keystores support entries containing arbitrary attributes. Use the PKCS12Attribute class to create the attributes. When creating the new keystore entry use a constructor method that accepts attributes. Finally, use the following method to add the entry to the keystore:

Getting Information from the Keystore

The getkey method returns the key associated with the given alias. The key is recovered using the given password:

```
final Key getKey(String alias, char[] password)
```

The following methods return the certificate, or certificate chain, respectively, associated with the given alias:

```
final Certificate getCertificate(String alias)
final Certificate[] getCertificateChain(String alias)
```

You can determine the name (alias) of the first entry whose certificate matches a given certificate via the following:

```
final String getCertificateAlias(Certificate cert)
```

PKCS #12 keystores support entries containing arbitrary attributes. Use the following method to retrieve an entry that may contain attributes:

```
final Entry getEntry(String alias, ProtectionParameter protParam)
```

and then use the KeyStore. Entry. getAttributes method to extract such attributes and use the methods of the KeyStore. Entry. Attribute interface to examine them.

Saving the KeyStore

The in-memory keystore can be saved via the store method:

```
final void store(OutputStream stream, char[] password)
```

The password is used to calculate an integrity checksum of the keystore data, which is appended to the keystore data.



A DKS keystore is stored by passing a <code>DomainLoadStoreParameter</code> to the alternative store method:

final void store(KeyStore.LoadStoreParameter param)

Algorithm Parameters Classes

Like Keys and Keyspecs, an algorithm's initialization parameters are represented by either AlgorithmParameterS Or AlgorithmParameterSpecS.

Depending on the use situation, algorithms can use the parameters directly, or the parameters might need to be converted into a more portable format for transmission or storage.

A *transparent* representation of a set of parameters (via AlgorithmParameterSpec) means that you can access each parameter value in the set individually. You can access these values through one of the get methods defined in the corresponding specification class (e.g., DSAParameterSpec defines getP, getQ, and getG methods, to access p, q, and q, respectively).

In contrast, the The AlgorithmParameters Class class supplies an *opaque* representation, in which you have no direct access to the parameter fields. You can only get the name of the algorithm associated with the parameter set (via getAlgorithm) and some kind of encoding for the parameter set (via getEncoded).

The AlgorithmParameterSpec Interface

AlgorithmParameterSpec is an interface to a transparent specification of cryptographic parameters. This interface contains no methods or constants. Its only purpose is to group (and provide type safety for) all parameter specifications. All parameter specifications must implement this interface.

The following are the algorithm parameter specification interfaces and classes in the java.security.spec and javax.crypto.spec packages:

- DHParameterSpec
- DHGenParameterSpec
- DSAParameterSpec
- ECGenParameterSpec
- ECParameterSpec
- GCMParameterSpec
- IvParameterSpec
- MGF1ParameterSpec
- OAEPParameterSpec
- PBEParameterSpec
- PSSParameterSpec
- RC2ParameterSpec
- RC5ParameterSpec
- RSAKeyGenParameterSpec



The following algorithm parameter specs are used specifically for XML digital signatures.

- Interface C14NMethodParameterSpec
- Interface DigestMethodParameterSpec
- Interface SignatureMethodParameterSpec
- Interface TransformParameterSpec
- Interface ExcC14NParameterSpec
- Interface HMACParameterSpec
- Interface XPathFilter2ParameterSpec
- Interface XPathFilterParameterSpec
- XSLTTransformParameterSpec

The AlgorithmParameters Class

The AlgorithmParameters class provides an opaque representation of cryptographic parameters.

The AlgorithmParameters Class

The AlgorithmParameters class is an Engine Classes and Algorithms . You can initialize the AlgorithmParameters class using a specific AlgorithmParameterSpec object, or by encoding the parameters in a recognized format. You can retrieve the resulting specification with the getParameterSpec method (see the following section).

Creating an AlgorithmParameters Object

AlgorithmParameters objects are obtained by using one of the AlgorithmParameters getInstance() static factory methods. For more information, see How Provider Implementations Are Requested and Supplied.

Initializing an AlgorithmParameters Object

Once an AlgorithmParameters object is instantiated, it must be initialized via a call to init, using an appropriate parameter specification or parameter encoding:

```
void init(AlgorithmParameterSpec paramSpec)
void init(byte[] params)
void init(byte[] params, String format)
```

In these init methods, params is an array containing the encoded parameters, and format is the name of the decoding format. In the init method with a params argument but no format argument, the primary decoding format for parameters is used. The primary decoding format is ASN.1, if an ASN.1 specification for the parameters exists.

Obtaining the Encoded Parameters

A byte encoding of the parameters represented in an AlgorithmParameters object may be obtained via a call to getEncoded:

```
byte[] getEncoded()
```



This method returns the parameters in their primary encoding format. The primary encoding format for parameters is ASN.1, if an ASN.1 specification for this type of parameters exists.

If you want the parameters returned in a specified encoding format, use

byte[] getEncoded(String format)

If format is null, the primary encoding format for parameters is used, as in the other getEncoded method.

Converting an AlgorithmParameters Object to a Transparent Specification

A transparent parameter specification for the algorithm parameters may be obtained from an AlgorithmParameters object via a call to getParameterSpec:

AlgorithmParameterSpec getParameterSpec(Class paramSpec)

paramspec identifies the specification class in which the parameters should be returned. The specification class could be, for example, DSAParameterSpec.class to indicate that the parameters should be returned in an instance of the DSAParameterSpec class. (This class is in the java.security.spec package.)

The AlgorithmParameterGenerator Class

The AlgorithmParameterGenerator class is an Engine Classes and Algorithms used to generate a set of **brand-new** parameters suitable for a certain algorithm (the algorithm is specified when an AlgorithmParameterGenerator instance is created). This object is used when you do not have an existing set of algorithm parameters, and want to generate one from scratch.

Creating an AlgorithmParameterGenerator Object

AlgorithmParameterGenerator objects are obtained by using one of the AlgorithmParameterGenerator getInstance() static factory methods. See How Provider Implementations Are Requested and Supplied.

Initializing an AlgorithmParameterGenerator Object

The AlgorithmParameterGenerator object can be initialized in two different ways: an algorithm-independent manner or an algorithm-specific manner.

The algorithm-independent approach uses the fact that all parameter generators share the concept of a "size" and a source of randomness. The measure of size is universally shared by all algorithm parameters, though it is interpreted differently for different algorithms. For example, in the case of parameters for the DSA algorithm, "size" corresponds to the size of the prime modulus, in bits. (See Java Security Standard Algorithm Names Specification to know more about the sizes for specific algorithms.) When using this approach, algorithm-specific parameter generation values--if any--default to some standard values. One <code>init</code> method that takes these two universally shared types of arguments:

void init(int size, SecureRandom random);



Another init method takes only a size argument and uses a system-provided source of randomness:

```
void init(int size)
```

A third approach initializes a parameter generator object using algorithm-specific semantics, which are represented by a set of algorithm-specific parameter generation values supplied in an AlgorithmParameterSpec object:

To generate Diffie-Hellman system parameters, for example, the parameter generation values usually consist of the size of the prime modulus and the size of the random exponent, both specified in number of bits.

Generating Algorithm Parameters

Once you have created and initialized an AlgorithmParameterGenerator object, you can use the generateParameters method to generate the algorithm parameters:

AlgorithmParameters generateParameters()

The CertificateFactory Class

The CertificateFactory class defines the functionality of a certificate factory, which is used to generate certificate and certificate revocation list (CRL) objects from their encoding.

The CertificateFactory class is an Engine Classes and Algorithms.

A certificate factory for X.509 must return certificates that are an instance of java.security.cert.X509Certificate, and CRLs that are an instance of java.security.cert.X509CRL.

Creating a CertificateFactory Object

CertificateFactory objects are obtained by using one of the <code>getInstance()</code> static factory methods. For more information, see How Provider Implementations Are Requested and Supplied.

Generating Certificate Objects

To generate a certificate object and initialize it with the data read from an input stream, use the <code>generateCertificate</code> method:

final Certificate generateCertificate(InputStream inStream)

To return a (possibly empty) collection view of the certificates read from a given input stream, use the <code>generateCertificates</code> method:

final Collection generateCertificates(InputStream inStream)



Generating CRL Objects

To generate a certificate revocation list (CRL) object and initialize it with the data read from an input stream, use the <code>generateCRL</code> method:

```
final CRL generateCRL(InputStream inStream)
```

To return a (possibly empty) collection view of the CRLs read from a given input stream, use the <code>generateCRLs</code> method:

final Collection generateCRLs(InputStream inStream)

Generating CertPath Objects

The certificate path builder and validator for PKIX is defined by the Internet X.509 Public Key Infrastructure Certificate and CRL Profile, RFC 5280.

A certificate store implementation for retrieving certificates and CRLs from Collection and LDAP directories, using the PKIX LDAP V2 Schema is also available from the IETF as RFC 2587.

To generate a <code>certPath</code> object and initialize it with data read from an input stream, use one of the following <code>generateCertPath</code> methods (with or without specifying the encoding to be used for the data):

To generate a CertPath object and initialize it with a list of certificates, use the following method:

```
final CertPath generateCertPath(List certificates)
```

To retrieve a list of the CertPath encoding supported by this certificate factory, you can call the getCertPathEncodings method:

```
final Iterator getCertPathEncodings()
```

The default encoding will be listed first.

How the JCA Might Be Used in a SSL/TLS Implementation

With an understanding of the JCA classes, consider how these classes might be combined to implement an advanced network protocol like SSL/TLS.

The SSL/TLS Overview section in the SSL, TLS, and DTLS Protocols describes at a high level how the protocols work. As asymmetric (public key) cipher operations are much slower than symmetric operations (secret key), public key cryptography is used to establish secret keys which are then used to protect the actual application data. Vastly simplified, the SSL/TLS handshake involves exchanging initialization data,



performing some public key operations to arrive at a secret key, and then using that key to encrypt further traffic.



The details presented here simply show how some of the above classes might be employed. This section will not present sufficient information for building a SSL/TLS implementation. For more information, see Java Secure Socket Extension (JSSE) Reference Guide and RFC 5246: The Transport Layer Security (TLS) Protocol, Version 1.2.

Assume that this SSL/TLS implementation will be made available as a JSSE provider. A concrete implementation of the Provider class is first written that will eventually be registered in the Security class' list of providers. This provider mainly provides a mapping from algorithm names to actual implementation classes. (that is: "SSLContext.TLS"->"com.foo.TLSImpl") When an application requests an "TLS" instance (via SSLContext.getInstance("TLS")), the provider's list is consulted for the requested algorithm, and an appropriate instance is created.

Before discussing details of the actual handshake, a quick review of some of the JSSE's architecture is needed. The heart of the JSSE architecture is the SSLContext. The context eventually creates end objects (SSLSocket and SSLEngine) which actually implement the SSL/TLS protocol. SSLContexts are initialized with two callback classes, KeyManager and TrustManager, which allow applications to first select authentication material to send and second to verify credentials sent by a peer.

A JSSE KeyManager is responsible for choosing which credentials to present to a peer. Many algorithms are possible, but a common strategy is to maintain a RSA or DSA public/private key pair along with a X509Certificate in a KeyStore backed by a disk file. When a KeyStore object is initialized and loaded from the file, the file's raw bytes are converted into PublicKey and PrivateKey objects using a KeyFactory, and a certificate chain's bytes are converted using a CertificateFactory. When a credential is needed, the KeyManager simply consults this KeyStore object and determines which credentials to present.

A KeyStore's contents might have originally been created using a utility such as keytool. keytool creates a RSA or DSA KeyPairGenerator and initializes it with an appropriate keysize. This generator is then used to create a KeyPair which keytool would store along with the newly-created certificate in the KeyStore, which is eventually written to disk.

A JSSE TrustManager is responsible for verifying the credentials received from a peer. There are many ways to verify credentials: one of them is to create a CertPath object, and let the JDK's built-in Public Key Infrastructure (PKI) framework handle the validation. Internally, the CertPath implementation might create a Signature object, and use that to verify that the each of the signatures in the certificate chain.

With this basic understanding of the architecture, we can look at some of the steps in the SSL/TLS handshake. The client begins by sending a ClientHello message to the server. The server selects a ciphersuite to use, and sends that back in a ServerHello message, and begins creating JCA objects based on the suite selection. We'll use server-only authentication in the following examples.



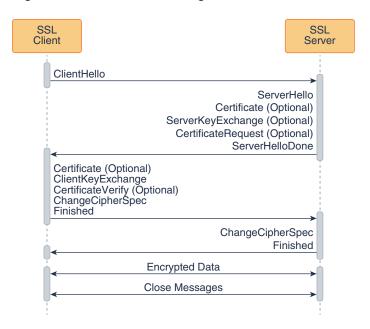


Figure 2-17 SSL/TLS Messages

Server-only authentication is described in the following examples. The examples are vastly simplified, but gives an idea of how the JSSE classes might be combined to create a higher level protocol:

Example 2-9 SSL/TLS Server Uses a RSA-based ciphersuite Such as TLS RSA WITH AES 128 CBC SHA

 ${\tt KeyManagerTrustManagerSecureRandomCipherPublicKeyPrivateKeyCipher}$

Example 2-10 Choose an Ephemeral Diffie-Hellman Key Agreement Algorithm Along with the DSA Signature Algorithm such as TLS DHE DSS WITH AES 128 CBC SHA

 ${\tt KeyPairGeneratorKeyFactoryDHPublicKeySpecKeyAgreementPrivateKeyServerKeyExchangeKeyFactoryKeyAgreement}$

Once the actual encryption keys have been established, the secret key is used to initialize a symmetric <code>cipher</code> object, and this cipher is used to protect all data in transit. To help determine if the data has been modified, a <code>MessageDigest</code> is created and receives a copy of the data destined for the network. When the packet is complete, the digest (hash) is appended to data, and the entire packet is encrypted by the <code>cipher</code>. If a block cipher such as AES is used, the data must be padded to make a complete block. On the remote side, the steps are simply reversed.

Cryptographic Strength Configuration

You can configure the cryptographic strength of the Java Cryptography Extension (JCE) architecture using jurisdiction policy files (see <u>Jurisdiction Policy File Format</u>) and the security properties file.

Prior to Oracle Java JDK 9, the default cryptographic strength allowed by Oracle implementations was "strong but limited" (for example AES keys limited to 128 bits). To remove this restriction, administrators could download and install a separate

"unlimited strength" Jurisdiction Policy Files bundle. The Jurisdiction Policy File mechanism was reworked for JDK 9. It now allows for much more flexible configuration. The Oracle JDK now ships with a default value of "unlimited" rather than "limited". As always, administrators and users must still continue to follow all import/export guidelines for their geographical locations. The active cryptographic strength is now determined using a Security Property (typically set in the <code>java.security</code> properties file), in combination with the jurisdiction policy files found in the configuration directory.

All the necessary JCE policy files to provide either unlimited cryptographic strength or strong but limited cryptographic strength are bundled with the JDK.

Cryptographic Strength Settings

Each directory under

// security/policy represents a set of policy
configurations defined by the jurisdiction policy files that they contain. You activate a
particular cryptographic strength setting represented by the policy files in a directory by
setting the crypto.policy Security Property (configured in the file java_home>/conf/
security/java.security) to point to that directory.

The JDK comes bundled with two such directories, limited and unlimited, each containing a number of policy files. By default, the <code>crypto.policy</code> Security Property is set to:

```
crypto.policy = unlimited
```

The overall value is the intersection of the files contained within the directory. These policy files settings are VM-wide, and affect all applications running on this VM. If you want to override cryptographic strength at the application level, see How to Make Applications Exempt from Cryptographic Restrictions.

Unlimited Directory Contents

The unlimited directory contains the following policy files:

<pre

Note:

As there are no current restrictions on export of cryptography from the United States, the default US export.policy file is set with no restrictions.

<p



Note:

Depending on the country, there may be local restrictions, but as this policy file is located in the unlimited directory, there are no restrictions listed here.

To select unlimited cryptographic strength as defined in these two files set crypto.policy = unlimited in the file < java_home > /conf/security/java.security.

Limited Directory Contents

The limited directory currently contains the following policy files:

<pre

Note:

Even though this is in the limited directory, as there are no current restrictions on export of cryptography from the United States, the default_US_export.policy file is set with no restrictions.

/java_home>/conf/security/limited/default_local.policy

Note:

This local policy file shows the default restrictions. It should be allowed by any country, including those that have import restrictions, but please obtain legal guidance.

```
// Some countries have import limits on crypto strength, but may allow for
// these exemptions if the exemption mechanism is used.
grant {
```



```
// There is no restriction to any algorithms if KeyRecovery is enforced.
permission javax.crypto.CryptoPermission *, "KeyRecovery";

// There is no restriction to any algorithms if KeyEscrow is enforced.
permission javax.crypto.CryptoPermission *, "KeyEscrow";

// There is no restriction to any algorithms if KeyWeakening is enforced.
permission javax.crypto.CryptoPermission *, "KeyWeakening";
};
```

Note:

Countries that have import restrictions should use "limited", but these restrictions could be relaxed if the exemption mechanism can be employed. See How to Make Applications Exempt from Cryptographic Restrictions. Please obtain legal guidance for your situation.

Custom Cryptographic Strength Settings

To set up restrictions to cryptographic strength that are different than the settings in the policy files in the limited or unlimited directory, you can create a new directory, parallel with limited and unlimited, and place your policy files there. For example, you may create a directory called custom. In this custom directory you include the files default_*export.policy and/or exempt_*local.policy.

To select cryptographic strength as defined in the files in the custom directory, set crypto.policy = custom in the file <java_home>/conf/security/java.security.

Jurisdiction Policy File Format

JCA represents its jurisdiction policy files as Java-style policy files with corresponding permission statements. As described in Cryptographic Strength Configuration, a Java policy file specifies what permissions are allowed for code from specified code sources. A permission represents access to a system resource. In the case of JCA, the "resources" are cryptography algorithms, and code sources do not need to be specified, because the cryptographic restrictions apply to all code.

A jurisdiction policy file consists of a very basic "grant entry" containing one or more "permission entries."

```
grant {
     <permission entries>;
};
```

The format of a permission entry in a jurisdiction policy file is:



```
];
```

A sample jurisdiction policy file that includes restricting the AES algorithm to maximum key sizes of 128 bits is:

```
grant {
    permission javax.crypto.CryptoPermission "AES", 128;
    // ...
};
```

A permission entry must begin with the word permission. Items that appear in a permission entry must appear in the specified order. An entry is terminated with a semicolon. Case is unimportant for the identifiers (grant, permission) but is significant for the <crypto permission class name> or for any string that is passed in as a value. An asterisk (*) can be used as a wildcard for any permission entry option. For example, an asterisk for an <alg_name> option means "all algorithms."

The following table describes a permission entry's options:

Table 2-1 Permission Entry Options

| Option | Description |
|-------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre><crypto class="" name="" permission=""></crypto></pre> | Specific permission class name, such as javax.crypto.CryptoPermission. Required. |
| | A crypto permission class reflects the ability of an application to use certain algorithms with certain key sizes in certain environments. There are two crypto permission classes: CryptoPermission and CryptoAllPermission. The special CryptoAllPermission class implies all cryptography-related permissions, that is, it specifies that there are no cryptography-related restrictions. |
| <alg_name></alg_name> | Quoted string specifying the standard name of
a cryptography algorithm, such as "AES" or
"RSA". Optional. |
| <exemption mechanism="" name=""></exemption> | Quoted string indicating an exemption mechanism which, if enforced, enables a reduction in cryptographic restrictions. Optional. |
| | Exemption mechanism names that can be used include "KeyRecovery" "KeyEscrow", and "KeyWeakening". |
| <maxkeysize></maxkeysize> | Integer specifying the maximum key size (in bits) allowed for the specified algorithm. Optional. |



Table 2-1 (Cont.) Permission Entry Options

Option	Description
<pre><algorithmparameterspec class="" name=""></algorithmparameterspec></pre>	Class name that specifies the strength of the algorithm. Optional.
	For some algorithms, it may not be sufficient to specify the algorithm strength in terms of just a key size. For example, in the case of the "RC5" algorithm, the number of rounds must also be considered. For algorithms whose strength needs to be expressed as more than a key size, use this option to specify the AlgorithmParameterSpec class name that does this (such as javax.crypto.spec.RC5ParameterSpec for the "RC5" algorithm).
<pre><parameters algorithmparameterspec="" an="" constructing="" for="" object=""></parameters></pre>	List of parameters for constructing the specified AlgorithmParameterSpec object. Required if <algorithmparameterspec class="" name=""> has been specified and requires parameters.</algorithmparameterspec>

How to Make Applications Exempt from Cryptographic Restrictions

NOT_SUPPORTED:

This section should be ignored by most application developers. It is only for people whose applications may be exported to those few countries whose governments mandate cryptographic restrictions, if it is desired that such applications have fewer cryptographic restrictions than those mandated.

By default, an application can use cryptographic algorithms of any strength. However, due to import control restrictions by the governments of a few countries, you may have to limit those algorithms' strength. The JCA framework includes an ability to enforce restrictions regarding the maximum strengths of cryptographic algorithms available to applications in different jurisdiction contexts (locations). You specify these restrictions in jurisdiction policy files. For more information about jurisdiction policy files and how to create and configure them, see Cryptographic Strength Configuration.

It is possible that the governments of some or all such countries may allow certain applications to become exempt from some or all cryptographic restrictions. For example, they may consider certain types of applications as "special" and thus exempt. Or they may exempt any application that utilizes an "exemption mechanism," such as key recovery. Applications deemed to be exempt could get access to stronger cryptography than that allowed for non-exempt applications in such countries.

In order for an application to be recognized as "exempt" at runtime, it must meet the following conditions:



- It must have a permission policy file bundled with it in a JAR file. The permission policy file specifies what cryptography-related permissions the application has, and under what conditions (if any).
- The JAR file containing the application and the permission policy file must have been signed using a code-signing certificate issued after the application was accepted as exempt.

Below are sample steps required in order to make an application exempt from some cryptographic restrictions. This is a basic outline that includes information about what is required by JCA in order to recognize and treat applications as being exempt. You will need to know the exemption requirements of the particular country or countries in which you would like your application to be able to be run but whose governments require cryptographic restrictions. You will also need to know the requirements of a JCA framework vendor that has a process in place for handling exempt applications. Consult such a vendor for further information.



The ${\tt SunJCE}$ provider does not supply an implementation of the ${\tt ExemptionMechanismSpi}$ class

- 1. Write and Compile Your Application Code
- 2. Create a Permission Policy File Granting Appropriate Cryptographic Permissions
- 3. Prepare for Testing
 - a. Apply for Government Approval From the Government Mandating Restrictions.
 - b. Get a Code-Signing Certificate
 - c. Bundle the Application and Permission Policy File into a JAR file
 - d. Step 7.1: Get a Code-Signing Certificate
 - e. Set Up Your Environment Like That of a User in a Restricted Country
 - f. (only for applications using exemption mechanisms) Install a Provider Implementing the Exemption Mechanism Specified by the entry in the Permission Policy File
- 4. Test Your Application
- 5. Apply for U.S. Government Export Approval If Required
- 6. Deploy Your Application

Special Code Requirements for Applications that Use Exemption Mechanisms

When an application has a permission policy file associated with it (in the same JAR file) and that permission policy file specifies an exemption mechanism, then when the Cipher getInstance method is called to instantiate a Cipher, the JCA code searches the installed providers for one that implements the specified exemption mechanism. If it finds such a provider, JCA instantiates an ExemptionMechanism API object associated with the provider's implementation, and then associates the ExemptionMechanism object with the Cipher returned by getInstance.

After instantiating a Cipher, and prior to initializing it (via a call to the Cipher init method), your code must call the following Cipher method:



```
public ExemptionMechanism getExemptionMechanism()
```

This call returns the ExemptionMechanism object associated with the Cipher. You must then initialize the exemption mechanism implementation by calling the following method on the returned ExemptionMechanism:

```
public final void init(Key key)
```

The argument you supply should be the same as the argument of the same types that you will subsequently supply to a Cipher init method.

Once you have initialized the ExemptionMechanism, you can proceed as usual to initialize and use the Cipher.

Permission Policy Files

In order for an application to be recognized at runtime as being "exempt" from some or all cryptographic restrictions, it must have a permission policy file bundled with it in a JAR file. The permission policy file specifies what cryptography-related permissions the application has, and under what conditions (if any).

The format of a permission entry in a permission policy file that accompanies an exempt application is the same as the format for a jurisdiction policy file downloaded with the JDK, which is:

See Jurisdiction Policy File Format.

Permission Policy Files for Exempt Applications

Some applications may be allowed to be completely unrestricted. Thus, the permission policy file that accompanies such an application usually just needs to contain the following:

```
grant {
    // There are no restrictions to any algorithms.
    permission javax.crypto.CryptoAllPermission;
};
```

If an application just uses a single algorithm (or several specific algorithms), then the permission policy file could simply mention that algorithm (or algorithms) explicitly, rather than granting CryptoAllPermission.

For example, if an application just uses the Blowfish algorithm, the permission policy file doesn't have to grant CryptoAllPermission to all algorithms. It could just specify that there is no cryptographic restriction if the Blowfish algorithm is used. In order to do this, the permission policy file would look like the following:



```
grant {
    permission javax.crypto.CryptoPermission "Blowfish";
};
```

Permission Policy Files for Applications Exempt Due to Exemption Mechanisms

If an application is considered "exempt" if an exemption mechanism is enforced, then the permission policy file that accompanies the application must specify one or more exemption mechanisms. At run time, the application will be considered exempt if any of those exemption mechanisms is enforced. Each exemption mechanism must be specified in a permission entry that looks like the following:

```
// No algorithm restrictions if specified
// exemption mechanism is enforced.
permission javax.crypto.CryptoPermission *,
    "<ExemptionMechanismName>";
```

where <ExemptionMechanismName> specifies the name of an exemption mechanism. The list of possible exemption mechanism names includes:

- KeyRecovery
- KeyEscrow
- KeyWeakening

As an example, suppose your application is exempt if either key recovery or key escrow is enforced. Then your permission policy file should contain the following:

```
grant {
    // No algorithm restrictions if KeyRecovery is enforced.
    permission javax.crypto.CryptoPermission *, "KeyRecovery";

    // No algorithm restrictions if KeyEscrow is enforced.
    permission javax.crypto.CryptoPermission *, "KeyEscrow";
};
```

Note:

Permission entries that specify exemption mechanisms should *not* also specify maximum key sizes. The allowed key sizes are actually determined from the installed exempt jurisdiction policy files, as described in the next section.

How Bundled Permission Policy Files Affect Cryptographic Permissions

At runtime, when an application instantiates a Cipher (via a call to its getInstance method) and that application has an associated permission policy file, JCA checks to see whether the permission policy file has an entry that applies to the algorithm specified in the getInstance call. If it does, and the entry grants CryptoAllPermission or does not specify that an exemption mechanism must be enforced, it means there is no cryptographic restriction for this particular algorithm.

If the permission policy file has an entry that applies to the algorithm specified in the <code>getInstance</code> call and the entry does specify that an exemption mechanism must be enforced, then the exempt jurisdiction policy file(s) are examined. If the exempt permissions include an entry for the relevant algorithm and exemption mechanism, and that entry is implied by the permissions in the permission policy file bundled with the application, and if there is an implementation of the specified exemption



mechanism available from one of the registered providers, then the maximum key size and algorithm parameter values for the Cipher are determined from the exempt permission entry.

If there is no exempt permission entry implied by the relevant entry in the permission policy file bundled with the application, or if there is no implementation of the specified exemption mechanism available from any of the registered providers, then the application is only allowed the standard default cryptographic permissions.

Standard Names

The Standard Names document contains information about the algorithm specifications.

Java Security Standard Algorithm Names describes the standard names for algorithms, certificate and keystore types that the JDK Security API requires and uses. It also contains more information about the algorithm specifications. Specific provider information can be found in the JDK Providers.

Cryptographic implementations in the JDK are distributed through several different providers primarily for historical reasons (Sun, SunJSSE, SunJCE, SunRsaSign). Note these providers may not be available on all JDK implementations, and therefore, truly portable applications should call getInstance() without specifying specific providers. Applications specifying a particular provider may not be able to take advantage of native providers tuned for an underlying operating environment (such as PKCS or Microsoft's CAPI).

The SunPKCS11 provider itself does not contain any cryptographic algorithms, but instead, directs requests into an underlying PKCS11 implementation. Consult PKCS#11 Reference Guide and the underlying PKCS11 implementation to determine if a desired algorithm will be available through the PKCS11 provider. Likewise, on Windows systems, the SunMSCAPI provider does not provide any cryptographic functionality, but instead routes requests to the underlying Operating System for handling.

Packaging Your Application

You can package an application in three different kinds of modules:

- Named or explicit module: A module that appears on the module path and contains module configuration information in the module-info.class file.
- Automatic module: A module that appears on the module path, but does not contain module configuration information in a module-info.class file (essentially a "regular" JAR file).
- Unnamed module: A module that appears on the class path. It may or may not have a module-info.class file; this file is ignored.

It is recommended that you package your applications in named modules as they provide better performance, stronger encapsulation, and simpler configuration. They also offer greater flexibility; you can use them with non-modular JDKs or even as unnamed modules by specifying them in a modular JDK's class path.

For more information about modules, see The State of the Module System and JEP 261: Module System



Additional JCA Code Samples

These examples illustrate use of several JCA mechanisms. See also Sample Programs for Diffie-Hellman Key Exchange, AES/GCM, and HMAC-SHA256

Topics

Computing a MessageDigest Object

Generating a Pair of Keys

Generating and Verifying a Signature Using Generated Keys

Generating/Verifying Signatures Using Key Specifications and KeyFactory

Determining If Two Keys Are Equal

Reading Base64-Encoded Certificates

Parsing a Certificate Reply

Using Encryption

Using Password-Based Encryption

Computing a MessageDigest Object

An example describing the procedure to compute a MessageDigest object.

1. Create the MessageDigest object, as in the following example:

```
MessageDigest sha = MessageDigest.getInstance("SHA-256");
```

This call assigns a properly initialized message digest object to the $_{\rm sha}$ variable. The implementation implements the Secure Hash Algorithm (SHA-256), as defined in the National Institute for Standards and Technology's (NIST) FIPS 180-2 document.

2. Suppose we have three byte arrays, i1, i2 and i3, which form the total input whose message digest we want to compute. This digest (or "hash") could be calculated via the following calls:

```
sha.update(i1);
sha.update(i2);
sha.update(i3);
byte[] hash = sha.digest();
```

3. Optional: An equivalent alternative series of calls would be:

```
sha.update(i1);
sha.update(i2);
byte[] hash = sha.digest(i3);
```

After the message digest has been calculated, the message digest object is automatically reset and ready to receive new data and calculate its digest. All former state (i.e., the data supplied to update calls) is lost.

Example 2-11 Hash Implementations Through Cloning

Some hash implementations may support intermediate hashes through cloning. Suppose we want to calculate separate hashes for:

```
i1i1 and i2i1, i2, and i3
```

The following is one way to calculate these hashes; however, this code works only if the SHA-256 implementation is cloneable:

```
/* compute the hash for i1 */
sha.update(i1);
byte[] i1Hash = sha.clone().digest();
/* compute the hash for i1 and i2 */
sha.update(i2);
byte[] i12Hash = sha.clone().digest();
/* compute the hash for i1, i2 and i3 */
sha.update(i3);
byte[] i123hash = sha.digest();
```

Example 2-12 Determine if the Hash Implementation is Cloneable or not Cloneable

```
MessageDigest

try {
    // try and clone it
    /* compute the hash for il */
    sha.update(il);
    byte[] ilHash = sha.clone().digest();
    // ...
    byte[] i123hash = sha.digest();
} catch (CloneNotSupportedException cnse) {
    // do something else, such as the code shown below
}
```

Example 2-13 Compute Intermediate Digests if the Hash Implementation is not Cloneable

```
MessageDigest md1 = MessageDigest.getInstance("SHA-256");
MessageDigest md2 = MessageDigest.getInstance("SHA-256");
MessageDigest md3 = MessageDigest.getInstance("SHA-256");
byte[] i1Hash = md1.digest(i1);
md2.update(i1);
byte[] i12Hash = md2.digest(i2);
md3.update(i1);
md3.update(i2);
byte[] i123Hash = md3.digest(i3);
```

Generating a Pair of Keys

In this example we will generate a public-private key pair for the algorithm named "DSA" (Digital Signature Algorithm), and use this keypair in future examples. We will

generate keys with a 2048-bit modulus. We don't care which provider supplies the algorithm implementation.

Creating the Key Pair Generator

The first step is to get a key pair generator object for generating keys for the DSA algorithm:

```
KeyPairGenerator keyGen = KeyPairGenerator.getInstance("DSA");
```

Initializing the Key Pair Generator

The next step is to initialize the key pair generator. In most cases, algorithm-independent initialization is sufficient, but in some cases, algorithm-specific initialization is used.

Algorithm-Independent Initialization

All key pair generators share the concepts of a keysize and a source of randomness. The KeyPairGenerator class initialization methods at a minimum needs a keysize. If the source of randomness is not explicitly provided, a SecureRandom implementation of the highest-priority installed provider will be used. Thus to generate keys with a keysize of 2048, simply call:

```
keyGen.initialize(2048);
```

The following code illustrates how to use a specific, additionally seeded SecureRandom object:

```
SecureRandom random = SecureRandom.getInstance("DRBG", "SUN");
random.setSeed(userSeed);
keyGen.initialize(2048, random);
```

Since no other parameters are specified when you call the above algorithm-independent initialize method, it is up to the provider what to do about the algorithm-specific parameters (if any) to be associated with each of the keys. The provider may use precomputed parameter values or may generate new values.

Algorithm-Specific Initialization

For situations where a set of algorithm-specific parameters already exists (such as "community parameters" in DSA), there are two <code>initialize</code> methods that have an <code>AlgorithmParameterSpec</code> argument. Suppose your key pair generator is for the "DSA" algorithm, and you have a set of DSA-specific parameters, <code>p</code>, <code>q</code>, and <code>g</code>, that you would like to use to generate your key pair. You could execute the following code to initialize your key pair generator (recall that <code>DSAParameterSpec</code> is an <code>AlgorithmParameterSpec</code>):

```
DSAParameterSpec dsaSpec = new DSAParameterSpec(p, q, g);
keyGen.initialize(dsaSpec);
```

Generating the Pair of Keys

The final step is actually generating the key pair. No matter which type of initialization was used (algorithm-independent or algorithm-specific), the same code is used to generate the KeyPair:

```
KeyPair pair = keyGen.generateKeyPair();
```



Generating and Verifying a Signature Using Generated Keys

Examples of generating and verifying a signature using generated keys.

The following signature generation and verification examples use the KeyPair generated in the Generating a Pair of Keys.

Generating a Signature

We first create a Signature Class object:

```
Signature dsa = Signature.getInstance("SHA256withDSA");
```

Next, using the key pair generated in the key pair example, we initialize the object with the private key, then sign a byte array called data.

```
/* Initializing the object with a private key */
PrivateKey priv = pair.getPrivate();
dsa.initSign(priv);

/* Update and sign the data */
dsa.update(data);
byte[] sig = dsa.sign();
```

Verifying a Signature

Verifying the signature is straightforward. (Note that here we also use the key pair generated in the key pair example.)

```
/* Initializing the object with the public key */
PublicKey pub = pair.getPublic();
dsa.initVerify(pub);

/* Update and verify the data */
dsa.update(data);
boolean verifies = dsa.verify(sig);
System.out.println("signature verifies: " + verifies);
```

Generating/Verifying Signatures Using Key Specifications and KeyFactory

Sample code that is used to generate and verify signatures using key specifications and <code>KeyFactory</code>.

Suppose that, rather than having a public/private key pair (as, for example, was generated in the Generating a Pair of Keys above), you simply have the components of your DSA private key: x (the private key), p (the prime), q (the sub-prime), and p (the base).

Further suppose you want to use your private key to digitally sign some data, which is in a byte array named <code>someData</code>. You would do the following steps, which also illustrate creating a key specification and using a key factory to obtain a <code>PrivateKey</code> from the key specification (initSign requires a <code>PrivateKey</code>):



```
DSAPrivateKeySpec dsaPrivKeySpec = new DSAPrivateKeySpec(x, p, q, g);

KeyFactory keyFactory = KeyFactory.getInstance("DSA");

PrivateKey privKey = keyFactory.generatePrivate(dsaPrivKeySpec);

Signature sig = Signature.getInstance("SHA256withDSA");

sig.initSign(privKey);

sig.update(someData);

byte[] signature = sig.sign();
```

Suppose Alice wants to use the data you signed. In order for her to do so, and to verify your signature, you need to send her three things:

- 1. The data
- 2. The signature
- 3. The public key corresponding to the private key you used to sign the data

You can store the someData bytes in one file, and the signature bytes in another, and send those to Alice.

For the public key, assume, as in the signing example above, you have the components of the DSA public key corresponding to the DSA private key used to sign the data. Then you can create a DSAPublicKeySpec from those components:

```
DSAPublicKeySpec dsaPubKeySpec = new DSAPublicKeySpec(y, p, q, g);
```

You still need to extract the key bytes so that you can put them in a file. To do so, you can first call the <code>generatePublic</code> method on the DSA key factory already created in the example above:

```
PublicKey pubKey = keyFactory.generatePublic(dsaPubKeySpec);
```

Then you can extract the (encoded) key bytes via the following:

```
byte[] encKey = pubKey.getEncoded();
```

You can now store these bytes in a file, and send it to Alice along with the files containing the data and the signature.

Now, assume Alice has received these files, and she copied the data bytes from the data file to a byte array named data, the signature bytes from the signature file to a byte array named signature, and the encoded public key bytes from the public key file to a byte array named encodedPubKey.

Alice can now execute the following code to verify the signature. The code also illustrates how to use a key factory in order to instantiate a DSA public key from its encoding (initVerify requires a PublicKey).

```
X509EncodedKeySpec pubKeySpec = new X509EncodedKeySpec(encodedPubKey);
KeyFactory keyFactory = KeyFactory.getInstance("DSA");
PublicKey pubKey = keyFactory.generatePublic(pubKeySpec);
Signature sig = Signature.getInstance("SHA256withDSA");
sig.initVerify(pubKey);
```



```
sig.update(data);
sig.verify(signature);
```



In the above, Alice needed to generate a PublicKey from the encoded key bits, since initVerify requires a PublicKey. Once she has a PublicKey, she could also use the KeyFactorygetKeySpec method to convert it to a DSAPublicKeySpec so that she can access the components, if desired, as in:

```
DSAPublicKeySpec dsaPubKeySpec =
   (DSAPublicKeySpec)keyFactory.getKeySpec(pubKey, DSAPublicKeySpec.class);
```

Now she can access the DSA public key components y, p, q, and g through the corresponding "get" methods on the DSAPublicKeySpec class (getY, getP, getQ, and getG).

Generating Random Numbers

The following code sample illustrates generating random numbers configured with different security strengths using a DRBG implementation of the SecureRandom class:

```
SecureRandom drbq;
byte[] buffer = new byte[32];
// Any DRBG can be provided
drbg = SecureRandom.getInstance("DRBG");
drbg.nextBytes(buffer);
SecureRandomParameters params = drbg.getParameters();
if (params instanceof DrbgParameters.Instantiation) {
   DrbqParameters.Instantiation ins = (DrbqParameters.Instantiation) params;
    if (ins.getCapability().supportsReseeding()) {
       drbq.reseed();
// The following call requests a weak DRBG instance. It is only
// guaranteed to support 112 bits of security strength.
drbg = SecureRandom.getInstance("DRBG",
   DrbgParameters.instantiation(112, NONE, null));
// Both the next two calls will likely fail, because drbg could be
// instantiated with a smaller strength with no prediction resistance
// support.
drbg.nextBytes(buffer,
    DrbgParameters.nextBytes(256, false, "more".getBytes()));
drbg.nextBytes(buffer,
    DrbgParameters.nextBytes(112, true, "more".getBytes()));
// The following call requests a strong DRBG instance, with a
// personalization string. If it successfully returns an instance,
// that instance is guaranteed to support 256 bits of security strength
// with prediction resistance available.
drbg = SecureRandom.getInstance("DRBG", DrbgParameters.instantiation(
```



Determining If Two Keys Are Equal

Example code for determining if two keys are equal.

In many cases you would like to know if two keys are equal; however, the default method <code>java.lang.Object.equals</code> may not give the desired result. The most provider-independent approach is to compare the encoded keys. If this comparison isn't appropriate (for example, when comparing an <code>RSAPrivateKey</code> and an <code>RSAPrivateCrtKey</code>), you should compare each component.

The following code demonstrates this idea:

```
static boolean keysEqual(Key key1, Key key2) {
    if (key1.equals(key2)) {
      return true;
    if (Arrays.equals(key1.getEncoded(), key2.getEncoded())) {
      return true;
 // More code for different types of keys here.
 // For example, the following code can check if
 // an RSAPrivateKey and an RSAPrivateCrtKey are equal:
 // if ((key1 instanceof RSAPrivateKey) &&
 //
       (key2 instanceof RSAPrivateKey)) {
 //
        if ((key1.getModulus().equals(key2.getModulus())) &&
 //
            (key1.getPrivateExponent().equals(
 //
                                         key2.getPrivateExponent()))) {
            return true;
 //
 // }
    return false;
```

Reading Base64-Encoded Certificates

The following example reads a file with Base64-encoded certificates, which are each bounded at the beginning by

```
and at the end by
----END CERTIFICATE----
```



We convert the FileInputStream (which does not support mark and reset) to a ByteArrayInputStream (which supports those methods), so that each call to generateCertificate consumes only one certificate, and the read position of the input stream is positioned to the next certificate in the file:

```
try (FileInputStream fis = new FileInputStream(filename);
    BufferedInputStream bis = new BufferedInputStream(fis)) {
    CertificateFactory cf = CertificateFactory.getInstance("X.509");
    while (bis.available() > 0) {
        Certificate cert = cf.generateCertificate(bis);
        System.out.println(cert.toString());
    }
}
```

Parsing a Certificate Reply

The following example parses a PKCS7-formatted certificate reply stored in a file and extracts all the certificates from it:

```
try (FileInputStream fis = new FileInputStream(filename)) {
   CertificateFactory cf = CertificateFactory.getInstance("X.509");

   Collection<? extends Certificate> c = cf.generateCertificates(fis);
   for (Certificate cert : c) {
        System.out.println(cert);
   }

   // Or use the aggregate operations below for the above for-loop
   // c.stream().forEach(e -> System.out.println(e));
}
```

Using Encryption

This section takes the user through the process of generating a key, creating and initializing a cipher object, encrypting a file, and then decrypting it. Throughout this example, we use the Advanced Encryption Standard (AES).

Generating a Key

To create an AES key, we have to instantiate a KeyGenerator for AES. We do not specify a provider, because we do not care about a particular AES key generation implementation. Since we do not initialize the KeyGenerator, a system-provided source of randomness and a default keysize will be used to create the AES key:

```
KeyGenerator keygen = KeyGenerator.getInstance("AES");
keygen.init(128);
SecretKey aesKey = keygen.generateKey();
```

After the key has been generated, the same KeyGenerator object can be re-used to create further keys.

Creating a Cipher

The next step is to create a Cipher instance. To do this, we use one of the getInstance factory methods of the Cipher class. We must specify the name of the

requested transformation, which includes the following components, separated by slashes (/):

- · the algorithm name
- the mode (optional)
- the padding scheme (optional)

In this example, we create an AES cipher in Cipher Block Chaining mode, with PKCS5-style padding. We do not specify a provider, because we do not care about a particular implementation of the requested transformation.

The standard algorithm name for AES is "AES", the standard name for the Cipher Block Chaining mode is "CBC", and the standard name for PKCS5-style padding is "PKCS5Padding":

```
Cipher aesCipher;
// Create the cipher
aesCipher = Cipher.getInstance("AES/CBC/PKCS5Padding");
```

We use the generated aeskey from above to initialize the Cipher object for encryption:

```
// Initialize the cipher for encryption
aesCipher.init(Cipher.ENCRYPT_MODE, aesKey);

// Our cleartext
byte[] cleartext = "This is just an example".getBytes();

// Encrypt the cleartext
byte[] ciphertext = aesCipher.doFinal(cleartext);

// Retrieve the parameters used during encryption to properly
// initialize the cipher for decryption
AlgorithmParameters params = aesCipher.getParameters();

// Initialize the same cipher for decryption
aesCipher.init(Cipher.DECRYPT_MODE, aesKey, params);

// Decrypt the ciphertext
byte[] cleartext1 = aesCipher.doFinal(ciphertext);
```

cleartext and cleartext1 are identical.

Using Password-Based Encryption

In this example, we prompt the user for a password from which we derive an encryption key.

It would seem logical to collect and store the password in an object of type <code>java.lang.String</code>. However, here's the caveat: Objects of type <code>string</code> are immutable, i.e., there are no methods defined that allow you to change (overwrite) or zero out the contents of a <code>string</code> after usage. This feature makes <code>string</code> objects unsuitable for storing security sensitive information such as user passwords. You should always collect and store security sensitive information in a char array instead. For that reason, the <code>javax.crypto.spec.PBEKeySpec</code> class takes (and returns) a password as a char array.



In order to use Password-Based Encryption (PBE) as defined in PKCS5, we have to specify a *salt* and an *iteration count*. The same salt and iteration count that are used for encryption must be used for decryption. Newer PBE algorithms use an iteration count of at least 1000.

```
PBEKeySpec pbeKeySpec;
PBEParameterSpec pbeParamSpec;
SecretKeyFactory keyFac;
// Salt
byte[] salt = new SecureRandom().nextBytes(salt);
// Iteration count
int count = 1000;
// Create PBE parameter set
pbeParamSpec = new PBEParameterSpec(salt, count);
// Prompt user for encryption password.
// Collect user password as char array, and convert
// it into a SecretKey object, using a PBE key
// factory.
char[] password = System.console.readPassword("Enter encryption password: ");
pbeKeySpec = new PBEKeySpec(password);
keyFac = SecretKeyFactory.getInstance("PBEWithHmacSHA256AndAES_256");
SecretKey pbeKey = keyFac.generateSecret(pbeKeySpec);
// Create PBE Cipher
Cipher pbeCipher = Cipher.getInstance("PBEWithHmacSHA256AndAES_256");
// Initialize PBE Cipher with key and parameters
pbeCipher.init(Cipher.ENCRYPT_MODE, pbeKey, pbeParamSpec);
// Our cleartext
byte[] cleartext = "This is another example".getBytes();
// Encrypt the cleartext
byte[] ciphertext = pbeCipher.doFinal(cleartext);
```

Sample Programs for Diffie-Hellman Key Exchange, AES/GCM, and HMAC-SHA256

The following are sample programs for Diffie-Hellman key exchange, AES/GCM, and HMAC-SHA256.

Topics

Diffie-Hellman Key Exchange between 2 Parties

Diffie-Hellman Key Exchange between 3 Parties

AES/GCM Example

HMAC-SHA256 Example



Diffie-Hellman Key Exchange between 2 Parties

The program runs the Diffie-Hellman key agreement protocol between 2 parties.

```
* Copyright (c) 1997, 2017, Oracle and/or its affiliates. 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 Oracle 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.
 */
import java.io.*;
import java.math.BigInteger;
import java.security.*;
import java.security.spec.*;
import java.security.interfaces.*;
import javax.crypto.*;
import javax.crypto.spec.*;
import javax.crypto.interfaces.*;
import com.sun.crypto.provider.SunJCE;
public class DHKeyAgreement2 {
    private DHKeyAgreement2() {}
    public static void main(String argv[]) throws Exception {
         * Alice creates her own DH key pair with 2048-bit key size
        System.out.println("ALICE: Generate DH keypair ...");
        KeyPairGenerator aliceKpairGen = KeyPairGenerator.getInstance("DH");
        aliceKpairGen.initialize(2048);
        KeyPair aliceKpair = aliceKpairGen.generateKeyPair();
        // Alice creates and initializes her DH KeyAgreement object
        System.out.println("ALICE: Initialization ...");
        KeyAgreement aliceKeyAgree = KeyAgreement.getInstance("DH");
```



```
aliceKeyAgree.init(aliceKpair.getPrivate());
        // Alice encodes her public key, and sends it over to Bob.
       byte[] alicePubKeyEnc = aliceKpair.getPublic().getEncoded();
        ^{\star} Let's turn over to Bob. Bob has received Alice's public key
        * in encoded format.
        \mbox{\scriptsize {\tt *}} He instantiates a DH public key from the encoded key material.
        KeyFactory bobKeyFac = KeyFactory.getInstance("DH");
        X509EncodedKeySpec x509KeySpec = new X509EncodedKeySpec(alicePubKeyEnc);
        PublicKey alicePubKey = bobKeyFac.generatePublic(x509KeySpec);
        * Bob gets the DH parameters associated with Alice's public key.
        * He must use the same parameters when he generates his own key
        * pair.
        */
        DHParameterSpec dhParamFromAlicePubKey =
((DHPublicKey)alicePubKey).getParams();
        // Bob creates his own DH key pair
       System.out.println("BOB: Generate DH keypair ...");
       KeyPairGenerator bobKpairGen = KeyPairGenerator.getInstance("DH");
       bobKpairGen.initialize(dhParamFromAlicePubKey);
       KeyPair bobKpair = bobKpairGen.generateKeyPair();
        // Bob creates and initializes his DH KeyAgreement object
       System.out.println("BOB: Initialization ...");
       KeyAgreement bobKeyAgree = KeyAgreement.getInstance("DH");
       bobKeyAgree.init(bobKpair.getPrivate());
        // Bob encodes his public key, and sends it over to Alice.
       byte[] bobPubKeyEnc = bobKpair.getPublic().getEncoded();
        * Alice uses Bob's public key for the first (and only) phase
        * of her version of the DH
        * protocol.
        * Before she can do so, she has to instantiate a DH public key
        * from Bob's encoded key material.
        */
       KeyFactory aliceKeyFac = KeyFactory.getInstance("DH");
       x509KeySpec = new X509EncodedKeySpec(bobPubKeyEnc);
       PublicKey bobPubKey = aliceKeyFac.generatePublic(x509KeySpec);
       System.out.println("ALICE: Execute PHASE1 ...");
       aliceKeyAgree.doPhase(bobPubKey, true);
        * Bob uses Alice's public key for the first (and only) phase
        * of his version of the DH
         * protocol.
        System.out.println("BOB: Execute PHASE1 ...");
        bobKeyAgree.doPhase(alicePubKey, true);
        * At this stage, both Alice and Bob have completed the DH key
        * agreement protocol.
```

```
* Both generate the (same) shared secret.
        * /
       try {
           byte[] aliceSharedSecret = aliceKeyAgree.generateSecret();
           int aliceLen = aliceSharedSecret.length;
           byte[] bobSharedSecret = new byte[aliceLen];
           int bobLen;
        } catch (ShortBufferException e) {
           System.out.println(e.getMessage());
                // provide output buffer of required size
       bobLen = bobKeyAgree.generateSecret(bobSharedSecret, 0);
       System.out.println("Alice secret: " +
                toHexString(aliceSharedSecret));
       System.out.println("Bob secret: " +
                toHexString(bobSharedSecret));
       if (!java.util.Arrays.equals(aliceSharedSecret, bobSharedSecret))
           throw new Exception("Shared secrets differ");
       System.out.println("Shared secrets are the same");
        * Now let's create a SecretKey object using the shared secret
        * and use it for encryption. First, we generate SecretKeys for the
        * "AES" algorithm (based on the raw shared secret data) and
        * Then we use AES in CBC mode, which requires an initialization
         * vector (IV) parameter. Note that you have to use the same IV
         * for encryption and decryption: If you use a different IV for
         * decryption than you used for encryption, decryption will fail.
        \mbox{\scriptsize \star} If you do not specify an IV when you initialize the Cipher
         ^{\star} object for encryption, the underlying implementation will generate
        * a random one, which you have to retrieve using the
        * javax.crypto.Cipher.getParameters() method, which returns an
        ^{\star} instance of java.security.AlgorithmParameters. You need to transfer
        * the contents of that object (e.g., in encoded format, obtained via
        * the AlgorithmParameters.getEncoded() method) to the party who will
        * do the decryption. When initializing the Cipher for decryption,
        * the (reinstantiated) AlgorithmParameters object must be explicitly
         * passed to the Cipher.init() method.
       System.out.println("Use shared secret as SecretKey object ...");
       SecretKeySpec bobAesKey = new SecretKeySpec(bobSharedSecret, 0, 16, "AES");
       SecretKeySpec aliceAesKey = new SecretKeySpec(aliceSharedSecret, 0, 16,
"AES");
        * Bob encrypts, using AES in CBC mode
       Cipher bobCipher = Cipher.getInstance("AES/CBC/PKCS5Padding");
       bobCipher.init(Cipher.ENCRYPT_MODE, bobAesKey);
       byte[] cleartext = "This is just an example".getBytes();
       byte[] ciphertext = bobCipher.doFinal(cleartext);
       // Retrieve the parameter that was used, and transfer it to Alice in
       // encoded format
       byte[] encodedParams = bobCipher.getParameters().getEncoded();
        * Alice decrypts, using AES in CBC mode
       // Instantiate AlgorithmParameters object from parameter encoding
```

```
// obtained from Bob
    AlgorithmParameters aesParams = AlgorithmParameters.getInstance("AES");
    aesParams.init(encodedParams);
    Cipher aliceCipher = Cipher.getInstance("AES/CBC/PKCS5Padding");
    aliceCipher.init(Cipher.DECRYPT_MODE, aliceAesKey, aesParams);
   byte[] recovered = aliceCipher.doFinal(ciphertext);
    if (!java.util.Arrays.equals(cleartext, recovered))
        throw new Exception("AES in CBC mode recovered text is " +
                "different from cleartext");
    System.out.println("AES in CBC mode recovered text is "
            "same as cleartext");
 * Converts a byte to hex digit and writes to the supplied buffer
private static void byte2hex(byte b, StringBuffer buf)
    char[] hexChars = { '0', '1', '2', '3', '4', '5', '6', '7', '8',
            '9', 'A', 'B', 'C', 'D', 'E', 'F' };
    int high = ((b \& 0xf0) >> 4);
    int low = (b \& 0x0f);
   buf.append(hexChars[high]);
   buf.append(hexChars[low]);
 * Converts a byte array to hex string
private static String toHexString(byte[] block) {
    StringBuffer buf = new StringBuffer();
    int len = block.length;
    for (int i = 0; i < len; i++) {
        byte2hex(block[i], buf);
        if (i < len-1) {
            buf.append(":");
   return buf.toString();
```

Diffie-Hellman Key Exchange between 3 Parties

The program runs the Diffie-Hellman key agreement protocol between 3 parties.

```
/*

* Copyright (c) 1997, 2017, Oracle and/or its affiliates. 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 Oracle 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.
import java.security.*;
import java.security.spec.*;
import javax.crypto.*;
import javax.crypto.spec.*;
import javax.crypto.interfaces.*;
    * This program executes the Diffie-Hellman key agreement protocol between
    * 3 parties: Alice, Bob, and Carol using a shared 2048-bit DH parameter.
    public class DHKeyAgreement3 {
       private DHKeyAgreement3() {}
       public static void main(String argv[]) throws Exception {
        // Alice creates her own DH key pair with 2048-bit key size
            System.out.println("ALICE: Generate DH keypair ...");
            KeyPairGenerator aliceKpairGen = KeyPairGenerator.getInstance("DH");
            aliceKpairGen.initialize(2048);
            KeyPair aliceKpair = aliceKpairGen.generateKeyPair();
        // This DH parameters can also be constructed by creating a
        // DHParameterSpec object using agreed-upon values
            DHParameterSpec dhParamShared =
((DHPublicKey)aliceKpair.getPublic()).getParams();
        // Bob creates his own DH key pair using the same params
            System.out.println("BOB: Generate DH keypair ...");
            KeyPairGenerator bobKpairGen = KeyPairGenerator.getInstance("DH");
           bobKpairGen.initialize(dhParamShared);
            KeyPair bobKpair = bobKpairGen.generateKeyPair();
        // Carol creates her own DH key pair using the same params
            System.out.println("CAROL: Generate DH keypair ...");
            KeyPairGenerator carolKpairGen = KeyPairGenerator.getInstance("DH");
            carolKpairGen.initialize(dhParamShared);
            KeyPair carolKpair = carolKpairGen.generateKeyPair();
        // Alice initialize
            System.out.println("ALICE: Initialize ...");
            KeyAgreement aliceKeyAgree = KeyAgreement.getInstance("DH");
            aliceKeyAgree.init(aliceKpair.getPrivate());
        // Bob initialize
            System.out.println("BOB: Initialize ...");
            KeyAgreement bobKeyAgree = KeyAgreement.getInstance("DH");
            bobKeyAgree.init(bobKpair.getPrivate());
        // Carol initialize
            System.out.println("CAROL: Initialize ...");
            KeyAgreement carolKeyAgree = KeyAgreement.getInstance("DH");
            carolKeyAgree.init(carolKpair.getPrivate());
        // Alice uses Carol's public key
            Key ac = aliceKeyAgree.doPhase(carolKpair.getPublic(), false);
```

```
// Bob uses Alice's public key
              Key ba = bobKeyAgree.doPhase(aliceKpair.getPublic(), false);
      // Carol uses Bob's public key
              Key cb = carolKeyAgree.doPhase(bobKpair.getPublic(), false);
      // Alice uses Carol's result from above
              aliceKeyAgree.doPhase(cb, true);
      // Bob uses Alice's result from above
              bobKeyAgree.doPhase(ac, true);
      // Carol uses Bob's result from above
               carolKeyAgree.doPhase(ba, true);
      // Alice, Bob and Carol compute their secrets
              byte[] aliceSharedSecret = aliceKeyAgree.generateSecret();
               System.out.println("Alice secret: " + toHexString(aliceSharedSecret));
              byte[] bobSharedSecret = bobKeyAgree.generateSecret();
              System.out.println("Bob secret: " + toHexString(bobSharedSecret));
              byte[] carolSharedSecret = carolKeyAgree.generateSecret();
              System.out.println("Carol secret: " + toHexString(carolSharedSecret));
      // Compare Alice and Bob
               if (!java.util.Arrays.equals(aliceSharedSecret, bobSharedSecret))
                       throw new Exception("Alice and Bob differ");
               System.out.println("Alice and Bob are the same");
      // Compare Bob and Carol
               if (!java.util.Arrays.equals(bobSharedSecret, carolSharedSecret))
                       throw new Exception("Bob and Carol differ");
              System.out.println("Bob and Carol are the same");
* Converts a byte to hex digit and writes to the supplied buffer
      private static void byte2hex(byte b, StringBuffer buf) {
              char[] \ hexChars = \{ \ '0', \ '1', \ '2', \ '3', \ '4', \ '5', \ '6', \ '7', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8', \ '8',
                                                          '9', 'A', 'B', 'C', 'D', 'E', 'F' };
               int high = ((b \& 0xf0) >> 4);
               int low = (b \& 0x0f);
              buf.append(hexChars[high]);
              buf.append(hexChars[low]);
   Converts a byte array to hex string
      private static String toHexString(byte[] block) {
              StringBuffer buf = new StringBuffer();
              int len = block.length;
               for (int i = 0; i < len; i++) {
                       byte2hex(block[i], buf);
                       if (i < len-1) {
                                buf.append(":");
              return buf.toString();
```

AES/GCM Example

The following is a sample program to demonstrate AES/GCM usage to encrypt/decrypt data.

```
/*
 * Copyright (c) 2017, Oracle and/or its affiliates. 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 Oracle 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.
import java.security.AlgorithmParameters;
import java.util.Arrays;
import javax.crypto.*;
public class AESGCMTest {
    public static void main(String[] args) throws Exception {
        // Slightly longer than 1 AES block (128 bits) to show PADDING
        // is "handled" by GCM.
        byte[] data = {
            0x00, 0x01, 0x02, 0x03, 0x04, 0x05, 0x06, 0x07,
            0x08, 0x09, 0x0a, 0x0b, 0x0c, 0x0d, 0x0e, 0x0f,
            0x10};
        // Create a 128-bit AES key.
        KeyGenerator kg = KeyGenerator.getInstance("AES");
        kg.init(128);
        SecretKey key = kg.generateKey();
        // Obtain a AES/GCM cipher to do the enciphering. Must obtain
        // and use the Parameters for successful decryption.
        Cipher encCipher = Cipher.getInstance("AES/GCM/NOPADDING");
        encCipher.init(Cipher.ENCRYPT_MODE, key);
        byte[] enc = encCipher.doFinal(data);
        AlgorithmParameters ap = encCipher.getParameters();
        // Obtain a similar cipher, and use the parameters.
        Cipher decCipher = Cipher.getInstance("AES/GCM/NOPADDING");
        decCipher.init(Cipher.DECRYPT_MODE, key, ap);
        byte[] dec = decCipher.doFinal(enc);
        if (Arrays.compare(data, dec) != 0) {
```

```
throw new Exception("Original data != decrypted data");
}
}
```

HMAC-SHA256 Example

The following is a sample program that demonstrates how to generate a secret-key object for HMAC-SHA256, and initialize a HMAC-SHA256 object with it.

Example 2-14 Generate a Secret-key Object for HMAC-SHA256

```
* Copyright (c) 1997, 2017, Oracle and/or its affiliates. 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 Oracle 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.
import java.security.*;
import javax.crypto.*;
    /**
     * This program demonstrates how to generate a secret-key object for
     * HMACSHA256, and initialize an HMACSHA256 object with it.
    public class initMac {
       public static void main(String[] args) throws Exception {
            // Generate secret key for HmacSHA256
           KeyGenerator kg = KeyGenerator.getInstance("HmacSHA256");
            SecretKey sk = kg.generateKey();
            // Get instance of Mac object implementing HmacSHA256, and
            // initialize it with the above secret key
```



```
Mac mac = Mac.getInstance("HmacSHA256");
    mac.init(sk);
    byte[] result = mac.doFinal("Hi There".getBytes());
}
```



3

How to Implement a Provider in the Java Cryptography Architecture

This document describes what you need to do in order to integrate your provider into Java SE so that algorithms and other services can be found when Java Security API clients request them.

Who Should Read This Document

Programmers that only need to use the Java Security APIs (see Core Classes and Interfaces in Java Cryptography Architecture (JCA) Reference Guide) to access existing cryptography algorithms and other services do *not* need to read this document.

This document is intended for experienced programmers wishing to create their own provider packages supplying cryptographic service implementations. It documents what you need to do in order to integrate your provider into Java so that your algorithms and other services can be found when Java Security API clients request them.

Notes on Terminology

Throughout this document, the terms *JCA* by itself refers to the JCA framework. Whenever this document notes a specific JCA provider, it will be referred to explicitly by the provider name.

- Prior to JDK 1.4, the JCE was an unbundled product, and as such, the JCA and JCE were regularly referred to as separate, distinct components. As JCE is now bundled in JDK, the distinction is becoming less apparent. Since the JCE uses the same architecture as the JCA, the JCE should be more properly thought of as a subset of the JCA.
- The JCA within the JDK includes two software components:
 - the framework that defines and supports cryptographic services for which providers supply implementations. This framework includes packages such as java.security, javax.crypto, javax.crypto.spec, and javax.crypto.interfaces.
 - the actual providers such as *Sun*, *SunRsaSign*, *SunJCE*, which contain the actual cryptographic implementations.

Introduction to Implementing Providers

The Java platform defines a set of APIs spanning major security areas, including cryptography, public key infrastructure, authentication, secure communication, and access control. These APIs allow developers to easily integrate security into their application code. They were designed around the following principles:

- Implementation independence: Applications do not need to implement security themselves. Rather, they can request security services from the Java platform.
 Security services are implemented in providers (see below), which are plugged into the Java platform via a standard interface. An application may rely on multiple independent providers for security functionality.
- **Implementation interoperability**: Providers are interoperable across applications. Specifically, an application is not bound to a specific provider, and a provider is not bound to a specific application.
- Algorithm extensibility: The Java platform includes a number of built-in providers
 that implement a basic set of security services that are widely used today.
 However, some applications may rely on emerging standards not yet
 implemented, or on proprietary services. The Java platform supports the
 installation of custom providers that implement such services.

A Cryptographic Service Provider (provider) refers to a package (or a set of packages) that supply a concrete implementation of a subset of the cryptography aspects of the JDK Security API.

The <code>java.security.Provider</code> class encapsulates the notion of a security provider in the Java platform. It specifies the provider's name and lists the security services it implements. Multiple providers may be configured at the same time, and are listed in order of preference. When a security service is requested, the highest priority provider that implements that service is selected. See Security Providers, which illustrates how a provider selects a requested security service.

Engine Classes and Corresponding Service Provider Interface Classes

An engine class defines a cryptographic service in an abstract fashion (without a concrete implementation). A cryptographic service is always associated with a particular algorithm or type.

A *cryptographic service* either provides cryptographic operations (like those for digital signatures or message digests, ciphers or key agreement protocols); generates or supplies the cryptographic material (keys or parameters) required for cryptographic operations; or generates data objects (keystores or certificates) that encapsulate cryptographic keys (which can be used in a cryptographic operation) in a secure fashion.

For example, here are four engine classes:

- signature class provides access to the functionality of a digital signature algorithm.
- A DSA KeyFactory class supplies a DSA private or public key (from its encoding or transparent specification) in a format usable by the initSign or initVerify methods, respectively, of a DSA Signature object.
- Cipher class provides access to the functionality of an encryption algorithm (such as AES)
- KeyAgreement class provides access to the functionality of a key agreement protocol (such as Diffie-Hellman)

The Java Cryptography Architecture encompasses the classes comprising the Security package that relate to cryptography, including the engine classes. Users of



the API request and utilize instances of the engine classes to carry out corresponding operations. The JDK defines the following engine classes:

- MessageDigest used to calculate the message digest (hash) of specified data.
- Signature used to sign data and verify digital signatures.
- KeyPairGenerator used to generate a pair of public and private keys suitable for a specified algorithm.
- KeyFactory used to convert opaque cryptographic keys of type Key into key specifications (transparent representations of the underlying key material), and vice versa.
- KeyStore used to create and manage a keystore. A keystore is a database of keys. Private keys in a keystore have a certificate chain associated with them, which authenticates the corresponding public key. A keystore also contains certificates from trusted entities.
- CertificateFactory used to create public key certificates and Certificate Revocation Lists (CRLs).
- AlgorithmParameters used to manage the parameters for a particular algorithm, including parameter encoding and decoding.
- AlgorithmParameterGenerator used to generate a set of parameters suitable for a specified algorithm.
- SecureRandom used to generate random or pseudo-random numbers.
- Cipher used to encrypt or decrypt some specified data.
- KeyAgreement used to execute a key agreement (key exchange) protocol between 2 or more parties.
- KeyGenerator used to generate a secret (symmetric) key suitable for a specified algorithm.
- Mac: used to compute the message authentication code of some specified data.
- SecretKeyFactory used to convert opaque cryptographic keys of type SecretKey
 into key specifications (transparent representations of the underlying key material),
 and vice versa.
- CertPathBuilder used to create public key certificates and Certificate Revocation Lists (CRLs).
- CertPathValidator used to validate certificate chains.
- CertStore used to retrieve Certificates and CRLs from a repository.
- ExemptionMechanism used to provide the functionality of an exemption mechanism such as key recovery, key weakening, key escrow, or any other (custom) exemption mechanism. Applications or applets that use an exemption mechanism may be granted stronger encryption capabilities than those which don't. However, please note that cryptographic restrictions are no longer required for most countries, and thus exemption mechanisms may only be useful in those few countries whose governments mandate restrictions.



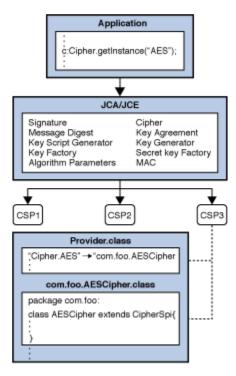
Note:

A *generator* creates objects with brand-new contents, whereas a *factory* creates objects from existing material (for example, an encoding).

An *engine* class provides the interface to the functionality of a specific type of cryptographic service (independent of a particular cryptographic algorithm). It defines *Application Programming Interface* (API) methods that allow applications to access the specific type of cryptographic service it provides. The actual implementations (from one or more providers) are those for specific algorithms. For example, the Signature engine class provides access to the functionality of a digital signature algorithm. The actual implementation supplied in a <code>signatureSpi</code> subclass (see next paragraph) would be that for a specific kind of signature algorithm, such as SHA256withDSA or SHA512withRSA.

The application interfaces supplied by an engine class are implemented in terms of a **Service Provider Interface (SPI)**. That is, for each engine class, there is a corresponding abstract SPI class, which defines the Service Provider Interface methods that cryptographic service providers must implement.

Figure 3-1 Engine Classes



An instance of an engine class, the "API object", encapsulates (as a private field) an instance of the corresponding SPI class, the "SPI object". All API methods of an API object are declared "final", and their implementations invoke the corresponding SPI methods of the encapsulated SPI object. An instance of an engine class (and of its corresponding SPI class) is created by a call to the <code>getInstance</code> factory method of the engine class.



The name of each SPI class is the same as that of the corresponding engine class, followed by "Spi". For example, the SPI class corresponding to the Signature engine class is the SignatureSpi class.

Each SPI class is abstract. To supply the implementation of a particular type of service and for a specific algorithm, a provider must subclass the corresponding SPI class and provide implementations for all the abstract methods.

Another example of an engine class is the MessageDigest class, which provides access to a message digest algorithm. Its implementations, in MessageDigestSpi subclasses, may be those of various message digest algorithms such as SHA256 or SHA384.

As a final example, the <code>KeyFactory</code> engine class supports the conversion from opaque keys to transparent key specifications, and vice versa. See <code>Key Specification</code> Interfaces and Classes Required by <code>Key Factories</code>. The actual implementation supplied in a <code>KeyFactorySpi</code> subclass would be that for a specific type of keys, e.g., DSA public and private keys.

Steps to Implement and Integrate a Provider

Follow these steps to implement a provider and integrate it into the JCA framework:

- Step 1: Write your Service Implementation Code
- Step 2: Give your Provider a Name
- Step 3: Write Your Master Class, a Subclass of Provider
- Step 4: Create a Module Declaration for Your Provider
- Step 5: Compile Your Code
- Step 6: Place Your Provider in a JAR File
- Step 7: Sign Your JAR File, If Necessary
- Step 8: Prepare for Testing
- Step 9: Write and Compile Your Test Programs
- Step 10: Run Your Test Programs
- Step 11: Apply for U.S. Government Export Approval If Required
- Step 12: Document Your Provider and Its Supported Services
- Step 13: Make Your Class Files and Documentation Available to Clients

Step 1: Write your Service Implementation Code

The first thing you need to do is to write the code that provides algorithm-specific implementations of the cryptographic services you want to support. Your provider may supply implementations of cryptographic services already available in one or more of the security components of the JDK.

For cryptographic services not defined in JCA (for example, signatures and message digests), see Engine Classes and Algorithms.

For each cryptographic service you wish to implement, create a subclass of the appropriate SPI class. JCA defines the following engine classes:

SignatureSpi



- MessageDigestSpi
- KeyPairGeneratorSpi
- SecureRandomSpi
- AlgorithmParameterGeneratorSpi
- AlgorithmParametersSpi
- KeyFactorySpi
- CertificateFactorySpi
- KeyStoreSpi
- CipherSpi
- KeyAgreementSpi
- KeyGeneratorSpi
- MacSpi
- SecretKeyFactorySpi
- ExemptionMechanismSpi

To know more about the JCA and other cryptographic classes, see Engine Classes and Corresponding Service Provider Interface Classes.

In the subclass, you need to:

- Supply implementations for the abstract methods, whose names usually begin with engine. See Further Implementation Details and Requirements.
- 2. Depending on how you write your provider and register its algorithms (using either String objects or the Provider. Service class), the provider either:
 - Ensure that there is a public constructor without any arguments. Here's why: When one of your services is requested, Java Security looks up the subclass implementing that service, as specified by a property in your "master class" (see Step 3: Write Your Master Class, a Subclass of Provider). Java Security then creates the class object associated with your subclass, and creates an instance of your subclass by calling the newInstance method on that class object. newInstance requires your subclass to have a public constructor without any parameters. (A default constructor without arguments will automatically be generated if your subclass doesn't have any constructors. But if your subclass defines any constructors, you must explicitly define a public constructor without arguments.)
 - Override the newInstance() method in the registered Provider.Service.

 This is the preferred mechanism in JDK 9 and later.

Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations

When instantiating a provider's implementation (class) of a Cipher, KeyAgreement, KeyGenerator, MAC, or SecretKey factory, the framework will determine the provider's codebase (JAR file) and verify its signature. In this way, JCA authenticates the provider and ensures that only providers signed by a trusted entity can be plugged into



the JCA. Thus, one requirement for encryption providers is that they must be signed, as described in later steps.

In order for provider classes to become unusable if instantiated by an application directly, bypassing JCA, providers should implement the following:

- All SPI implementation classes in a provider package should be declared final (so that they cannot be subclassed), and their (SPI) implementation methods should be declared protected.
- All crypto-related helper classes in a provider package should have packageprivate scope, so that they cannot be accessed from outside the provider package.

For providers that may be exported outside the U.S., <code>cipherSpi</code> implementations must include an implementation of the <code>engineGetKeySize</code> method which, given a <code>Key</code>, returns the key size. If there are restrictions on available cryptographic strength specified in jurisdiction policy files, each <code>cipher</code> initialization method calls <code>engineGetKeySize</code> and then compares the result with the maximum allowable key size for the particular location and circumstances of the applet or application being run. If the key size is too large, the initialization method throws an exception.

Additional *optional* features that providers may implement are:

- Optional: The engineWrap and engineUnwrap methods of CipherSpi. Wrapping a key
 enables secure transfer of the key from one place to another. Information about
 wrapping and unwrapping keys is provided in the wrap.
- Optional: One or more exemption mechanisms. An exemption mechanism is something such as key recovery, key escrow, or key weakening which, if implemented and enforced, may enable reduced cryptographic restrictions for an application (or applet) that uses it. To know more about the requirements for apps that utilize exemption mechanisms, see How to Make Applications Exempt from Cryptographic Restrictions.

Step 2: Give your Provider a Name

Decide on a unique name for your provider. This is the name to be used by client applications to refer to your provider, and it must not conflict with any other provider names.

Step 3: Write Your Master Class, a Subclass of Provider

Create a subclass of the <code>java.security.Provider</code> class. This is essentially a lookup table that advertises the algorithms that your provider implements.

You can use the following coding styles to subclass the Provider class:

- Create a provider that registers its services with String objects to store algorithm names and their associated implementation class name. These are stored in the Hashtable<Object,Object> superclass of java.security.Provider.
- Create a provider that uses the Provider.Service class, which uses a different method to store algorithm names and create new objects. The Provider.Service class enables you customize how the JCE framework requests services from your provider, such as how the framework creates new instances of your provider's services. This coding style is recommended, especially when using modules.



A provider can use either style, or even use both styles at the same time. Regardless of which style you choose, your subclass should be final.

Step 3.1: Create a Provider That Uses String Objects to Register Its Services

The following is an example of a provider that uses String objects to store implemented algorithm names:

To create a provider with this coding style, do the following:

Call super, specifying the provider name (see Step 2: Give your Provider a Name) version number, and a string of information about the provider and algorithms it supports.

• Set the values of various properties that are required for the Java Security API to look up the cryptographic services implemented by the provider.

For each service implemented by the provider, there must be a property whose name is the type of service followed by a period and the name of the algorithm to which the service applies. The property value must specify the fully qualified name of the class implementing the service.

For example, this following statement sets a property named <code>cipher.MyCypher</code> whose value is <code>com.my.crypto.provider.MyCipher</code>, a class that extends <code>CipherSPI</code>:

```
put("Cipher.MyCipher", "com.my.crypto.provider.MyCipher");
```

The following list shows the various types of JCA services, where the actual algorithm name is substituted for <code>algName</code>:

- Signature.algName
- MessageDigest.algName
- KeyPairGenerator.algName
- SecureRandom.algName
- AlgorithmParameterGenerator.algName
- AlgorithmParameters.algName
- KeyFactory.algName
- CertificateFactory.algName
- KeyStore.algName



- Cipher.algName: algName may actually represent a transformation, and may be composed of an algorithm name, a particular mode, and a padding scheme.
 See Java Security Standard Algorithm Names Specification
- KeyAgreement.algName
- KeyGenerator.algName
- Mac.alqName
- SecretKeyFactory.algName
- ExemptionMechanism.algName: algName refers to the name of the exemption mechanism, which can be one of the following: KeyRecovery, KeyEscrow, or KeyWeakening. Case does not matter.

In all of these except ExemptionMechanism and Cipher, algName is the "standard" name of the algorithm, certificate type, or keystore type. See Java Security Standard Algorithm Names Specification for the standard names that should be used.

The value of each property must be the fully qualified name of the class implementing the specified algorithm, certificate type, or keystore type. That is, it must be the package name followed by the class name, where the two are separated by a period.

As an example, the default provider named *SUN* implements the Digital Signature Algorithm (whose standard name is <code>SHA256withDSA</code>) in a class named <code>DSA</code> in the <code>sun.security.provider</code> package. Its subclass of <code>Provider</code> (which is the Sun class in the <code>sun.security.provider</code> package) sets the <code>Signature.SHA256withDSA</code> property to have the <code>value sun.security.provider.DSA</code> via the following:

```
put("Signature.SHA256withDSA", "sun.security.provider.DSA")
```

The list below shows more properties that can be defined for the various types of services, where the actual algorithm name is substituted for *algName*, certificate type for *certType*, keystore type for *storeType*, and attribute name for *attrName*:

- Signature.algName [one or more spaces] attrName
- MessageDigest.algName [one or more spaces] attrName
- KeyPairGenerator.algName [one or more spaces] attrName
- SecureRandom.algName [one or more spaces] attrName
- KeyFactory.algName [one or more spaces] attrName
- CertificateFactory.certType [one or more spaces] attrName
- KeyStore.storeType [one or more spaces] attrName
- AlgorithmParameterGenerator.algName [one or more spaces] attrName
- AlgorithmParameters.algName [one or more spaces] attrName
- Cipher.algName [one or more spaces] attrName
- KeyAgreement.algName [one or more spaces] attrName
- KeyGenerator.algName [one or more spaces] attrName
- Mac.algName [one or more spaces] attrName
- SecretKeyFactory.algName [one or more spaces] attrName
- ExemptionMechanism.algName [one or more spaces] attrName



In each of these, attrName is the "standard" name of the algorithm, certificate type, keystore type, or attribute. (See Java Security Standard Algorithm Names Specification for the standard names that should be used.)

For a property in the above format, the value of the property must be the value for the corresponding attribute. (See Java Security Standard Algorithm Names Specification for the definition of each standard attribute.)

For further master class property setting examples, see the JDK 9 source code for the sun.security.provider.Sun and com.sun.crypto.provider.SunJCE classes. They show how the Sun and SunJCE providers set properties.

As an example, the default provider named SUN implements the SHA256withDSA Digital Signature Algorithm in software. The class sun.security.provider.Sun calls the method SunEntries.putEntries, which sets the properties for the SUN provider, including setting the property Signature.SHA256withDSA ImplementedIn to have the value Software:

```
put("Signature.SHA256withDSA ImplementedIn", "Software");
```

Note:

For examples of this coding style, see the source code for sun.security.provider.Sun and sun.security.provider.SunEntries Classes.

Step 3.2: Create a Provider That Uses Provider. Service

The following is an example of a provider that uses a Provider . Service class:

```
package p;
public final class MyProvider extends Provider {
    public MyProvider() {
        super("MyProvider", "1.0",
            "Some info about my provider and which algorithms it supports");
        putService(new ProviderService(this, "Cipher", "MyCipher", "p.MyCipher"));
    private static final class ProviderService extends Provider.Service {
        ProviderService(Provider p, String type, String algo, String cn) {
            super(p, type, algo, cn, null, null);
        @Override
        public Object newInstance(Object ctrParamObj)
            throws NoSuchAlgorithmException {
            String type = getType();
            String algo = getAlgorithm();
                if (type.equals("Cipher")) {
                    if (algo.equals("MyCipher")) {
                        return new MyCipher();
```



To create a provider with this coding style, do the following:

 For each algorithm your provider supports, call putService with an instance of Provider.Service; the arguments of the Provider.Service constructor represent a supported algorithm.

The following statement adds a service named MyCipher of type Cipher; the name of the class implementing this service is p.MyCipher. The argument of putService is a subclass of Provider. Service:

```
putService(new ProviderService(this, "Cipher", "MyCipher",
"p.MyCipher"));
```

This example uses a subclass of Provider.Service named ProviderService (rather than Provider.Service itself) as it customizes how the JCE framework instantiates services. If you don't need to customize the behavior of Provider.Service, then you can call the Provider.Service constructor directly:

Note that this example is essentially the same as the example described in Step 3.1: Create a Provider That Uses String Objects to Register Its Services.

 Override any method in Provider. Service, such as newInstance, to customize how the JCE framework handles the services in your provider.

The example at the beginning of this section overrides the method Provider.Service.newInstance. The method returns an instance of MyCipher only if the requested service is MyCipher. If not, it throws a NoSuchAlgorithmException and a ProviderException.

For more information about other methods you can override, see The Provider.Service Class.



For examples of this coding style, see the JDK 9 source code contained in the sun.security.mscapi package.



Step 3.3: Specify Additional Information for Cipher Implementations

As mentioned above, in the case of a cipher property, algName may actually represent a transformation. A transformation is a string that describes the operation (or set of operations) to be performed by a cipher object on some given input. A transformation always includes the name of a cryptographic algorithm (e.g., AES), and may be followed by a mode and a padding scheme.

A transformation is of the form:

- algorithm/mode/padding, or
- algorithm

(In the latter case, provider-specific default values for the mode and padding scheme are used). For example, the following is a valid transformation:

```
Cipher c = Cipher.getInstance("AES/CBC/PKCS5Padding");
```

When requesting a block cipher in stream cipher mode (for example; AES in CFB or OFB mode), a client may optionally specify the number of bits to be processed at a time, by appending this number to the mode name as shown in the following sample transformations:

```
Cipher c1 = Cipher.getInstance("AES/CFB8/NoPadding");
Cipher c2 = Cipher.getInstance("AES/OFB32/PKCS5Padding");
```

If a number does not follow a stream cipher mode, a provider-specific default is used. (For example, the *SunJCE* provider uses a default of 128 bits.)

A provider may supply a separate class for each combination of algorithm/mode/padding. Alternatively, a provider may decide to provide more generic classes representing sub-transformations corresponding to algorithm or algorithm/mode or algorithm//padding (note the double slashes); in this case the requested mode and/or padding are set automatically by the getInstance methods of Cipher, which invoke the engineSetMode and engineSetPadding methods of the provider's subclass of CipherSpi.

That is, a cipher property in a provider master class may have one of the formats shown in the table below.

Table 3-1 Cipher Property Format

Ginken Proporty Format	Description
Cipher Property Format	Description
Cipher.algName	A provider's subclass of CipherSpi implements algName with pluggable mode and padding
Cipher.algName/mode	A provider's subclass of CipherSpi implements algName in the specified mode, with pluggable padding
Cipher.algName//padding	A provider's subclass of CipherSpi implements algName with the specified padding, with pluggable mode
Cipher.algName/mode/padding	A provider's subclass of CipherSpi implements algName with the specified mode and padding

(See Java Security Standard Algorithm Names Specification for the standard algorithm names, modes, and padding schemes that should be used.)



For example, a provider may supply a subclass of <code>cipherSpi</code> that implements <code>AES/ECB/PKCS5Padding</code>, one that implements <code>AES/CBC/PKCS5Padding</code>, one that implements <code>AES/CFB/PKCS5Padding</code>, and yet another one that implements <code>AES/OFB/PKCS5Padding</code>. That provider would have the following <code>cipher</code> properties in its master class:

- Cipher.AES/ECB/PKCS5Padding
- Cipher.AES/CBC/PKCS5Padding
- Cipher.AES/CFB/PKCS5Padding
- Cipher.AES/OFB/PKCS5Padding

Another provider may implement a class for each of the above modes (i.e., one class for *ECB*, one for *CBC*, one for *CFB*, and one for *OFB*), one class for *PKCS5Padding*, and a generic *AES* class that subclasses from <code>cipherSpi</code>. That provider would have the following <code>cipher</code> properties in its master class:

- Cipher.AES
- Cipher.AES SupportedModes
 - Example: "ECB|CBC|CFB|OFB"
- Cipher.AES SupportedPaddings
 - Example: "NOPADDING | PKCS5Padding"

The getInstance factory method of the Cipher engine class follows these rules in order to instantiate a provider's implementation of CipherSpi for a transformation of the form "algorithm":

- Check if the provider has registered a subclass of CipherSpi for the specified "algorithm".
 - If the answer is YES, instantiate this class, for whose mode and padding scheme default values (as supplied by the provider) are used.
 - If the answer is NO, throw a NoSuchAlgorithmException exception.
- 2. The getInstance factory method of the Cipher engine class follows these rules in order to instantiate a provider's implementation of CipherSpi for a transformation of the form "algorithm/mode/padding":
 - a. Check if the provider has registered a subclass of CipherSpi for the specified "algorithm/mode/padding" transformation.
 - If the answer is YES, instantiate it.
 - If the answer is NO, go to the next step.
 - **b.** Check if the provider has registered a subclass of CipherSpi for the subtransformation "algorithm/mode".
 - If the answer is YES, instantiate it, and call engineSetPadding(padding) on the new instance.
 - If the answer is NO, go to the next step.
 - **c.** Check if the provider has registered a subclass of <code>cipherSpi</code> for the subtransformation "algorithm//padding" (note the double slashes)
 - If the answer is YES, instantiate it, and call engineSetMode(mode) on the new instance.



- If the answer is NO, go to the next step.
- **d.** Check if the provider has registered a subclass of CipherSpi for the subtransformation "algorithm".
 - If the answer is YES, instantiate it, and call <code>engineSetMode(mode)</code> and <code>engineSetPadding(padding)</code> on the new instance.
 - If the answer is NO, throw a NoSuchAlgorithmException exception.

Step 4: Create a Module Declaration for Your Provider

This step is optional but recommended; it enables you to package your provider in a named module. A modular JDK can then locate your provider in the module path as opposed to the class path. The module system can more thoroughly check for dependencies in modules in the module path. Note that you can use named modules in a non-modular JDK; the module declaration will be ignored. Also, you can still package your providers in unnamed or automatic modules.

Create a module declaration for your provider and save it in a file named module-info.java. This module declaration includes the following:

- The name of your module.
- Any module upon which your provider depends.
- A provides directive if your module provides a service implementation.

The following example module declaration defines a module named <code>com.foo.MyProvider.p.MyProvider</code> is the fully qualified class name of a service implementation. Suppose that, in this example, <code>p.MyProvider</code> uses API in the package <code>javax.security.auth.kerberos</code>, which is in the module <code>java.security.jgss</code>. Thus, the directive <code>requires java.security.jgss</code> appears in the module declaration.

```
module com.foo.MyProvider {
    provides java.security.Provider with p.MyProvider;
    requires java.security.jgss;
}
```

You can package a provider in three different kinds of modules:

- Named or explicit module: A module that appears on the module path and contains module configuration information in the module-info.class file.
 - The JCE framework can use the ServiceLoader class (which simplifies provider configuration) to search for providers in explicit modules without any additional changes to the module. See Step 8.1: Configure the Provider and Step 10: Run Your Test Programs.
- Automatic module: A module that appears on the module path, but does not contain module configuration information in a module-info.class file (essentially a "regular" JAR file).
- Unnamed module: A module that appears on the class path. It may or may not have a module-info.class file; this file is ignored.

It is recommended that you package your providers in named modules as they provide better performance, stronger encapsulation, simpler configuration and greater flexibility.



You have a lot of flexibility when it comes to packaging and configuring your providers. However, this impacts how you start applications that use them. For example, you might have to specify additional <code>--add-exports</code> or <code>--add-modules</code> options. Named modules, in general, require fewer of these additional options. In addition named modules offer more flexibility. You can use them with non-modular JDKs or even as unnamed modules by specifying them in a modular JDK's class path. For more information about modules, see The State of the Module System and JEP 261: Module System.

Step 5: Compile Your Code

After you have created your implementation code (Step 1: Write your Service Implementation Code), given your provider a name (Step 2: Give your Provider a Name), created the master class (Step 3: Write Your Master Class, a Subclass of Provider), and created a module declaration (Step 4: Create a Module Declaration for Your Provider), use the Java compiler to compile your files.

Step 6: Place Your Provider in a JAR File

Add the File java.security.Provider to Use the ServiceLoader Class to Search for Providers

If your provider is packaged in an automatic or unnamed module (you did not create a module declaration as described in Step 4: Create a Module Declaration for Your Provider) and you want the use the <code>java.util.ServiceLoader</code> to search for your providers, then add the file META-INF/services/java.security.Provider to the JAR file and ensure that the file contains the fully qualified class name of your provider implementation.

The security provider loading mechanism uses the ServiceLoader class to search for providers before consulting the class path.

For example, if the fully qualified class name of your provider is p.Provider and all the compiled code of your provider is in the directory classes, then create a file named classes/META-INF/services/java.security.Provider that contains the following line:

p.MyProvider

Run the jar Command to Create a JAR File

The following command creates a JAR file named MyProvider.jar. All the compiled code for the module JAR file is in the directory classes. In addition, the module descriptor, module-info.class, is in the directory classes:

jar --create --file MyProvider.jar --module-version 1.0 -C classes



The module-info.class file and the --module-version option are optional. However, the module-info.class file is required if you want to create a modular JAR file. (A modular JAR file is a regular JAR file that has a module-info.class file in its top-level directory.)



See jar in Java Platform, Standard Edition Tools Reference.

Step 7: Sign Your JAR File, If Necessary

If your provider is supplying encryption algorithms through the Cipher, KeyAgreement, KeyGenerator, Mac, Or SecretKeyFactory classes, you must sign your JAR file so that the JCA can authenticate the code at run time; see Step 1.1: Consider Additional JCA Provider Requirements and Recommendations for Encryption Implementations. If you are not providing an implementation of this type, then you can skip this step.

Step 7.1: Get a Code-Signing Certificate

The next step is to request a code-signing certificate so that you can use it to sign your provider prior to testing. The certificate will be good for both testing and production. It will be valid for 5 years.

Below are the steps you should use to get a code-signing certificate. See keytool in the Java Platform, Standard Edition Tools Reference.

1. Use **keytool** to generate a RSA keypair, using RSA algorithm as an example:

```
keytool -genkeypair -alias <alias> \
    -keyalg RSA -keysize 2048 \
    -dname "cn=<Company Name>, \
    ou=Java Software Code Signing, \
    o=Oracle Corporation" \
    -keystore <keystore file name> \
    -storepass <keystore password>
```

This will generate a DSA keypair (a public key and an associated private key) and store it in an entry in the specified keystore. The public key is stored in a self-signed certificate. The keystore entry can subsequently be accessed using the specified alias.

The option values in angle brackets ("<" and ">") represent the actual values that must be supplied. For example, <alias> must be replaced with whatever alias name you wish to be used to refer to the newly-generated keystore entry in the future, and <keystore file name> must be replaced with the name of the keystore to be used.



Tip:

Do not surround actual values with angle brackets. For example, if you want your alias to be myTestAlias, specify the -alias option as follows:

```
-alias myTestAlias
```

If you specify a keystore that doesn't yet exist, it will be created.



Note:

If command lines you type are not allowed to be as long as the <code>keytool -genkeypair</code> command you want to execute (for example, if you are typing to a Microsoft Windows DOS prompt), you can create and execute a plain-text batch file containing the command. That is, create a new text file that contains nothing but the full <code>keytool -genkeypair</code> command. (Remember to type it all on one line.) Save the file with a .bat extension. Then in your DOS window, type the file name (with its path, if necessary). This will cause the command in the batch file to be executed.

2. Use **keytool** to generate a certificate signing request.

```
keytool -certreq -alias <alias> \
    -file <csr file name> \
    -keystore <keystore file name> \
    -storepass <keystore password>
```

Here, <alias> is the alias for the DSA keypair entry created in the previous step. This command generates a Certificate Signing Request (CSR), using the PKCS#10 format. It stores the CSR in the file whose name is specified in <csr file name>.

- Send the CSR, contact information, and other required documentation to the JCA Code Signing Certification Authority. See JCA Code Signing Certification Authority for contact information.
- 4. After the JCA Code Signing Certification Authority has received your email message, they will send you a request number via email. Once you receive this request number, you should print, fill out and send the Certification Form for CSPs. See Sending Certification Form for CSPs for contact information.
- 5. Use **keytool** to import the certificates received from the CA.

Once you have received the two certificates from the JCA Code Signing Certification Authority, you can use **keytool** to import them into your keystore. First import the CA's certificate as a "trusted certificate":

```
keytool -import -alias <alias for the CA cert> \
    -file <CA cert file name> \
    -keystore <keystore file name> \
    -storepass <keystore password>
```

Then import the code-signing certificate:

```
keytool -import -alias <alias> \
    -file <code-signing cert file name> \
    -keystore <keystore file name> \
    -storepass <keystore password>
```

<alias> is the same alias as that which you created in Step 1 where you generated a DSA keypair. This command replaces the self-signed certificate in the keystore entry specified by <alias> with the one signed by the JCA Code Signing Certification Authority.

Now that you have in your keystore a certificate from an entity trusted by JCA (the JCA Code Signing Certification Authority), you can place your provider code in a JAR file



(Step 6: Place Your Provider in a JAR File) and then use that certificate to sign the JAR file (Step 7.2: Sign Your Provider).

Step 7.2: Sign Your Provider

Sign the JAR file created in Step 6: Place Your Provider in a JAR File with the codesigning certificate obtained in Step 7.1: Get a Code-Signing Certificate. See jarsigner in Java Platform, Standard Edition Tools Reference.

```
jarsigner -keystore <keystore file name> \
    -storepass <keystore password> \
    <JAR file name> <alias>
```

Here, <alias> is the alias into the keystore for the entry containing the code-signing certificate received from the JCA Code Signing Certification Authority (the same alias as that specified in the commands in Step 7.1: Get a Code-Signing Certificate).

You can test verification of the signature via the following:

```
jarsigner -verify <JAR file name>
```

The text "jar verified" will be displayed if the verification was successful.

Note:

- If you bundle a signed JCE provider as part of an RIA (applet or webstart application), for the best user experience, you should apply a second signature to the JCE provider JAR with the same certificate/key that you used to sign the applet or webstart application. See Deployment Configuration File and Properties to know about deploying RIAs, and jarsigner in Java Platform, Standard Edition Tools Reference for applying multiple signatures to a JAR file.
- You cannot package signed providers in JMOD files.
- Providers don't need to be signed.
- You can link a provider in a custom runtime image with the jlink command as long as it doesn't have a Cipher, KeyAgreement, or MAC implementation.

Step 8: Prepare for Testing

The next steps describe how to install and configure your new provider so that it is available via the JCA.

Step 8.1: Configure the Provider

Register your provider so that the JCE framework can find your provider, either with the ServiceLoader class or in the class path or module path.

- Open the java.security file in an editor:
 - Solaris, Linux, or macOS: <java-home>/conf/security/java.security



- Windows: <java-home>\conf\security\java.security
- 2. In the java.security file, find the section where standard providers such as SUN, SunRsaSign, and SunJCE are configured as static providers; it looks like the following:

```
security.provider.1=SUN
security.provider.2=SunRsaSign
security.provider.3=SunEC
security.provider.4=SunJSSE
security.provider.5=SunJCE
security.provider.6=SunJGSS
security.provider.7=SunSASL
security.provider.8=XMLDSig
security.provider.9=SunPCSC
security.provider.10=JdkLDAP
security.provider.11=JdkSASL
security.provider.12=SunMSCAPI
security.provider.13=SunPKCS11
```

Each line in this section has the following form:

```
security.provider.n=provName | className
```

This declares a provider, and specifies its preference order n. The preference order is the order in which providers are searched for requested algorithms when no specific provider is requested. The order is 1-based; 1 is the most preferred, followed by 2, and so on.

provName is the provider's name and className is the fully qualified class name of the provider. You can use either of these two names.

3. Register your provider by adding to the java.security file a line with the form security.provider.n=provName | className.

If you configured your provider so that the ServiceLoader class can search for it (because you packaged the provider in a named module as described in Step 4: Create a Module Declaration for Your Provider or added a java.security.Provider file as described in Add the File java.security.Provider to Use the ServiceLoader Class to Search for Providers), then specify just the provider's name.

If you have not configured your provider so that ServiceLoader class can search for it, which means that the JCE framework will search for it in the class path or module path, then specify the fully qualified class name of your provider.

For example, the highlighted line registers the provider MyProvider (whose fully qualified class name is p.MyProvider and has been configured so that the ServiceLoader class can search for it) as the 14th preferred provider:

```
# ...
security.provider.11=JdkSASL
security.provider.12=SunMSCAPI
security.provider.13=SunPKCS11
security.provider.14=MyProvider
```

If you are not sure if the ServiceLoader mechanism will be used, or if you'll be deploying on a non-modular system, then you can also register the provider again, this time using the full class name:

```
{\tt security.provider.15=p.MyProvider}
```



Alternatively, you can register providers dynamically. To do so, a program (such as your test program, to be written in Step 9: Write and Compile Your Test Programs) call either the addProvider or insertProviderAt method in the security class:

```
ServiceLoader<Provider> sl = ServiceLoader.load(java.security.Provider.class);
for (Provider p : sl) {
    System.out.println(p);
    if (p.getName().equals("MyProvider")) {
        Security.addProvider(p);
    }
}
```

This type of registration is not persistent and can only be done by code which is granted the following permission:

```
java.security.SecurityPermission "insertProvider.cprovider name>"
```

For example, if the provider name is MyJCE, and if the provider's code is in the myjce_provider.jar file in the /localWork directory, then the following is a sample policy file that contains a grant statement that grants that permission:

```
grant codeBase "file:/localWork/myjce_provider.jar" {
    permission java.security.SecurityPermission
         "insertProvider.MyJCE";
};
```

Step 8.2: Set Provider Permissions

Permissions must be granted for when applications are run while a security manager is installed. A security manager may be installed for an application either through code in the application itself or through a command-line argument.

- Your provider may need the following permissions granted to it in the client environment:
 - java.lang.RuntimePermission to get class protection domains. The provider may need to get its own protection domain in the process of doing self-integrity checking.
 - java.security.SecurityPermission to set provider properties.
- To ensure your provider works when a security manager is installed, you need to test such an installation and execution environment. In addition, prior to testing your need to grant appropriate permissions to your provider and to any other providers it uses.

For example, a sample statement granting permissions to a provider whose name is MyJCE and whose code is in <code>myjce_provider.jar</code> appears below. Such a statement could appear in a policy file. In this example, the <code>myjce_provider.jar</code> file is assumed to be in the <code>/localWork</code> directory.

```
grant codeBase "file:/localWork/myjce_provider.jar" {
   permission java.lang.RuntimePermission "getProtectionDomain";
   permission java.security.SecurityPermission
        "putProviderProperty.MyJCE";
};
```



Step 9: Write and Compile Your Test Programs

Write and compile one or more test programs that test your provider's incorporation into the Security API as well as the correctness of its algorithm(s). Create any supporting files needed, such as those for test data to be encrypted.

 The first tests your program should perform are ones to ensure that your provider is found, and that its name, version number, and additional information is as expected.

To do so, you could write code like the following, substituting your provider name for MyPro:

```
import java.security.*;
Provider p = Security.getProvider("MyPro");
System.out.println("MyPro provider name is " + p.getName());
System.out.println("MyPro provider version # is " + p.getVersion());
System.out.println("MyPro provider info is " + p.getInfo());
```

2. You should ensure that your services are found.

For instance, if you implemented the AES encryption algorithm, you could check to ensure it's found when requested by using the following code (again substituting your provider name for "MyPro"):

```
Cipher c = Cipher.getInstance("AES", "MyPro");
System.out.println("My Cipher algorithm name is " + c.getAlgorithm());
```

- 3. Optional: If you don't specify a provider name in the call to <code>getInstance</code>, all registered providers will be searched, in preference order (see Step 8.1: Configure the Provider), until one implementing the algorithm is found.
- 4. Optional: If your provider implements an exemption mechanism, you should write a test applet or application that uses the exemption mechanism. Such an applet/ application also needs to be signed, and needs to have a "permission policy file" bundled with it.

See How to Make Applications Exempt from Cryptographic Restrictions for complete information on creating and testing such an application.

Step 10: Run Your Test Programs

When you run your test applications, the required <code>java</code> command options will vary depending on factors such as whether you packaged your provider as a named, automatic, or unnamed module and if you configured it so that the <code>ServiceLoader</code> class can search for it.

If you packaged your provider as a named module and have configured it so that the ServiceLoader class can search for it (by registering it with its name in the java.security as described in Step 8.1: Configure the Provider), then run your test program with the following command:

```
java --module-path "jars" <other java options>
```

The directory jars contains your provider.



You may require more options depending on your provider code style (see Step 3.1: Create a Provider That Uses String Objects to Register Its Services and Step 3.2: Create a Provider That Uses Provider.Service), if you packaged your provider in a different kind of module, or if you have not configured it for the ServiceLoader class. The following table describes these options.

For the java commands, the name of the provider is MyProvider, its fully qualified class name is p.MyProvider, and it is packaged in the file com.foo.MyProvider.jar, which is in the directory jars.

Table 3-2 Expected Java Runtime Options for Various Provider Implementation Styles

Modu le Type	Provider Code Style	Configur ed for ServiceL oader Class?	Provider Name Used in java.securit y File	java Command
Unna med	String objects or Provider.Se rvice	No	Fully qualified class name	<pre>java -cp "jars/ com.foo.MyProvider.jar" <other java="" options=""></other></pre>
Unna med	String objects or Provider.Se rvice	Yes	Fully qualified class name or provider name	<pre>java -cp "jars/ com.foo.MyProvider.jar" <other java="" options=""></other></pre>
Auto matic	String objects or Provider.Se rvice	No	Fully qualified class name	<pre>javamodule-path "jars/ com.foo.MyProvider.jar"add- modules=com.foo.MyProvider <other java="" options=""></other></pre>
Auto matic	String objects or Provider.Se rvice	Yes	Fully qualified class name or provider name	<pre>javamodule-path "jars/ com.foo.MyProvider.jar" <other java="" options=""></other></pre>
Name d	String objects or Provider.Se rvice	No	Fully qualified class name	javamodule-path "jars"add- modules=com.foo.MyProvideradd- exports=com.foo.MyProvider/ p=java.base <other java="" options=""> You can remove theadd-exports option if you add exports p in the module declaration.</other>
Name d	String objects	Yes	Fully qualified class name	javamodule-path "jars"add- exports=com.foo.MyProvider/ p=java.base <other java="" options=""> You can remove theadd-exports option if you add exports p in the module declaration.</other>
Name d	String objects	Yes	Provider name	javamodule-path "jars"add- exports=com.foo.MyProvider/ p=java.base <other java="" options=""> You can remove theadd-exports option if you add exports p in the module declaration.</other>



Table 3-2 (Cont.) Expected Java Runtime Options for Various Provider Implementation Styles

Modu le Type	Provider Code Style	Configur ed for ServiceL oader Class?	Provider Name Used in java.securit y File	java Command
Name d	Provider.Se rvice	Yes	Fully qualified class name	javamodule-path "jars"add- exports=com.foo.MyProvider/ p=java.base <other java="" options=""> You can remove theadd-exports option if you add exports p in the module declaration.</other>
Name d	Provider.Se rvice	Yes	Provider name	javamodule-path "jars" <other java="" options=""></other>

Once you have determined the proper java options for your test programs, run them. Debug your code and continue testing as needed. If the Java runtime cannot seem to find one of your algorithms, review the previous steps and ensure that they are all completed.

Be sure to include testing of your programs using different installation options (for example, configured to use the ServiceLoader class or to be found in the class path or module path) and execution environments (with or without a security manager running).

- 1. Optional: If you find during testing that your code needs modification, make the changes and recompile Step 5: Compile Your Code.
- 2. Place the updated provider code in a JAR file (Step 6: Place Your Provider in a JAR File).
- 3. Sign the JAR file (Step 7: Sign Your JAR File, If Necessary).
- 4. Re-configure the provider (Step 8.1: Configure the Provider).
- Optional: If needed, fix or add to the permissions (Step 8.2: Set Provider Permissions).
- 6. Run your programs.
- Optional: If required, repeat steps 1 to 6.

Step 11: Apply for U.S. Government Export Approval If Required

All U.S. vendors whose providers may be exported outside the U.S. should apply to the Bureau of Industry and Security in the U.S. Department of Commerce for export approval.

Please consult your export counsel for more information.



Note:

If your provider calls <code>cipher.getInstance()</code> and the returned <code>cipher</code> object needs to perform strong cryptography regardless of what cryptographic strength is allowed by the user's downloaded jurisdiction policy files, you should include a copy of the <code>cryptoPerms</code> permission policy file which you intend to bundle in the JAR file for your provider and which specifies an appropriate permission for the required cryptographic strength. The necessity for this file is just like the requirement that applets and applications "exempt" from cryptographic restrictions must include a <code>cryptoPerms</code> permission policy file in their JAR file. See How to Make Applications Exempt from Cryptographic Restrictions.

Here are two URLs that may be useful:

- US Department of Commerce
- Bureau of Industry and Security

Step 12: Document Your Provider and Its Supported Services

The next step is to write documentation for your clients. At the minimum, you need to specify:

The name programs should use to refer to your provider.

Note:

As of this writing, provider name searches are **case-sensitive**. That is, if your master class specifies your provider name as "CryptoX" but a user requests "CRYPTOx", your provider will not be found. This behavior may change in the future, but for now be sure to warn your clients to use the exact case you specify.

- The types of algorithms and other services implemented by your provider.
- Instructions for installing the provider, similar to those provided in Step 8.1:
 Configure the Provider, except that the information and examples should be specific to your provider.
- The permissions your provider will require if it is not installed as an installed extension and if a security manager is run, as described in Step 8.2: Set Provider Permissions.

In addition, your documentation should specify anything else of interest to clients, such as any default algorithm parameters.

Step 12.1: Indicate Whether Your Implementation is Cloneable for Message Digests and MACs

For each Message Digest and MAC algorithm, indicate whether or not your implementation is cloneable. This is not technically necessary, but it may save clients

some time and coding by telling them whether or not intermediate Message Digests or MACs may be possible through cloning.

Clients who do not know whether or not a MessageDigest or Mac implementation is cloneable can find out by attempting to clone the object and catching the potential exception, as illustrated by the following example:

```
try {
    // try and clone it
    /* compute the MAC for i1 */
    mac.update(i1);
    byte[] i1Mac = mac.clone().doFinal();

    /* compute the MAC for i1 and i2 */
    mac.update(i2);
    byte[] i12Mac = mac.clone().doFinal();

    /* compute the MAC for i1, i2 and i3 */
    mac.update(i3);
    byte[] i123Mac = mac.doFinal();
} catch (CloneNotSupportedException cnse) {
    // have to use an approach not involving cloning
}
```

Where,

mad

Indicates the MAC object they received when they requested one via a call to Mac.getInstance

i1, i2 and i3

Indicates input byte arrays, and they want to calculate separate hashes for:

- 11
- i1 and i2
- i1, i2, and i3

Key Pair Generators

For a key pair generator algorithm, in case the client does not explicitly initialize the key pair generator (via a call to an <code>initialize</code> method), each provider must supply and document a default initialization.

For example, the Diffie-Hellman key pair generator supplied by the *SunJCE* provider uses a default prime modulus size (keysize) of 2048 bits.

Key Factories

A provider should document all the key specifications supported by its (secret-)key factory.

Algorithm Parameter Generators

In case the client does not explicitly initialize the algorithm parameter generator (via a call to an init method in the AlgorithmParameterGenerator engine class), each provider must supply and document a default initialization.



For example, the *SunJCE* provider uses a default prime modulus size (keysize) of 2048 bits for the generation of Diffie-Hellman parameters, the *Sun* provider a default modulus prime size of 2048 bits for the generation of DSA parameters.

Signature Algorithms

If you implement a signature algorithm, you should document the format in which the signature (generated by one of the sign methods) is encoded.

For example, the SHA256withDSA signature algorithm supplied by the "SUN" provider encodes the signature as a standard ASN.1 SEQUENCE of two integers, r and s.

Random Number Generation (SecureRandom) Algorithms

For a random number generation algorithm, provide information regarding how "random" the numbers generated are, and the quality of the seed when the random number generator is self-seeding. Also note what happens when a <code>SecureRandom</code> object (and its encapsulated <code>SecureRandomSpi</code> implementation object) is deserialized: If subsequent calls to the <code>nextBytes</code> method (which invokes the <code>engineNextBytes</code> method of the encapsulated <code>SecureRandomSpi</code> object) of the restored object yield the exact same (random) bytes as the original object would, then let users know that if this behavior is undesirable, they should seed the restored random object by calling its <code>setSeed</code> method.

Certificate Factories

A provider should document what types of certificates (and their version numbers, if relevant), can be created by the factory.

Keystores

A provider should document any relevant information regarding the keystore implementation, such as its underlying data format.

Step 13: Make Your Class Files and Documentation Available to Clients

After writing, configuring, testing, installing and documenting your provider software, make documentation available to your customers.

Further Implementation Details and Requirements

This section provides additional information about alias names, service interdependencies, algorithm parameter generators and algorithm parameters.

Alias Names

In the JDK, the aliasing scheme enables clients to use aliases when referring to algorithms or types, rather than the standard names.

For many cryptographic algorithms and types, there is a single official "standard name" defined in the Java Security Standard Algorithm Names.



For example, "SHA-256" is the standard name for the SHA-256 Message Digest algorithm defined in RFC 1321. DiffieHellman is the standard for the Diffie-Hellman key agreement algorithm defined in PKCS3.

In the JDK, there is an aliasing scheme that enables clients to use aliases when referring to algorithms or types, rather than their standard names.

For example, the "SUN" provider's master class (Sun.java) defines the alias "SHA1/DSA" for the algorithm whose standard name is "SHA1withDSA". Thus, the following statements are equivalent:

```
Signature sig = Signature.getInstance("SHAlwithDSA", "SUN");
Signature sig = Signature.getInstance("SHA1/DSA", "SUN");
```

Aliases can be defined in your "master class" (see Step 3: Write Your Master Class, a Subclass of Provider). To define an alias, create a property named

```
Alg.Alias.engineClassName.aliasName
```

where *engineClassName* is the name of an engine class (e.g., signature), and *aliasName* is your alias name. The *value* of the property must be the standard algorithm (or type) name for the algorithm (or type) being aliased.

As an example, the "SUN" provider defines the alias "SHA1/DSA" for the signature algorithm whose standard name is "SHA1withDSA" by setting a property named Alg.Alias.Signature.SHA1/DSA to have the value SHA1withDSA via the following:

```
put("Alq.Alias.Signature.SHA1/DSA", "SHA1withDSA");
```



The aliases defined by one provider are available only to that provider and not to any other providers. Thus, aliases defined by the *SunJCE* provider are available only to the *SunJCE* provider.

Service Interdependencies

Some algorithms require the use of other types of algorithms. For example, a PBE algorithm usually needs to use a message digest algorithm in order to transform a password into a key.

If you are implementing one type of algorithm that requires another, you can do one of the following:

- Provide your own implementations for both.
- Let your implementation of one algorithm use an instance of the other type of algorithm, as supplied by the default Sun provider that is included with every Java SE Platform installation. For example, if you are implementing a PBE algorithm that requires a message digest algorithm, you can obtain an instance of a class implementing the SHA256 message digest algorithm by calling:

```
MessageDigest.getInstance("SHA256", "SUN")
```



- Let your implementation of one algorithm use an instance of the other type of algorithm, as supplied by another specific provider. This is only appropriate if you are sure that all clients who will use your provider will also have the other provider installed.
- Let your implementation of one algorithm use an instance of the other type of algorithm, as supplied by another (unspecified) provider. That is, you can request an algorithm by name, but without specifying any particular provider, as in:

```
MessageDigest.getInstance("SHA256")
```

This is only appropriate if you are sure that there will be at least one implementation of the requested algorithm (in this case, SHA256) installed on each Java platform where your provider will be used.

Here are some common types of algorithm interdependencies:

Signature and Message Digest Algorithms

A signature algorithm often requires use of a message digest algorithm. For example, the SHA256withDSA signature algorithm requires the SHA256 message digest algorithm.

Signature and (Pseudo-)Random Number Generation Algorithms

A signature algorithm often requires use of a (pseudo-)random number generation algorithm. For example, such an algorithm is required in order to generate a DSA signature.

Key Pair Generation and Message Digest Algorithms

A key pair generation algorithm often requires use of a message digest algorithm. For example, DSA keys are generated using the SHA-256 message digest algorithm.

Algorithm Parameter Generation and Message Digest Algorithms

An algorithm parameter generator often requires use of a message digest algorithm. For example, DSA parameters are generated using the SHA-256 message digest algorithm.

Keystores and Message Digest Algorithms

A keystore implementation will often utilize a message digest algorithm to compute keyed hashes (where the key is a user-provided password) to check the integrity of a keystore and make sure that the keystore has not been tampered with.

Key Pair Generation Algorithms and Algorithm Parameter Generators

A key pair generation algorithm sometimes needs to generate a new set of algorithm parameters. It can either generate the parameters directly, or use an algorithm parameter generator.

Key Pair Generation, Algorithm Parameter Generation, and (Pseudo-)Random Number Generation Algorithms

A key pair generation algorithm may require a source of randomness in order to generate a new key pair and possibly a new set of parameters associated with the keys. That source of randomness is represented by a SecureRandom object. The implementation of the key pair generation algorithm may generate the key parameters



itself, or may use an algorithm parameter generator to generate them, in which case it may or may not initialize the algorithm parameter generator with a source of randomness.

Algorithm Parameter Generators and Algorithm Parameters

An algorithm parameter generator's engineGenerateParameters method must return an AlgorithmParameters instance.

Signature and Key Pair Generation Algorithms or Key Factories

If you are implementing a signature algorithm, your implementation's engineInitSign and engineInitVerify methods will require passed-in keys that are valid for the underlying algorithm (e.g., DSA keys for the DSS algorithm). You can do one of the following:

- Also create your own classes implementing appropriate interfaces (e.g. classes implementing the DSAPrivateKey and DSAPublicKey interfaces from the package java.security.interfaces), and create your own key pair generator and/or key factory returning keys of those types. Require the keys passed to engineInitSign and engineInitVerify to be the types of keys you have implemented, that is, keys generated from your key pair generator or key factory. Or you can,
- Accept keys from other key pair generators or other key factories, as long as they
 are instances of appropriate interfaces that enable your signature implementation
 to obtain the information it needs (such as the private and public keys and the key
 parameters). For example, the engineInitSign method for a DSS Signature
 class could accept any private keys that are instances of
 java.security.interfaces.DSAPrivateKey.

Keystores and Key and Certificate Factories

A keystore implementation will often utilize a key factory to parse the keys stored in the keystore, and a certificate factory to parse the certificates stored in the keystore.

Default Initialization

In case the client does not explicitly initialize a key pair generator or an algorithm parameter generator, each provider of such a service must supply (and document) a default initialization.

Sun

Default Key Pair Generator Parameter Requirements

If you implement a key pair generator, your implementation should supply default parameters that are used when clients don't specify parameters.

The documentation you supply (Step 12: Document Your Provider and Its Supported Services) should state what the default parameters are.

For example, the DSA key pair generator in the *Sun* provider supplies a set of precomputed p, q, and g default values for the generation of 512, 768, 1024, and 2048-bit key pairs. The following p, q, and g values are used as the default values for the generation of 1024-bit DSA key pairs:



```
p = fd7f5381 1d751229 52df4a9c 2eece4e7 f611b752 3cef4400 c31e3f80 b6512669 455d4022 51fb593d 8d58fabf c5f5ba30 f6cb9b55 6cd7813b 801d346f f26660b7 6b9950a5 a49f9fe8 047b1022 c24fbba9 d7feb7c6 1bf83b57 e7c6a8a6 150f04fb 83f6d3c5 lec30235 54135a16 9132f675 f3ae2b61 d72aeff2 2203199d d14801c7

q = 9760508f 15230bcc b292b982 a2eb840b f0581cf5

g = f7e1a085 d69b3dde cbbcab5c 36b857b9 7994afbb fa3aea82 f9574c0b 3d078267 5159578e bad4594f e6710710 8180b449 167123e8 4c281613 b7cf0932 8cc8a6e1 3c167a8b 547c8d28 e0a3ae1e 2bb3a675 916ea37f 0bfa2135 62f1fb62 7a01243b cca4f1be a8519089 a883dfe1 5ae59f06 928b665e 807b5525 64014c3b fecf492a
```

(The $\tt p$ and $\tt q$ values given here were generated by the prime generation standard, using the 160-bit

```
SEED: 8d515589 4229d5e6 89ee01e6 018a237e 2cae64cd
```

With this seed, the algorithm found p and q when the counter was at 92.)

The Provider Service Class

Provider. Service class offers an alternative way for providers to advertise their services and supports additional features.

Since its introduction, security providers have published their service information via appropriately formatted key-value String pairs they put in their Hashtable entries. While this mechanism is simple and convenient, it limits the amount customization possible. As a result, JDK 5.0 introduced a second option, the Provider. Service class. It offers an alternative way for providers to advertise their services and supports additional features as described below. Note that this addition is fully compatible with the older method of using String valued Hashtable entries. A provider on JDK 5.0 can choose either method as it prefers, or even use both at the same time.

A Provider.Service object encapsulates all information about a service. This is the provider that offers the service, its type (e.g. MessageDigest or Signature), the algorithm name, and the name of the class that implements the service. Optionally, it also includes a list of alternate algorithm names for this service (aliases) and attributes, which are a map of (name, value) String pairs. In addition, it defines the methods newInstance() and supportsParameter(). They have default implementations, but can be overridden by providers if needed, as may be the case with providers that interface with hardware security tokens.

The newInstance() method is used by the security framework when it needs to construct new implementation instances. The default implementation uses reflection to invoke the standard constructor for the respective type of service. For all standard services except CertStore, this is the no-args constructor. The constructorParameter to newInstance() must be null in theses cases. For services of type CertStore, the constructor that takes a CertStoreParameters object is invoked, and constructorParameter must be a non-null instance of CertStoreParameters. A security provider can override the newInstance() method to implement instantiation as appropriate for that implementation. It could use direct invocation or call a constructor that passes additional information specific to the Provider instance or token. For example, if multiple Smartcard readers are present on the system, it might pass information about which reader the newly created service is to be associated with.



However, despite customization all implementations must follow the conventions about constructorParameter described above.

The supportsParameter () tests whether the Service can use the specified parameter. It returns false if this service cannot use the parameter. It returns true if this service can use the parameter, if a fast test is infeasible, or if the status is unknown. It is used by the security framework with some types of services to quickly exclude non-matching implementations from consideration. It is currently only defined for the following standard services: Signature, Cipher, Mac, and KeyAgreement. The parameter must be an instance of Key in these cases. For example, for Signature services, the framework tests whether the service can use the supplied Key before instantiating the service. The default implementation examines the attributes SupportedKeyFormats and SupportedKeyClasses as described below. Again, a provider may override this methods to implement additional tests.

The SupportedKeyFormats attribute is a list of the supported formats for encoded keys (as returned by key.getFormat()) separated by the "|" (pipe) character. For example, x. 509|PKCS#8. The SupportedKeyClasses attribute is a list of the names of classes of interfaces separated by the "|" character. A key object is considered to be acceptable if it is assignable to at least one of those classes or interfaces named. In other words, if the class of the key object is a subclass of one of the listed classes (or the class itself) or if it implements the listed interface. An example value is

```
"java.security.interfaces.RSAPrivateKey|java.security.interfaces.RSAPublicKey".
```

Four methods have been added to the Provider class for adding and looking up Services. As mentioned earlier, the implementation of those methods and also of the existing Properties methods have been specifically designed to ensure compatibility with existing Provider subclasses. This is achieved as follows:

If legacy Properties methods are used to add entries, the Provider class makes sure that the property strings are parsed into equivalent Service objects prior to lookup via <code>getService()</code>. Similarly, if the <code>putService()</code> method is used, equivalent property strings are placed into the provider's hashtable at the same time. If a provider implementation overrides any of the methods in the Provider class, it has to ensure that its implementation does not interfere with this conversion. To avoid problems, we recommend that implementations do not override any of the methods in the <code>Provider</code> class.

Signature Formats

The signature algorithm should specify the format in which the signature is encoded.

If you implement a signature algorithm, the documentation you supply (Step 12: Document Your Provider and Its Supported Services) should specify the format in which the signature (generated by one of the sign methods) is encoded.

For example, the *SHA1withDSA* signature algorithm supplied by the *Sun* provider encodes the signature as a standard ASN.1 sequence of two ASN.1 INTEGER values: r and s, in that order:

```
SEQUENCE ::= {
    r INTEGER,
    s INTEGER }
```



DSA Interfaces and their Required Implementations

The Java Security API contains interfaces (in the <code>java.security.interfaces</code> package) for the convenience of programmers implementing DSA services.

The Java Security API contains the following interfaces:

- Interface DSAKey
- Interface DSAKeyPairGenerator
- Interface DSAParams
- Interface DSAPrivateKey
- Interface DSAPublicKey

The following sections discuss requirements for implementations of these interfaces.

DSAKeyPairGenerator

The interface Interface DSAKeyPairGenerator is obsolete. It used to be needed to enable clients to provide DSA-specific parameters to be used rather than the default parameters your implementation supplies. However, in Java it is no longer necessary; a new KeyPairGenerator initialize method that takes an AlgorithmParameterSpec parameter enables clients to indicate algorithm-specific parameters.

DSAParams Implementation

If you are implementing a DSA key pair generator, you need a class implementing Interface DSAParams for holding and returning the p, q, and q parameters.

A DSAParams implementation is also required if you implement the DSAPrivateKey and DSAPublicKey interfaces. DSAPublicKey and DSAPrivateKey both extend the DSAKey interface, which contains a getParams method that must return a DSAParams object.



There is a DSAParams implementation built into the JDK: the java.security.spec.DSAParameterSpec Class.

DSAPrivateKey and DSAPublicKey Implementations

If you implement a DSA key pair generator or key factory, you need to create classes implementing the Interface DSAPrivateKey and Interface DSAPublicKey interfaces.

If you implement a DSA key pair generator, your <code>generateKeyPair</code> method (in your <code>KeyPairGeneratorSpi</code> subclass) will return instances of your implementations of those interfaces.

If you implement a DSA key factory, your <code>engineGeneratePrivate</code> method (in your <code>KeyFactorySpi</code> subclass) will return an instance of your <code>DSAPrivateKey</code> implementation, and your <code>engineGeneratePublic</code> method will return an instance of your <code>DSAPublicKey</code> implementation.

Also, your engineGetKeySpec and engineTranslateKey methods will expect the passed-in key to be an instance of a DSAPrivateKey or DSAPublicKey implementation. The



getParams method provided by the interface implementations is useful for obtaining and extracting the parameters from the keys and then using the parameters, for example as parameters to the DSAParameterSpec constructor called to create a parameter specification from parameter values that could be used to initialize a KeyPairGenerator object for DSA.

If you implement a DSA signature algorithm, your engineInitSign method (in your SignatureSpi subclass) will expect to be passed a DSAPrivateKey and your engineInitVerify method will expect to be passed a DSAPublicKey.

Please note: The DSAPublicKey and DSAPrivateKey interfaces define a very generic, provider-independent interface to DSA public and private keys, respectively. The engineGetKeySpec and engineTranslateKey methods (in your KeyFactorySpi subclass) could additionally check if the passed-in key is actually an instance of their provider's own implementation of DSAPrivateKey or DSAPublicKey, e.g., to take advantage of provider-specific implementation details. The same is true for the DSA signature algorithm engineInitSign and engineInitVerify methods (in your SignatureSpi subclass).

To see what methods need to be implemented by classes that implement the DSAPublicKey and DSAPrivateKey interfaces, first note the following interface signatures:

In the java.security.interfaces package:

In order to implement the DSAPrivateKey and DSAPublicKey interfaces, you must implement the methods they define as well as those defined by interfaces they extend, directly or indirectly.

Thus, for private keys, you need to supply a class that implements

- The getx method from the Interface DSAPrivateKey interface.
- The getParams method from the Interface DSAKey interface, since DSAPrivateKey extends DSAKey. Note: The getParams method returns a DSAParams object, so you must also have a DSAParams implementation.
- The getAlgorithm, getEncoded, and getFormat methods from the Interface Key interface, Since DSAPrivateKey extends java.security.PrivateKey, and PrivateKey extends Key.
 - Similarly, for public DSA keys, you need to supply a class that implements:
 - The gety method from the Interface DSAPublicKey interface.



The getParams method from the Interface DSAKey interface, since DSAPublicKey extends DSAKey.



The getParams method returns a DSAParams object, so you must also have a DSAParams Implementation.

 The getAlgorithm, getEncoded, and getFormat methods from the Interface Key, since DSAPublicKey extends java.security.PublicKey, and PublicKey extends Key.

RSA Interfaces and their Required Implementations

The Java Security API contains the interfaces (in the <code>java.security.interfaces</code> package) for the convenience of programmers implementing RSA services.

- Interface RSAPrivateKey
- Interface RSAPrivateCrtKey
- Interface RSAPublicKey

The following sections discuss requirements for implementations of these interfaces.

RSAPrivateKey, RSAPrivateCrtKey, and RSAPublicKey Implementations

If you implement an RSA key pair generator or key factory, you need to create classes implementing the Interface RSAPublicKey (and/or Interface RSAPrivateCrtKey) and Interface RSAPublicKey interfaces. (RSAPrivateCrtKey is the interface to an RSA private key, using the *Chinese Remainder Theorem* (CRT) representation.)

If you implement an RSA key pair generator, your <code>generateKeyPair</code> method (in your <code>KeyPairGeneratorSpi</code> subclass) will return instances of your implementations of those interfaces.

If you implement an RSA key factory, your engineGeneratePrivate method (in your KeyFactorySpi subclass) will return an instance of your RSAPrivateKey (or RSAPrivateCrtKey) implementation, and your engineGeneratePublic method will return an instance of your RSAPublicKey implementation.

Also, your engineGetKeySpec and engineTranslateKey methods will expect the passed-in key to be an instance of an RSAPrivateKey, RSAPrivateCrtKey, Or RSAPublicKey implementation.

If you implement an RSA signature algorithm, your <code>engineInitSign</code> method (in your <code>SignatureSpi</code> subclass) will expect to be passed either an <code>RSAPrivateKey</code> or an <code>RSAPrivateCrtKey</code>, and your <code>engineInitVerify</code> method will expect to be passed an <code>RSAPublicKey</code>.

Please note: The RSAPublicKey, RSAPrivateKey, and RSAPrivateCrtKey interfaces define a very generic, provider-independent interface to RSA public and private keys. The engineGetKeySpec and engineTranslateKey methods (in your KeyFactorySpi subclass) could additionally check if the passed-in key is actually an instance of their provider's own implementation of RSAPrivateKey, RSAPrivateCrtKey, Or RSAPublicKey, e.g., to take advantage of provider-specific implementation details. The same is true for the RSA



Signature algorithm engineInitSign and engineInitVerify methods (in your SignatureSpi Subclass).

To see what methods need to be implemented by classes that implement the RSAPublicKey, RSAPrivateKey, and RSAPrivateCrtKey interfaces, first note the following interface signatures:

In the java.security.interfaces package:

```
public interface RSAPrivateKey extends java.security.PrivateKey
public interface RSAPrivateCrtKey extends RSAPrivateKey
public interface RSAPublicKey extends java.security.PublicKey
```

In the java.security package:

```
public interface PrivateKey extends Key
public interface PublicKey extends Key
public interface Key extends java.io.Serializable
```

In order to implement the RSAPrivateKey, RSAPrivateCrtKey, and RSAPublicKey interfaces, you must implement the methods they define as well as those defined by interfaces they extend, directly or indirectly.

Thus, for RSA private keys, you need to supply a class that implements:

- The getModulus and getPrivateExponent methods from the Interface RSAPrivateKey interface.
- The getAlgorithm, getEncoded, and getFormat methods from the Interface Key interface, since RSAPrivateKey extends java.security.PrivateKey, and PrivateKey extends Key.

Similarly, for RSA private keys using the *Chinese Remainder Theorem* (CRT) representation, you need to supply a class that implements:

- All the methods listed above for RSA private keys, since RSAPrivateCrtKey extends java.security.interfaces.RSAPrivateKey.
- The getPublicExponent, getPrimeP, getPrimeQ, getPrimeExponentP, getPrimeExponentQ, and getCrtCoefficient methods from the Interface RSAPrivateKey interface.

For public RSA keys, you need to supply a class that implements:

- The getModulus and getPublicExponent methods from the Interface RSAPublicKey interface.
- The getAlgorithm, getEncoded, and getFormat methods from the Interface Key interface, since RSAPublicKey extends java.security.PublicKey, and PublicKey extends Key.

JCA contains a number of AlgorithmParameterSpec implementations for the most frequently used cipher and key agreement algorithm parameters. If you are operating on algorithm parameters that should be for a different type of algorithm not provided by JCA, you will need to supply your own AlgorithmParameterSpec implementation appropriate for that type of algorithm.



Diffie-Hellman Interfaces and their Required Implementations

JCA contains interfaces (in the <code>javax.crypto.interfaces</code> package) for the convenience of programmers implementing Diffie-Hellman services.

- Interface DHPublicKey
- Interface DHKey
- Interface DHPrivateKey

The following sections discuss requirements for implementations of these interfaces.

DHPrivateKey and DHPublicKey Implementations

If you implement a Diffie-Hellman key pair generator or key factory, you need to create classes implementing the Interface DHPrivateKey and Interface DHPublicKey interfaces.

If you implement a Diffie-Hellman key pair generator, your <code>generateKeyPair</code> method (in your <code>KeyPairGeneratorSpi</code> subclass) will return instances of your implementations of those interfaces.

If you implement a Diffie-Hellman key factory, your engineGeneratePrivate method (in your KeyFactorySpi subclass) will return an instance of your DHPrivateKey implementation, and your engineGeneratePublic method will return an instance of your DHPublicKey implementation.

Also, your engineGetKeySpec and engineTranslateKey methods will expect the passed-in key to be an instance of a DHPrivateKey or DHPublicKey implementation. The getParams method provided by the interface implementations is useful for obtaining and extracting the parameters from the keys. You can then use the parameters, for example, as parameters to the DHParameterSpec constructor called to create a parameter specification from parameter values used to initialize a KeyPairGenerator object for Diffie-Hellman.

If you implement the Diffie-Hellman key agreement algorithm, your <code>engineInit</code> method (in your <code>KeyAgreementSpi</code> subclass) will expect to be passed a <code>DHPrivateKey</code> and your <code>engineDoPhase</code> method will expect to be passed a <code>DHPublicKey</code>.

Note:

The DHPublicKey and DHPrivateKey interfaces define a very generic, provider-independent interface to Diffie-Hellman public and private keys, respectively. The engineGetKeySpec and engineTranslateKey methods (in your KeyFactorySpi subclass) could additionally check if the passed-in key is actually an instance of their provider's own implementation of DHPrivateKey or DHPublicKey, e.g., to take advantage of provider-specific implementation details. The same is true for the Diffie-Hellman algorithm engineInit and engineDoPhase methods (in your KeyAgreementSpi subclass).

To see what methods need to be implemented by classes that implement the <code>DHPublicKey</code> and <code>DHPrivateKey</code> interfaces, first note the following interface signatures:

In the javax.crypto.interfaces package:



```
public interface DHPrivateKey extends DHKey, java.security.PrivateKey
public interface DHPublicKey extends DHKey, java.security.PublicKey
public interface DHKey
```

In the java.security package:

```
public interface PrivateKey extends Key
public interface PublicKey extends Key
public interface Key extends java.io.Serializable
```

To implement the <code>DHPrivateKey</code> and <code>DHPublicKey</code> interfaces, you must implement the methods they define as well as those defined by interfaces they extend, directly or indirectly.

Thus, for private keys, you need to supply a class that implements:

- The getx method from the Interface DHPrivateKey interface.
- The getParams method from the Interface DHKey interface, since DHPrivateKey extends DHKey.
- The getAlgorithm, getEncoded, and getFormat methods from the Interface Key
 interface, Since DHPrivateKey extends java.security.PrivateKey, and PrivateKey
 extends Key.

Similarly, for public Diffie-Hellman keys, you need to supply a class that implements:

- The gety method from the Interface DHPublicKey interface.
- The getParams method from the Interface DHKey interface, since DHPublicKey extends DHKey.
- The getAlgorithm, getEncoded, and getFormat methods from the Interface Key
 interface, Since DHPublicKey extends java.security.PublicKey, and PublicKey
 extends Key.

Interfaces for Other Algorithm Types

As noted above, the Java Security API contains interfaces for the convenience of programmers implementing services like DSA, RSA and ECC. If there are services without API support, you need to define your own APIs.

If you are implementing a key pair generator for a different algorithm, you should create an interface with one or more <code>initialize</code> methods that clients can call when they want to provide algorithm-specific parameters to be used rather than the default parameters your implementation supplies. Your subclass of <code>KeyPairGeneratorSpi</code> should implement this interface.

For algorithms without direct API support, it is recommended that you create similar interfaces and provide implementation classes. Your public key interface should extend the Interface PublicKey interface. Similarly, your private key interface should extend the Interface PrivateKey interface.



Algorithm Parameter Specification Interfaces and Classes

An algorithm parameter specification is a transparent representation of the sets of parameters used with an algorithm.

A *transparent* representation of parameters means that you can access each value individually, through one of the *get* methods defined in the corresponding specification class (e.g., DSAParameterSpec defines getP, getQ, and getG methods, to access the p, q, and g parameters, respectively).

This is contrasted with an *opaque* representation, as supplied by the AlgorithmParameters engine class, in which you have no direct access to the key material values; you can only get the name of the algorithm associated with the parameter set (via getAlgorithm) and some kind of encoding for the parameter set (via getEncoded).

If you supply an AlgorithmParametersSpi, AlgorithmParameterGeneratorSpi, or KeyPairGeneratorSpi implementation, you must utilize the AlgorithmParameterSpec interface, since each of those classes contain methods that take an AlgorithmParameterSpec parameter. Such methods need to determine which actual implementation of that interface has been passed in, and act accordingly.

JCA contains a number of AlgorithmParameterSpec implementations for the most frequently used signature, cipher and key agreement algorithm parameters. If you are operating on algorithm parameters that should be for a different type of algorithm not provided by JCA, you will need to supply your own AlgorithmParameterSpec implementation appropriate for that type of algorithm.

Java defines the following algorithm parameter specification interfaces and classes in the <code>java.security.spec</code> and <code>javax.crypto.spec</code> packages:

The AlgorithmParameterSpec Interface

AlgorithmParameterSpec is an interface to a transparent specification of cryptographic parameters.

This interface contains no methods or constants. Its only purpose is to group (and provide type safety for) all parameter specifications. All parameter specifications must implement this interface.

The DSAParameterSpec Class

This class (which implements the AlgorithmParameterSpec and DSAParams interfaces) specifies the set of parameters used with the DSA algorithm. It has the following methods:

```
public BigInteger getP()
public BigInteger getQ()
public BigInteger getG()
```

These methods return the DSA algorithm parameters: the prime ${\tt p}$, the sub-prime ${\tt q}$, and the base ${\tt g}$.



Many types of DSA services will find this class useful - for example, it is utilized by the DSA signature, key pair generator, algorithm parameter generator, and algorithm parameters classes implemented by the *Sun* provider. As a specific example, an algorithm parameters implementation must include an implementation for the getParameterSpec method, which returns an AlgorithmParameterSpec. The DSA algorithm parameters implementation supplied by *Sun* returns an instance of the DSAParameterSpec class.

The IvParameterSpec Class

This class (which implements the AlgorithmParameterSpec interface) specifies the initialization vector (IV) used with a cipher in feedback mode.

Table 3-3 Method in IvParameterSpec

Method	Description
byte[] getIV()	Returns the initialization vector (IV).

The OAEPParameterSpec Class

This class specifies the set of parameters used with OAEP Padding, as defined in the PKCS #1 standard.

Table 3-4 Methods in OAEPParameterSpec

Method	Description
String getDigestAlgorithm()	Returns the message digest algorithm name.
String getMGFAlgorithm()	Returns the mask generation function algorithm name.
AlgorithmParameterSpec getMGFParameters()	Returns the parameters for the mask generation function.
PSource getPSource()	Returns the source of encoding input P.

The PBEParameterSpec Class

This class (which implements the AlgorithmParameterSpec interface) specifies the set of parameters used with a password-based encryption (PBE) algorithm.

Table 3-5 Methods in PBEParameterSpec

Method	Description
int getIterationCount()	Returns the iteration count.
<pre>byte[] getSalt()</pre>	Returns the salt.

The RC2ParameterSpec Class

This class (which implements the AlgorithmParameterSpec interface) specifies the set of parameters used with the RC2 algorithm.



Table 3-6 Methods in RC2ParameterSpec

Method	Description
boolean equals(Object obj)	Tests for equality between the specified object and this object.
<pre>int getEffectiveKeyBits()</pre>	Returns the effective key size in bits.
<pre>byte[] getIV()</pre>	Returns the IV or null if this parameter set does not contain an IV.
<pre>int hashCode()</pre>	Calculates a hash code value for the object.

The RC5ParameterSpec Class

This class (which implements the AlgorithmParameterSpec interface) specifies the set of parameters used with the RC5 algorithm.

Table 3-7 Methods in RC5ParameterSpec

Method	Description
boolean equals(Object obj)	Tests for equality between the specified object and this object.
<pre>byte[] getIV()</pre>	Returns the IV or null if this parameter set does not contain an IV.
<pre>int getRounds()</pre>	Returns the number of rounds.
<pre>int getVersion()</pre>	Returns the version.
int getWordSize()	Returns the word size in bits.
<pre>int hashCode()</pre>	Calculates a hash code value for the object.

The DHParameterSpec Class

This class (which implements the AlgorithmParameterSpec interface) specifies the set of parameters used with the Diffie-Hellman algorithm.

Table 3-8 Methods in DHParameterSpec

Method	Description
BigInteger getG()	Returns the base generator g.
<pre>int getL()</pre>	Returns the size in bits, 1, of the random exponent (private value).
BigInteger getP()	Returns the prime modulus p.

Many types of Diffie-Hellman services will find this class useful; for example, it is used by the Diffie-Hellman key agreement, key pair generator, algorithm parameter generator, and algorithm parameters classes implemented by the "SunJCE" provider. As a specific example, an algorithm parameters implementation must include an implementation for the <code>getParameterSpec</code> method, which returns an <code>AlgorithmParameterSpec</code>. The Diffie-Hellman algorithm parameters implementation supplied by "SunJCE" returns an instance of the <code>DHParameterSpec</code> class.



Key Specification Interfaces and Classes Required by Key Factories

A key factory provides bi-directional conversions between opaque keys (of type $\kappa_{\rm ey}$) and key specifications. If you implement a key factory, you thus need to understand and utilize key specifications. In some cases, you also need to implement your own key specifications.

Key specifications are transparent representations of the key material that constitutes a key. If the key is stored on a hardware device, its specification may contain information that helps identify the key on the device.

A *transparent* representation of keys means that you can access each key material value individually, through one of the *get* methods defined in the corresponding specification class. For example, <code>java.security.spec.DSAPrivateKeySpec</code> defines <code>getX</code>, <code>getP</code>, <code>getQ</code>, and <code>getG</code> methods, to access the private key x, and the DSA algorithm parameters used to calculate the key: the prime p, the sub-prime q, and the base g.

This is contrasted with an *opaque* representation, as defined by the Key interface, in which you have no direct access to the parameter fields. In other words, an "opaque" representation gives you limited access to the key - just the three methods defined by the Key interface: getAlgorithm, getFormat, and getEncoded.

A key may be specified in an algorithm-specific way, or in an algorithm-independent encoding format (such as ASN.1). For example, a DSA private key may be specified by its components x, p, q, and g (see DSAPrivateKeySpec), or it may be specified using its DER encoding (see PKCS8EncodedKeySpec).

Java defines the following key specification interfaces and classes in the java.security.spec and javax.crypto.spec packages:

The KeySpec Interface

This interface contains no methods or constants. Its only purpose is to group (and provide type safety for) all key specifications. All key specifications must implement this interface.

Java supplies several classes implementing the KeySpec interface:

- DSAPrivateKeySpec
- DSAPublicKeySpec
- RSAPrivateKeySpec
- RSAPublicKeySpec
- EncodedKeySpec
- PKCS8EncodedKeySpec
- X509EncodedKeySpec

If your provider uses key types (e.g., Your_PublicKey_type and Your_PrivateKey_type) for which the JDK does not already provide corresponding KeySpec classes, there are two possible scenarios, one of which requires that you implement your own key specifications:

1. If your users will never have to access specific key material values of your key type, you will not have to provide any KeySpec classes for your key type.



In this scenario, your users will always create Your_PublicKey_type and Your_PrivateKey_type keys through the appropriate KeyPairGenerator supplied by your provider for that key type. If they want to store the generated keys for later usage, they retrieve the keys' encodings (using the getEncoded method of the Key interface). When they want to create an Your_PublicKey_type or Your_PrivateKey_type key from the encoding (e.g., in order to initialize a Signature object for signing or verification), they create an instance of X509EncodedKeySpec or PKCS8EncodedKeySpec from the encoding, and feed it to the appropriate KeyFactory supplied by your provider for that algorithm, whose generatePublic and generatePrivate methods will return the requested PublicKey (an instance of Your_PrivateKey_type) Or PrivateKey (an instance of Your_PrivateKey_type) Object, respectively.

2. If you anticipate a need for users to access specific key material values of your key type, or to construct a key of your key type from key material and associated parameter values, rather than from its encoding (as in the above case), you have to specify new KeySpec classes (classes that implement the KeySpec interface) with the appropriate constructor methods and *get* methods for returning key material fields and associated parameter values for your key type. You will specify those classes in a similar manner as is done by the DSAPrivateKeySpec and DSAPublicKeySpec classes. You need to ship those classes along with your provider classes, for example, as part of your provider JAR file.

The DSAPrivateKeySpec Class

This class (which implements the KeySpec Interface) specifies a DSA private key with its associated parameters. It has the following methods:

Table 3-9 Methods in DSAPrivateKeySpec

Method in DSAPrivateKeySpec	Description
<pre>public BigInteger getX()</pre>	Returns the private key x.
<pre>public BigInteger getP()</pre>	Returns the prime p.
<pre>public BigInteger getQ()</pre>	Returns the sub-prime q.
<pre>public BigInteger getG()</pre>	Returns the base g.

These methods return the private key x, and the DSA algorithm parameters used to calculate the key: the prime p, the sub-prime q, and the base g.

The DSAPublicKeySpec Class

This class (which implements the KeySpec Interface) specifies a DSA public key with its associated parameters. It has the following methods:

Table 3-10 Methods in DSAPublicKeySpec

Method in DSAPublicKeySpec	Description
<pre>public BigInteger getY()</pre>	returns the public key y.
<pre>public BigInteger getP()</pre>	Returns the prime p.
<pre>public BigInteger getQ()</pre>	Returns the sub-prime q.
<pre>public BigInteger getG()</pre>	Returns the base g.



The RSAPrivateKeySpec Class

This class (which implements the KeySpec Interface) specifies an RSA private key. It has the following methods:

Table 3-11 Methods in RSAPrivateKeySpec

Method in RSAPrivateKeySpec	Description
<pre>public BigInteger getModulus()</pre>	Returns the modulus.
<pre>public BigInteger getPrivateExponent()</pre>	Returns the private exponent.

These methods return the RSA modulus $\tt n$ and private exponent $\tt d$ values that constitute the RSA private key.

The RSAPrivateCrtKeySpec Class

This class (which extends the RSAPrivateKeySpec class) specifies an RSA private key, as defined in the PKCS#1 standard, using the *Chinese Remainder Theorem* (CRT) information values. It has the following methods (in addition to the methods inherited from its superclass RSAPrivateKeySpec):

Table 3-12 Methods in RSAPrivateCrtKeySpec

Method in RSAPrivateCrtKeySpec	Description
<pre>public BigInteger getPublicExponent()</pre>	Returns the public exponent.
<pre>public BigInteger getPrimeP()</pre>	Returns the prime P.
<pre>public BigInteger getPrimeQ()</pre>	Returns the prime Q.
<pre>public BigInteger getPrimeExponentP()</pre>	Returns the primeExponentP.
<pre>public BigInteger getPrimeExponentQ()</pre>	Returns the primeExponentQ.
<pre>public BigInteger getCrtCoefficient()</pre>	Returns the crtCoefficient.

These methods return the public exponent e and the CRT information integers: the prime factor p of the modulus n, the prime factor q of n, the exponent $d \mod (p-1)$, the exponent $d \mod (q-1)$, and the Chinese Remainder Theorem coefficient (inverse of q) mod p.

An RSA private key logically consists of only the modulus and the private exponent. The presence of the CRT values is intended for efficiency.

The RSAPublicKeySpec Class

This class (which implements the KeySpec Interface) specifies an RSA public key. It has the following methods:

Table 3-13 Methods in RSAPublicKeySpec

Method in RSAPublicKeySpec	Description
<pre>public BigInteger getModulus()</pre>	Returns the modulus.
<pre>public BigInteger getPublicExponent()</pre>	Returns the public exponent.



The EncodedKeySpec Class

This abstract class (which implements the KeySpec Interface) represents a public or private key in encoded format.

Table 3-14 Methods in EncodedKeySpec

Method in EncodedKeySpec	Description
<pre>public abstract byte[] getEncoded()</pre>	Returns the encoded key.
<pre>public abstract String getFormat()</pre>	Returns the name of the encoding format.

The JDK supplies two classes implementing the <code>EncodedKeySpec</code> interface: <code>PKCS8EncodedKeySpec</code> and <code>X509EncodedKeySpec</code>. If desired, you can supply your own <code>EncodedKeySpec</code> implementations for those or other types of key encodings.

The PKCS8EncodedKeySpec Class

This class, which is a subclass of <code>EncodedKeySpec</code>, represents the DER encoding of a private key, according to the format specified in the PKCS #8 standard.

Its getEncoded method returns the key bytes, encoded according to the PKCS #8 standard. Its getFormat method returns the string "PKCS#8".

The X509EncodedKeySpec Class

This class, which is a subclass of <code>EncodedKeySpec</code>, represents the DER encoding of a public or private key, according to the format specified in the X.509 standard.

Its getEncoded method returns the key bytes, encoded according to the X.509 standard. Its getFormat method returns the string "X.509".DHPrivateKeySpec, DHPublicKeySpec, DESKeySpec, DESedeKeySpec, PBEKeySpec, and SecretKeySpec.

The DHPrivateKeySpec Class

This class (which implements the KeySpec interface) specifies a Diffie-Hellman private key with its associated parameters.

Table 3-15 Methods in DHPrviateKeySpec

Method in DHPrivateKeySpec	Description
BigInteger getG()	Returns the base generator g.
<pre>BigInteger getP()</pre>	Returns the prime modulus p.
<pre>BigInteger getX()</pre>	Returns the private value \mathbf{x} .

The DHPublicKeySpec Class

Table 3-16 Methods in DHPublicKeySpec

Method in DHPublicKeySpec	Description
BigInteger getG()	Returns the base generator g.
<pre>BigInteger getP()</pre>	Returns the prime modulus p.
BigInteger getY()	Returns the public value y.



The DESKeySpec Class

This class (which implements the KeySpec interface) specifies a DES key.

Table 3-17 Methods in DESKeySpec

Method in DESKeySpec	Description
byte[] getKey()	Returns the DES key bytes.
<pre>static boolean isParityAdjusted(byte[] key, int offset)</pre>	Checks if the given DES key material is parity-adjusted.
<pre>static boolean isWeak(byte[] key, int offset)</pre>	Checks if the given DES key material is weak or semi-weak.

The DESedeKeySpec Class

This class (which implements the KeySpec interface) specifies a DES-EDE (Triple DES) key.

Table 3-18 Methods in DESedeKeySpec

Method in DESedeKeySpec	Description
<pre>byte[] getKey()</pre>	Returns the DES-EDE key.
<pre>static boolean isParityAdjusted(byte[] key, int offset)</pre>	Checks if the given DES-EDE key is parity-adjusted.

The PBEKeySpec Class

This class implements the κ_{eySpec} interface. A user-chosen password can be used with password-based encryption (PBE); the password can be viewed as a type of raw key material. An encryption mechanism that uses this class can derive a cryptographic key from the raw key material.

Table 3-19 Methods in PBEKeySpec

Method in PBEKeySpec	Description
void clearPassword	Clears the internal copy of the password.
int getIterationCount	Returns the iteration count or 0 if not specified.
int getKeyLength	Returns the to-be-derived key length or 0 if not specified.
char[] getPassword	Returns a copy of the password.
byte[] getSalt	Returns a copy of the salt or null if not specified.

The SecretKeySpec Class

This class implements the <code>KeySpec</code> interface. Since it also implements the <code>SecretKey</code> interface, it can be used to construct a <code>SecretKey</code> object in a provider-independent fashion, i.e., without having to go through a provider-based <code>SecretKeyFactory</code>.



Method in SecretKeySpec	Description
boolean equals (Object obj)	Indicates whether some other object is "equal to" this one.
String getAlgorithm()	Returns the name of the algorithm associated with this secret key.
<pre>byte[] getEncoded()</pre>	Returns the key material of this secret key.
String getFormat()	Returns the name of the encoding format for this secret key.
<pre>int hashCode()</pre>	Calculates a hash code value for the object.

Secret-Key Generation

If you provide a secret-key generator (subclass of <code>javax.crypto.KeyGeneratorSpi</code>) for a particular secret-key algorithm, you may return the generated secret-key object.

The generated secret-key object (which must be an instance of <code>javax.crypto.SecretKey</code>, see engineGenerateKey) can be returned in one of the following ways:

- You implement a class whose instances represent secret-keys of the algorithm
 associated with your key generator. Your key generator implementation returns
 instances of that class. This approach is useful if the keys generated by your key
 generator have provider-specific properties.
- Your key generator returns an instance of SecretKeySpec, which already implements the javax.crypto.SecretKey interface. You pass the (raw) key bytes and the name of the secret-key algorithm associated with your key generator to the SecretKeySpec constructor. This approach is useful if the underlying (raw) key bytes can be represented as a byte array and have no key-parameters associated with them.

Adding New Object Identifiers

The following information applies to providers who supply an algorithm that is not listed as one of the standard algorithms in Java Security Standard Algorithm Names Specification.

Mapping from OID to Name

Sometimes the JCA needs to instantiate a cryptographic algorithm implementation from an algorithm identifier (for example, as encoded in a certificate), which by definition includes the object identifier (OID) of the algorithm. For example, in order to verify the signature on an X.509 certificate, the JCA determines the signature algorithm from the signature algorithm identifier that is encoded in the certificate, instantiates a Signature object for that algorithm, and initializes the Signature object for verification.

For the JCA to find your algorithm, you must provide the object identifier of your algorithm as an alias entry for your algorithm in the provider master file.



Note that if your algorithm is known under more than one object identifier, you need to create an alias entry for each object identifier under which it is known.

An example of where the JCA needs to perform this type of mapping is when your algorithm ("Foo") is a signature algorithm and users run the keytool command and specify your (signature) algorithm alias.

```
% keytool -genkeypair -sigalg 1.2.3.4.5.6.7.8
```

In this case, your provider master file should contain the following entries:

```
put("Signature.Foo", "com.xyz.MyFooSignatureImpl");
put("Alq.Alias.Signature.1.2.3.4.5.6.7.8", "Foo");
```

Other examples of where this type of mapping is performed are (1) when your algorithm is a keytype algorithm and your program parses a certificate (using the X. 509 implementation of the SUN provider) and extracts the public key from the certificate in order to initialize a Signature object for verification, and (2) when keytool users try to access a private key of your keytype (for example, to perform a digital signature) after having generated the corresponding keypair. In these cases, your provider master file should contain the following entries:

```
put("KeyFactory.Foo", "com.xyz.MyFooKeyFactoryImpl");
put("Alg.Alias.KeyFactory.1.2.3.4.5.6.7.8", "Foo");
```

Mapping from Name to OID

If the JCA needs to perform the inverse mapping (that is, from your algorithm name to its associated OID), you need to provide an alias entry of the following form for one of the OIDs under which your algorithm should be known:

```
put("Alg.Alias.Signature.OID.1.2.3.4.5.6.7.8", "MySigAlg");
```

If your algorithm is known under more than one object identifier, prefix the preferred one with "OID."

An example of where the JCA needs to perform this kind of mapping is when users run keytool in any mode that takes a -sigalg option. For example, when the -genkeypair and -certreq commands are invoked, the user can specify your (signature) algorithm with the -sigalg option.

Ensuring Exportability

A key feature of JCA is the exportability of the JCA framework and of the provider cryptography implementations if certain conditions are met.

By default, an application can use cryptographic algorithms of any strength. However, due to import regulations in some countries, you may have to limit those algorithms' strength. You do this with jurisdiction policy files; see Cryptographic Strength Configuration. The JCA framework will enforce the restrictions specified in the installed jurisdiction policy files.



As noted elsewhere, you can write just one version of your provider software, implementing cryptography of maximum strength. It is up to JCA, not your provider, to enforce any jurisdiction policy file-mandated restrictions regarding the cryptographic algorithms and maximum cryptographic strengths available to applets/applications in different locations.

The conditions that must be met by your provider in order to enable it to be plugged into JCA are the following:

- The provider code should be written in such a way that provider classes become
 unusable if instantiated by an application directly, bypassing JCA. See Step 1:
 Write your Service Implementation Code in Steps to Implement and Integrate a
 Provider.
- The provider package must be signed by an entity trusted by the JCA framework. (See Step 7.1: Get a Code-Signing Certificate through Step 7.2: Sign Your Provider.) U.S. vendors whose providers may be exported outside the U.S. first need to apply for U.S. government export approval. (See Step 11: Apply for U.S. Government Export Approval If Required.)

Sample Code for MyProvider

The following is the complete source code for an example provider, MyProvider. It's a portable provider; you can specify it in a class or module path. It consists of two modules:

- com.example.MyProvider: Contains an example provider that demonstrate how to
 write a provider with the Provider.Service mechanism. You must compile,
 package, and sign the provider, then specify it in your class or module path as
 described in Steps to Implement and Integrate a Provider.
- com.example.MyApp: Contains a sample application that uses the MyProvider provider. It finds and loads this provider with the ServiceLoader mechanism, and then registers it dynamically with the Security.addProvider() method.

This example consists of the following files:

- src/com.example.MyProvider/module-info.java
- src/com.example.MyProvider/com/example/MyProvider/MyProvider.java
- src/com.example.MyProvider/com/example/MyProvider/MyCipher.java
- src/com.example.MyProvider/META-INF/services/java.security.Provider
- src/com.example.MyApp/module-info.java
- src/com.example.MyApp/com/example/MyApp/MyApp.java
- RunTest.sh

src/com.example.MyProvider/module-info.java

See Step 4: Create a Module Declaration for Your Provider for information about the module declaration, which is specified in module-info.java.

```
module com.example.MyProvider {
    provides java.security.Provider with com.example.MyProvider.MyProvider;
}
```



src/com.example.MyProvider/com/example/MyProvider/MyProvider.java

The MyProvider class is an example of a provider that uses the Provider. Service class. See Step 3.2: Create a Provider That Uses Provider. Service.

```
package com.example.MyProvider;
import java.security.*;
import java.util.*;
 * Test JCE provider.
 ^{\star} Registers services using Provider.Service and overrides newInstance().
public final class MyProvider extends Provider {
    public MyProvider() {
        super("MyProvider", "1.0", "My JCE provider");
        final Provider p = this;
        AccessController.doPrivileged((PrivilegedAction<Void>) () -> {
            putService(new ProviderService(p, "Cipher",
                    "MyCipher", "com.example.MyProvider.MyCipher"));
            return null;
        });
    private static final class ProviderService extends Provider.Service {
        ProviderService(Provider p, String type, String algo, String cn) {
            super(p, type, algo, cn, null, null);
        ProviderService(Provider p, String type, String algo, String cn,
                String[] aliases, HashMap<String, String> attrs) {
            super(p, type, algo, cn,
                    (aliases == null ? null : Arrays.asList(aliases)), attrs);
        }
        @Override
        public Object newInstance(Object ctrParamObj)
                throws NoSuchAlgorithmException {
            String type = getType();
            if (ctrParamObj != null) {
                throw new InvalidParameterException(
                        "constructorParameter not used with " + type
                        + " engines");
            String algo = getAlgorithm();
            try {
                if (type.equals("Cipher")) {
                    if (algo.equals("MyCipher")) {
                        return new MyCipher();
            } catch (Exception ex) {
```



src/com.example.MyProvider/com/example/MyProvider/MyCipher.java

The MyCipher class extends the CipherSPI, which is a Server Provider Interface (SPI). Each cryptographic service that a provider implements has a subclass of the appropriate SPI. See Step 1: Write your Service Implementation Code.



This code is only a stub provider that demonstrates how to write a provider; it's missing the actual cryptographic algorithm implementation. The MyCipher class would contain an actual cryptographic algorithm implementation if MyProvider were a real security provider.

```
package com.example.MyProvider;
import java.security.*;
import java.security.spec.*;
import javax.crypto.*;
 * Implementation represents a test Cipher.
 * All are stubs.
public class MyCipher extends CipherSpi {
    @Override
    protected byte[] engineDoFinal(byte[] input, int inputOffset, int inputLen)
            throws IllegalBlockSizeException, BadPaddingException {
        return null;
    }
    @Override
    protected int engineDoFinal(byte[] input, int inputOffset, int inputLen,
            byte[] output, int outputOffset) throws ShortBufferException,
            IllegalBlockSizeException, BadPaddingException {
        return 0;
    @Override
    protected int engineGetBlockSize() {
```

```
return 0;
@Override
protected byte[] engineGetIV() {
   return null;
@Override
protected int engineGetOutputSize(int inputLen) {
   return 0;
@Override
protected AlgorithmParameters engineGetParameters() {
   return null;
@Override
protected void engineInit(int opmode, Key key, SecureRandom random)
        throws InvalidKeyException {
@Override
protected void engineInit(int opmode, Key key,
        AlgorithmParameterSpec params, SecureRandom random)
        throws InvalidKeyException, InvalidAlgorithmParameterException {
@Override
protected void engineInit(int opmode, Key key, AlgorithmParameters params,
        SecureRandom random) throws InvalidKeyException,
        InvalidAlgorithmParameterException {
@Override
protected void engineSetMode(String mode) throws NoSuchAlgorithmException {
@Override
protected void engineSetPadding(String padding)
        throws NoSuchPaddingException {
@Override
protected int engineGetKeySize(Key key)
        throws InvalidKeyException {
    return 0;
protected byte[] engineUpdate(byte[] input, int inputOffset, int inputLen) {
    return null;
@Override
protected int engineUpdate(byte[] input, int inputOffset, int inputLen,
       byte[] output, int outputOffset) throws ShortBufferException {
   return 0;
```

src/com.example.MyProvider/META-INF/services/java.security.Provider

The java.security.Provider file enables automatic or unnamed modules to use the ServiceLoader class to search for your providers. See Step 6: Place Your Provider in a JAR File.

com.example.MyProvider.MyProvider

src/com.example.MyApp/module-info.java

This file contains a uses directive, which specifies a service that the module requires. This directive helps the module system locate providers and ensure that they run reliably. This is the complement to the provides directive in the MyProvider module definition.

```
module com.example.MyApp {
    uses java.security.Provider;
}
```

src/com.example.MyApp/com/example/MyApp/MyApp.java

```
package com.example.MyApp;
import java.util.*;
import java.security.*;
import javax.crypto.*;
 * A simple JCE test client to access a simple test Provider/Cipher
 * implementation in a signed modular jar.
public class MyApp {
    private static final String PROVIDER = "MyProvider";
    private static final String CIPHER = "MyCipher";
    public static void main(String[] args) throws Exception {
         * Registers MyProvider dynamically.
         \mbox{\ensuremath{^{\star}}} Could do statically by editing the java.security file.
         * Use the first form if using ServiceLoader ("uses" or
         ^{\star} META-INF/service), the second if using the traditional class
         * lookup method. Both if provider could be deployed to either.
         * security.provider.14=MyProvider
         * security.provider.15=com.example.MyProvider.MyProvider
        ServiceLoader<Provider> sl =
            ServiceLoader.load(java.security.Provider.class);
        for (Provider p : sl) {
            if (p.getName().equals(PROVIDER)) {
                System.out.println("Registering the Provider");
                Security.addProvider(p);
        }
         * Get a MyCipher from MyProvider and initialize it.
```



```
* /
       Cipher cipher = Cipher.getInstance(CIPHER, PROVIDER);
       cipher.init(Cipher.ENCRYPT_MODE, (Key) null);
        * What Provider did we get?
        * /
       Provider p = cipher.getProvider();
       Class c = p.getClass();
       Module m = c.getModule();
       System.out.println(p.getName() + ": version "
          + p.getVersionStr() + "\n"
          + p.getInfo() + "\n
          + ((m.getName() == null) ? "<UNNAMED>" : m.getName())
          + "/" + c.getName());
}
RunTest.sh
#!/bin/sh
# A simple example to show how a JCE provider could be developed in a
# modular JDK, for deployment as either Named/Unnamed modules.
# Edit as appropriate
JDK_DIR=d:/java/jdk9
KEYSTORE=YourKeyStore
STOREPASS=YourStorePass
SIGNER=YourAlias
echo "----"
echo "Clean/Init"
echo "----"
rm -rf mods jars
mkdir mods jars
echo "----"
echo "Compiling MyProvider"
echo "-----"
${JDK_DIR}/bin/javac.exe \
   --module-source-path src \
   $(find src/com.example.MyProvider -name '*.java' -print)
echo "-----"
echo "Packaging com.example.MyProvider.jar"
echo "-----"
{\DK_DIR}/bin/jar.exe --create \
   --file jars/com.example.MyProvider.jar \
   --verbose \
   --module-version=1.0 \
   -C mods/com.example.MyProvider . \
   -C src/com.example.MyProvider META-INF/services
echo "-----"
echo "Signing com.example.MyProvider.jar"
```



```
echo "-----"
${JDK_DIR}/bin/jarsigner.exe \
   -keystore ${KEYSTORE} \
   -storepass ${STOREPASS} \
   jars/com.example.MyProvider.jar ${SIGNER}
echo "----"
echo "Compiling MyApp"
echo "-----"
${JDK_DIR}/bin/javac.exe \
   --module-source-path src \
   -d mods \
   $(find src/com.example.MyApp -name '*.java' -print)
echo "Packaging com.example.MyApp.jar"
echo "-----"
${JDK_DIR}/bin/jar.exe --create \
   --file jars/com.example.MyApp.jar \
   --verbose \
   --module-version=1.0 \
   -C mods/com.example.MyApp .
echo "-----"
echo "Test1
echo "Named Provider/Named App"
echo "-----"
${JDK_DIR}/bin/java.exe \
   --module-path 'jars' \
   -m com.example.MyApp/com.example.MyApp.MyApp
echo "-----"
echo "Test2
echo "Named Provider/Unnamed App"
echo "-----"
${JDK_DIR}/bin/java.exe \
   --module-path 'jars/com.example.MyProvider.jar' \
   --class-path 'jars/com.example.MyApp.jar' \
   com.example.MyApp.MyApp
echo "-----"
echo "Test3
echo "Unnamed Provider/Named App"
echo "-----"
${JDK_DIR}/bin/java.exe \
   --module-path 'jars/com.example.MyApp.jar' \
   --class-path 'jars/com.example.MyProvider.jar' \
   -m com.example.MyApp/com.example.MyApp.MyApp
echo "-----"
echo "Test4
echo "Unnamed Provider/Unnamed App"
echo "-----"
${JDK_DIR}/bin/java.exe \
      'jars/com.example.MyProvider.jar;jars/com.example.MyApp.jar' \
   com.example.MyApp.MyApp
```



4

JDK Providers Documentation

This document contains the technical details of the providers that are included in the JDK. It is assumed that readers have a strong understanding of the Java Cryptography Architecture and Provider Architecture.



The Java Security Standard Algorithm Names Specification contains more information about the standard names used in this document.

Topics

Introduction to JDK Providers

Import Limits on Cryptographic Algorithms

Cipher Transformations

SecureRandom Implementations

The SunPKCS11 Provider

The SUN Provider

The SunRsaSign Provider

The SunJSSE Provider

The SunJCE Provider

The SunJGSS Provider

The SunSASL Provider

The XMLDSig Provider

The SunPCSC Provider

The SunMSCAPI Provider

The SunEC Provider

The OracleUcrypto Provider

The Apple Provider

The JdkLDAP Provider

The JdkSASL Provider



Introduction to JDK Providers

The Java platform defines a set of APIs spanning major security areas, including cryptography, public key infrastructure, authentication, secure communication, and access control. These APIs enable developers to easily integrate security mechanisms into their application code.

The Java Cryptography Architecture (JCA) and its Provider Architecture are core concepts of the Java Development Kit (JDK). It is assumed that readers have a solid understanding of this architecture.

Reminder: Cryptographic implementations in the JDK are distributed through several different providers ("SUN", "SunJSSE", "SunJCE", "SunRsaSign") for both historical reasons and by the types of services provided. General purpose applications **SHOULD NOT** request cryptographic services from specific providers. That is:

```
getInstance("...", "SunJCE"); // not recommended
versus
getInstance("..."); // recommended
```

Otherwise, applications are tied to specific providers that may not be available on other Java implementations. They also might not be able to take advantage of available optimized providers (for example, hardware accelerators via PKCS11 or native OS implementations such as Microsoft's MSCAPI) that have a higher preference order than the specific requested provider.

The following table lists the modules and the supported Java Cryptographic Service Providers:

Table 4-1 Modules and the Java Cryptographic Service Providers

Module	Provider(s)
java.base	SUN, SunRsaSign, SunJSSE, SunJCE [1], Apple
java.naming	JdkLDAP
java.security.jgss	SunJGSS
java.security.sasl	SunSASL
java.smartcardio	SunPCSC
java.xml.crypto	XMLDSig
jdk.crypto.cryptoki	SunPKCS11 [1]
jdk.crypto.ec	SunEC [1]
jdk.crypto.mscapi	SunMSCAPI [1]
jdk.crypto.ucrypto	OracleUcrypto [1]
jdk.security.jgss	JdkSASL

Footnote 1: Indicates JCE crypto providers previously distributed as signed JAR files (JCE providers contain Cipher/KeyAgreement/KeyGenerator/Mac/SecretKeyFactory implementations).



Import Limits on Cryptographic Algorithms

By default, an application can use cryptographic algorithms of any strength. However, due to import regulations in some locations, you may have to limit the strength of those algorithms. The JDK provides two different sets of jurisdiction policy files in the directory < java_home>/conf/security/policy that determine the strength of cryptographic algorithms. Information about jurisdiction policy files and how to activate them is available in Cryptographic Strength Configuration.

Consult your export/import control counsel or attorney to determine the exact requirements for your location.

For the "limited" configuration, the following table lists the maximum key sizes allowed by the "limited" set of jurisdiction policy files:

Table 4-2 Maximum Keysize of Cryptographic Algorithms

Algorithm	Maximum Keysize
DES	64
DESede	*
RC2	128
RC4 RC5	128
RC5	128
RSA	*
all others	128

Cipher Transformations

The <code>javax.crypto.Cipher.getInstance(String transformation)</code> factory method generates <code>cipher</code> objects using transformations of the form <code>algorithm/mode/padding</code>. If the mode/padding are omitted, the <code>SunJCE</code> and <code>SunPKCS11</code> providers use ECB as the default mode and <code>PKCS5Padding</code> as the default padding for many symmetric ciphers.

It is recommended to use transformations that fully specify the algorithm, mode, and padding instead of relying on the defaults. The defaults are provider specific and can vary among providers.



ECB works well for single blocks of data and can be parallelized, but absolutely should not be used for multiple blocks of data.

SecureRandom Implementations

The following table lists the default preference order of the available <code>SecureRandom</code> implementations.



Table 4-3 Default SecureRandom Implementations

os	Algorithm Name	Provider Name
Solaris	1. PKCS11 [1] [4]	SunPKCS11
	2. NativePRNG [2]	SUN
	3. DRBG	SUN
	4. SHA1PRNG [2]	SUN
	5. NativePRNGBlocking	SUN
	6. NativePRNGNonBlocking	SUN
Linux	1. NativePRNG [2]	SUN
	2. DRGB	SUN
	3. SHA1PRNG [2]	SUN
	4. NativePRNGBlocking	SUN
	5. NativePRNGNonBlocking	SUN
macOS	1. NativePRNG [2]	SUN
	2. DRGB	SUN
	3. SHA1PRNG [2]	SUN
	4. NativePRNGBlocking	SUN
	5. NativePRNGNonBlocking	SUN
Windows	1. DRGB	SUN
	2. SHA1PRNG	SUN
	3. Windows-PRNG [3]	SunMSCAPI

Footnote 1: The SunPKCS11 provider is available on all platforms, but is only enabled by default on Solaris as it is the only OS with a native PKCS11 implementation automatically installed and configured. On other platforms, applications or deployers must specifically install and configure a native PKCS11 library, and then configure and enable the SunPKCS11 provider to use it.

Footnote 2: On Solaris, Linux, and OS X, if the *entropy gathering device* in java.security is set to file:/dev/urandom or file:/dev/random, then NativePRNG is preferred to SHA1PRNG. Otherwise, SHA1PRNG is preferred.

Footnote 3: There is currently no NativePRNG on Windows. Access to the equivalent functionality is via the SunMSCAPI provider.

Footnote 4: The PKCS11 <code>SecureRandom</code> implementation for Solaris has been disabled due to the performance overhead of small-sized requests. Edit <code>sunpkcs11-solaris.cfg</code> to reenable.

If no SecureRandom implementations are registered in the JCA framework, java.security.SecureRandom uses the hardcoded SHA1PRNG.

The SunPKCS11 Provider

The Cryptographic Token Interface Standard (PKCS#11) provides native programming interfaces to cryptographic mechanisms, such as hardware cryptographic accelerators and Smart Cards. When properly configured, the SumPKCS11 provider enables applications to use the standard JCA/JCE APIs to access native PKCS#11 libraries. The SumPKCS11 provider itself does not contain cryptographic functionality, it is



simply a conduit between the Java environment and the native PKCS11 providers. The Java PKCS#11 Reference Guide has a much more detailed treatment of this provider.

The SUN Provider

JDK 1.1 introduced the Provider architecture. The first JDK provider was named SUN, and contained two types of cryptographic services (MessageDigestand Signature). In later releases, other mechanisms were added (SecureRandom, KeyPairGenerator, KeyFactory, and so on).

United States export regulations in effect at the time placed significant restrictions on the type of cryptographic functionality that could be made available internationally in the JDK. For this reason, the ${\tt SUN}$ provider has historically contained cryptographic engines that did not directly encrypt or decrypt data.

The following algorithms are available in the SUN provider:

Table 4-4 Algorithms in SUN provider

Engine	Algorithm Names
AlgorithmParameterGenerator	DSA
AlgorithmParameters	DSA
CertificateFactory	X.509
CertPathBuilder	PKIX
CertPathValidator	PKIX
CertStore	Collection
Configuration	JavaLoginConfig
KeyFactory	DSA
KeyPairGenerator	DSA
KeyStore	PKCS12
	JKS
	DKS
	CaseExactJKS
MessageDigest	MD2
	MD5
	SHA-1
	SHA-224
	SHA-256
	SHA-384
	SHA-512
	SHA-512/224
	SHA-512/256
	SHA3-224
	SHA3-256
	SHA3-384
	SHA3-512
Policy	JavaPolicy



Table 4-4 (Cont.) Algorithms in SUN provider

Engine	Algorithm Names
SecureRandom	DRBG
	(The following mechanisms and algorithms are supported: Hash_DRBG and HMAC_DRBG with SHA-224, SHA-512/224, SHA-256, SHA-512/256, SHA-384 and SHA-512. CTR_DRBG (both use derivation function and not use) with AES-128, AES-192 and AES-256. Prediction resistance and reseeding supported for each combination, and security strength can be requested from 112 up to the highest strength one supports.)
	SHA1PRNG
	(Initial seeding is currently done via a combination of system attributes and the java.security entropy gathering device.) NativePRNG
	<pre>(nextBytes() uses /dev/urandom, generateSeed() uses /dev/random)</pre>
	NativePRNGBlocking
	<pre>(nextBytes() and generateSeed() use /dev/ random)</pre>
	NativePRNGNonBlocking
	(nextBytes() and generateSeed() use $/\mbox{dev}/\mbox{urandom}$
Signature	NONEwithDSA
	SHA1withDSA
	SHA224withDSA
	SHA256withDSA
	NONEwithDSAinP1363Format
	SHA1withDSAinP1363Format
	SHA224withDSAinP1363Format
	SHA256withDSAinP1363Format

The following table lists OIDs associated with SHA Message Digests:

Table 4-5 OIDs associated with SHA Message Digests

SHA Message Digest	OID
SHA-224	2.16.840.1.101.3.4.2.4
SHA-256	2.16.840.1.101.3.4.2.1
SHA-384	2.16.840.1.101.3.4.2.2
SHA-512	2.16.840.1.101.3.4.2.3
SHA-512/224	2.16.840.1.101.3.4.2.5
SHA-512/256	2.16.840.1.101.3.4.2.6
SHA3-224	2.16.840.1.101.3.4.2.7
SHA3-256	2.16.840.1.101.3.4.2.8
	-



Table 4-5 (Cont.) OIDs associated with SHA Message Digests

SHA Message Digest	OID
SHA3-384	2.16.840.1.101.3.4.2.9
SHA3-512	2.16.840.1.101.3.4.2.10

The following table lists OIDs associated with DSA Signatures:

Table 4-6 OIDs associated with DSA Signatures

DSA Signature	OID
SHA1withDSA	1.2.840.10040.4.3
	1.3.14.3.2.13
	1.3.14.3.2.27
SHA224withDSA	2.16.840.1.101.3.4.3.1
SHA256withDSA	2.16.840.1.101.3.4.3.2

Keysize Restrictions

The SUN provider uses the following default keysizes (in bits) and enforces the following restrictions:

KeyPairGenerator

Alg. Name	Default Keysize	Restrictions/Comments
DSA	1024	Keysize must be a multiple of 64, ranging from 512 to 1024, plus 2048 and 3072.

AlgorithmParameterGenerator

Alg. Name	Default Keysize	Restrictions/Comments
DSA	1024	Keysize must be a multiple of 64, ranging from 512 to 1024, plus 2048 and 3072.

CertificateFactory/CertPathBuilder/CertPathValidator/CertStore Implementations

Additional details on the SUN provider implementations for CertificateFactory, CertPathBuilder, CertPathValidator and CertStore are documented in Appendix B: CertPath Implementation in SUN Provider of the PKI Programmer's Guide.

The SunRsaSign Provider

The SunRsaSign provider was introduced in JDK 1.3 as an enhanced replacement for the RSA signature in the SunJSSE provider.

The following algorithms are available in the SunRsaSign provider:



Table 4-7 The SunRsaSign Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
KeyFactory	RSA
KeyPairGenerator	RSA
Signature	MD2withRSA
	MD5withRSA
	SHA1withRSA
	SHA224withRSA
	SHA256withRSA
	SHA384withRSA
	SHA512withRSA

Keysize Restrictions

The SunRsaSign provider uses the following default keysize (in bits) and enforces the following restriction:

KeyPairGenerator

Table 4-8 The SunRsaSign Provider Keysize Restrictions

Alg. Name	Default Keysize	Restrictions/Comments
RSA	2048	Keysize must range between 512 and 65536 bits

The SunJSSE Provider

The Java Secure Socket Extension (JSSE) was originally released as a separate "Optional Package" (also briefly known as a "Standard Extension"), and was available for JDK 1.2.*n* and 1.3.*n*. The SunJSSE provider was introduced as part of this release.

In earlier JDK releases, there were no RSA signature providers available in the JDK, therefore <code>sunJsse</code> had to provide its own RSA implementation in order to use commonly available RSA-based certificates. JDK 5 introduced the <code>SunRsaSign</code> provider, which provides all the functionality (and more) of the <code>SunJsse</code> provider. Applications targeted at JDK 5.0 and later should request instances of the <code>SunRsaSign</code> provider instead. For backward compatibility, the RSA algorithms are still available through this provider, but are actually implemented in the <code>SunRsaSign</code> provider.

Algorithms

The following algorithms are available in the <code>sunJsse</code> provider:

Engine	Algorithm Name(s)
KeyFactory	RSA
	Note: The SunJSSE provider is for backwards compatibility with older releases, and should no longer be used for KeyFactory.



Engine	Algorithm Name(s)
KeyManagerFactory	PKIX: A factory for X509ExtendedKeyManager instances that manage X.509 certificate-based key pairs for local side authentication according to the rules defined by the IETF PKIX working group in RFC 5280. This KeyManagerFactory currently supports initialization using a KeyStore object or javax.net.ssl.KeyStoreBuilderParameters.
	SunX509: A factory for X509ExtendedKeyManager instances that manage X.509 certificate-based key pairs for local side authentication, but with less strict checking of certificate usage/validity and chain verification. This KeyManagerFactory supports initialization using a Keystore object, but does not currently support initialization using the class javax.net.ssl.ManagerFactoryParameters. Note: The SunX509 factory is for backwards
	compatibility with older releases, and should no longer be used.
KeyPairGenerator	RSA
	Note: The SunJSSE provider is for backwards compatibility with older releases, and should no longer be used for KeyPairGenerator.
KeyStore	PKCS12
	Note: The SunJSSE provider is for backwards compatibility with older releases, and should no longer be used for KeyStore.
Signature	MD2withRSA
	MD5withRSA
	SHA1withRSA
	Note: The SunJSSE provider is for backwards compatibility with older releases, and should no longer be used for Signature.
SSLContext	SSLv3
	TLSv1
	TLSv1.1
	TLSv1.2
	DTLSv1.0
	DTLSv1.2



Engine	Algorithm Name(s)
TrustManagerFactory	PKIX: A factory for X509ExtendedTrustManager instances that validate certificate chains according to the rules defined by the IETF PKIX working group in RFC 5280. This TrustManagerFactory currently supports initialization using a KeyStore object or javax.net.ssl.CertPathTrustManagerParame ters.
	SunX509: A factory for X509ExtendedTrustManager instances that validate certificate chains, but with less strict checking of certificate usage/validity and chain verification. This TrustManagerFactory supports initialization using a Keystore object, but does not currently support initialization using the class javax.net.ssl.ManagerFactoryParameters.
	Note: The SunX509 factory is for backwards compatibility with older releases, and should no longer be used.

SunJSSE Provider Protocol Parameters

The <code>SumJSSE</code> provider supports the following <code>protocol</code> parameters:

Table 4-9 SunJSSE Provider Protocol Parameters

Protocol	Enabled by Default for Client	Enabled by Default for Server
SSLv3	No (Unavailable [3])	No (Unavailable [3])
TLSv1	Yes	Yes
TLSv1.1	Yes	Yes
TLSv1.2	Yes	Yes
SSLv2Hello [1]	No	Yes
DTLSv1.0	Yes	Yes
DTLSv1.2 [2]	Yes	Yes

Footnote 1: The SSLv3, TLSv1, TLSv1.1 and TLSv1.2 protocols allow you to send SSLv3, TLSv1, TLSv1.1 and TLSv1.2 ClientHellos encapsulated in an SSLv2 format hello by using the SSLv2Hello psuedo-protocol.

Footnote 2: Both DTLSv1.0 and DTLSv1.2 are enabled.

Footnote 3: SSLv3 is enabled:

Starting with JDK 8u31, the SSLv3 protocol (Secure Socket Layer) has been deactivated and is not available by default. See the java.security.Security property jdk.tls.disabledAlgorithms in the <java_home>/conf/security/java.security file.



- If SSLv3 is absolutely required, the protocol can be reactivated by removing "SSLv3" from the jdk.tls.disabledAlgorithms property in the java.security file or by dynamically setting this Security Property before JSSE is initialized.
- To enable SSLv3 protocol at deploy level, after following the above steps, edit the deployment.properties file and add the following:

 deployment.security.SSLv3=true

The following table illustrates which connection combinations are possible when using SSLv2Hellos:

Table 4-10 Connections Possible Using SSLv2Hellos

Client	Server	Connection Possible?
Enabled	Enabled	Yes
Not enabled	Enabled	Yes (most interoperable: SunJSSE default)
Enabled	Not enabled	No
Not enabled	Not enabled	Yes

Cipher Suites

SunJSSE supports a large number of cipher suites. The tables Table 4-11 and Table 4-12 show the cipher suites supported by SunJSSE in preference order and the release in which they were introduced.

Table 4-11 lists the cipher suites that are enabled by default. Table 4-12 shows cipher suites that are supported by SunJSSE but not enabled by default.



According to DTLS Version 1.0 and DTLS Version 1.2, RC4 cipher suites must not be used with DTLS connections.

Cipher Suites That Are Enabled by Default

Table 4-11 Enabled Cipher Suites

Cipher Suite	JDK 6	JDK 7	JDK 8	JDK 9
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA 384		X [1]	Х	X
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA38 4		X [1]	Х	X
TLS_RSA_WITH_AES_256_CBC_SHA256		X [1]	Х	Х
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA3 84		X [1]	Х	Х
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384		X [1]	Х	Х
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256		X [1]	Х	Х
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	Х	Х	Х	Х
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	Х	Х	Х	Х



Table 4-11 (Cont.) Enabled Cipher Suites

			ı	1
Cipher Suite	JDK 6	JDK 7	JDK 8	JDK 9
TLS_RSA_WITH_AES_256_CBC_SHA	Х	Х	Х	Х
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA	Х	Х	Х	Х
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA	Х	Х	Х	Х
TLS_DHE_RSA_WITH_AES_256_CBC_SHA	Х	Х	Х	Х
TLS_DHE_DSS_WITH_AES_256_CBC_SHA	Х	Х	Х	Х
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA 256		X [1]	X	X
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA25 6		X [1]	Х	Х
TLS_RSA_WITH_AES_128_CBC_SHA256		X [1]	Х	Х
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA2 56		X [1]	Х	Х
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256		X [1]	Х	Х
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256		X [1]	Х	Х
TLS_DHE_DSS_WITH_AES_128_CBC_SHA256		X [1]	Х	Х
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	Х	Х	Х	Х
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	Х	Х	Х	Х
TLS_RSA_WITH_AES_128_CBC_SHA	Х	Х	Х	Х
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA	Х	Х	Х	Х
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA	Х	Х	Х	Х
TLS_DHE_RSA_WITH_AES_128_CBC_SHA	Х	Х	Х	Х
TLS_DHE_DSS_WITH_AES_128_CBC_SHA	Х	Х	Х	Х
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	Х	Х	Х	
TLS_ECDHE_RSA_WITH_RC4_128_SHA	Х	Х	Х	Х
SSL_RSA_WITH_RC4_128_SHA	Х	Х	Х	
TLS_ECDH_ECDSA_WITH_RC4_128_SHA	Х	Х	Х	Х
TLS_ECDH_RSA_WITH_RC4_128_SHA	Х	Х	Х	
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA 384			Х	Х
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA 256			Х	Х
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA38 4			Х	
TLS_RSA_WITH_AES_256_GCM_SHA384			Х	Х
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA3 84			Х	Х
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384			Х	Х
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384			Х	Х
TLS_DHE_DSS_WITH_AES_256_GCM_SHA384			Х	Х
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA25			X	Х
TLS_RSA_WITH_AES_128_GCM_SHA256			Х	х



Table 4-11 (Cont.) Enabled Cipher Suites

Cipher Suite	JDK 6	JDK 7	JDK 8	JDK 9
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA2 56			Х	Х
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256			Х	Х
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256			Х	Х
TLS_DHE_DSS_WITH_AES_128_GCM_SHA256			Х	Х
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SH A	X	Х	Х	Х
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	Х	Х	Х	Х
SSL_RSA_WITH_3DES_EDE_CBC_SHA	Х	Х	Х	Х
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA	Х	Х	Х	Х
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA	Х	Х	Х	Х
SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA	Х	Х	Х	Х
SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA	Х	Х	Х	Х
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA 384				Х
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256				Х
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256				Х
SSL_RSA_WITH_RC4_128_MD5	Х	Х	Х	
TLS_EMPTY_RENEGOTIATION_INFO_SCSV[2]	u22+	Х	Х	
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA 384				X
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256				Х
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256				Х

Footnote 1: Cipher suites with SHA384 and SHA256 are available only for TLS 1.2 or later.

Footnote 2: TLS_EMPTY_RENEGOTIATION_INFO_SCSV is a new pseudo-cipher suite to support RFC 5746. See Transport Layer Security (TLS) Renegotiation Issue in Java Secure Socket Extension (JSSE) Reference Guide.

Cipher Suites That Are Not Enabled by Default

Table 4-12 Not Enabled Cipher Suites

Cipher Suite	JDK 6	JDK 7	JDK 8	JDK 9
TLS_DH_anon_WITH_AES_256_GCM_SHA384			Х	Х
TLS_DH_anon_WITH_AES_128_GCM_SHA256			Х	Х
TLS_DH_anon_WITH_AES_256_CBC_SHA256		Х	Х	X
TLS_ECDH_anon_WITH_AES_256_CBC_SHA	Х	Х	Х	Х
TLS_DH_anon_WITH_AES_256_CBC_SHA	Х	Х	Х	Х
TLS_DH_anon_WITH_AES_128_CBC_SHA256		Х	Х	Х
TLS_ECDH_anon_WITH_AES_128_CBC_SHA	Х	Х	Х	Х
TLS_DH_anon_WITH_AES_128_CBC_SHA	Х	Х	Х	Х



Table 4-12 (Cont.) Not Enabled Cipher Suites

			_	
Cipher Suite	JDK 6	JDK 7	JDK 8	JDK 9
TLS_ECDH_anon_WITH_RC4_128_SHA	Х	Х	Х	Х
SSL_DH_anon_WITH_RC4_128_MD5	Х	Х	Х	Х
TLS_ECDH_anon_WITH_3DES_EDE_CBC_SHA	Х	Х	Х	Х
SSL_DH_anon_WITH_3DES_EDE_CBC_SHA	Х	Х	Х	Х
TLS_RSA_WITH_NULL_SHA256		Х	Х	Х
TLS_ECDHE_ECDSA_WITH_NULL_SHA	Х	Х	Х	Х
TLS_ECDHE_RSA_WITH_NULL_SHA	Х	Х	Х	Х
SSL_RSA_WITH_NULL_SHA	Х	Х	Х	Х
TLS_ECDH_ECDSA_WITH_NULL_SHA	Х	Х	Х	Х
TLS_ECDH_RSA_WITH_NULL_SHA	Х	Х	Х	Х
TLS_ECDH_anon_WITH_NULL_SHA	Х	Х	Х	Х
SSL_RSA_WITH_NULL_MD5	Х	Х	Х	Х
SSL_RSA_WITH_DES_CBC_SHA	Х	X [1]	Х	Х
SSL_DHE_RSA_WITH_DES_CBC_SHA	Х	X [1]	Х	Х
SSL_DHE_DSS_WITH_DES_CBC_SHA	Х	X [1]	Х	Х
SSL_DH_anon_WITH_DES_CBC_SHA	Х	X [1]	Х	Х
SSL_RSA_EXPORT_WITH_RC4_40_MD5	Х	X [2]	Х	Х
SSL_DH_anon_EXPORT_WITH_RC4_40_MD5	Х	X [2]	Х	Х
SSL_RSA_EXPORT_WITH_DES40_CBC_SHA	Х	X [2]	Х	Х
SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA	Х	X [2]	Х	Х
SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA	Х	X [2]	Х	Х
SSL_DH_anon_EXPORT_WITH_DES40_CBC_SHA	Х	X [2]	Х	Х
TLS_KRB5_WITH_RC4_128_SHA	Х	Х	Х	Х
TLS_KRB5_WITH_RC4_128_MD5	Х	Х	Х	Х
TLS_KRB5_WITH_3DES_EDE_CBC_SHA	Х	Х	Х	Х
TLS_KRB5_WITH_3DES_EDE_CBC_MD5	Х	Х	Х	Х
TLS_KRB5_WITH_DES_CBC_SHA	Х	X [1]	Х	Х
TLS_KRB5_WITH_DES_CBC_MD5	Х	X [1]	Х	Х
TLS_KRB5_EXPORT_WITH_RC4_40_SHA	Х	X [2]	Х	Х
TLS_KRB5_EXPORT_WITH_RC4_40_MD5	Х	X [2]	Х	Х
TLS_KRB5_EXPORT_WITH_DES_CBC_40_SHA	Х	X [2]	Х	Х
TLS_KRB5_EXPORT_WITH_DES_CBC_40_MD5	Х	X [2]	Х	Х
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA				Х
TLS_ECDHE_RSA_WITH_RC4_128_SHA				Х
SSL_RSA_WITH_RC4_128_MD5				Х

Footnote 1: RFC 5246 TLS 1.2 forbids the use of these suites. These can be used in the SSLv3/TLS1.0/TLS1.1 protocols, but cannot be used in TLS 1.2 and later.

Footnote 2: RFC 4346 TLS 1.1 forbids the use of these suites. These can be used in the SSLv3/TLS1.0 protocols, but cannot be used in TLS 1.1 and later.



Cipher suites that use AES_256 require the appropriate Java Cryptography Extension (JCE) unlimited strength jurisdiction policy file set, which is included in the JDK. By default, the active cryptography policy is unlimited. See Import Limits on Cryptographic Algorithms.

Cipher suites that use Elliptic Curve Cryptography (ECDSA, ECDH, ECDHE, ECDH_anon) require a JCE cryptographic provider that meets the following requirements:

- The provider must implement ECC as defined by the classes and interfaces in the packages java.security.spec and java.security.interfaces. The getAlgorithm() method of elliptic curve key objects must return the string "EC".
- The provider must support the Signature algorithms SHA1withECDSA and NONEwithECDSA, the KeyAgreement algorithm ECDH, and a KeyPairGenerator and a KeyFactory for algorithm EC. If one of these algorithms is missing, SunJSSE does not allow EC cipher suites to be used.
- The provider must support all the SECG curves referenced in RFC 4492 specification, section 5.1.1 (see also appendix A). In certificates, points should be encoded using the uncompressed form and curves should be encoded using the namedCurve choice, that is, using an object identifier.

If these requirements are not met, EC cipher suites may not be negotiated correctly.

Tighter Checking of EncryptedPreMasterSecret Version Numbers

Prior to the JDK 7 release, the SSL/TLS implementation did not check the version number in PreMasterSecret, and the SSL/TLS client did not send the correct version number by default. Unless the system property <code>com.sun.net.ssl.rsaPreMasterSecretFix</code> is set to <code>true</code>, the TLS client sends the active negotiated version, but not the expected maximum version supported by the client.

For compatibility, this behavior is preserved for SSL version 3.0 and TLS version 1.0. However, for TLS version 1.1 or later, the implementation tightens checking the PreMasterSecret version numbers as required by RFC 5246. Clients always send the correct version number, and servers check the version number strictly. The system property, com.sun.net.ssl.rsaPreMasterSecretFix, is not used in TLS 1.1 or later.

The SunJCE Provider

As described briefly in The SUN Provider, US export regulations at the time restricted the type of cryptographic functionality that could be available in the JDK. A separate API and reference implementation was developed that allowed applications to encrypt/decrypt date. The Java Cryptographic Extension (JCE) was released as a separate "Optional Package" (also briefly known as a "Standard Extension"), and was available for JDK 1.2x and 1.3x. During the development of JDK 1.4, regulations were relaxed enough that JCE (and SunJSSE) could be bundled as part of the JDK.

The following algorithms are available in the SunJCE provider:

Table 4-13 The SunJCE Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
AlgorithmParameterGenerator	DiffieHellman



Table 4-13 (Cont.) The SunJCE Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
AlgorithmParameters	AES
	Blowfish
	DES
	DESede
	DiffieHellman
	GCM
	OAEP
	PBE
	PBES2
	PBEWithHmacSHA1AndAES_128
	PBEWithHmacSHA224AndAES_128
	PBEWithHmacSHA256AndAES_128
	PBEWithHmacSHA384AndAES_128
	PBEWithHmacSHA512AndAES_128
	PBEWithHmacSHA1AndAES_256
	PBEWithHmacSHA224AndAES_256
	PBEWithHmacSHA256AndAES_256
	PBEWithHmacSHA384AndAES_256
	PBEWithHmacSHA512AndAES_256
	PBEWithMD5AndDES
	PBEWithMD5AndTripleDES
	PBEWithSHA1AndDESede
	PBEWithSHA1AndRC2_40
	PBEWithSHA1AndRC2_128
	PBEWithSHA1AndRC4_40
	PBEWithSHA1AndPC4_128
	RC2
Cipher	See Table 4-14
KeyAgreement	DiffieHellman
(eyFactory	DiffieHellman
KeyGenerator	AES
	ARCFOUR
	Blowfish
	DES
	DESede
	HmacMD5
	HmacMD5 HmacSHA1
	HmacSHA1
	HmacSHA1 HmacSHA224
	HmacSHA1 HmacSHA224 HmacSHA256
	HmacSHA1 HmacSHA224 HmacSHA256 HmacSHA384



Table 4-13 (Cont.) The SunJCE Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
KeyStore	JCEKS
Mac	HmacMD5
	HmacSHA1
	HmacSHA224
	HmacSHA384
	HmacSHA512
	HmacSHA256
	HmacSHA512/224
	HmacSHA512/256
	HmacPBESHA1
	PBEWithHmacSHA1
	PBEWithHmacSHA224
	PBEWithHmacSHA256
	PBEWithHmacSHA384
	PBEWithHmacSHA512
SecretKeyFactory	DES
	DESede
	PBEWithMD5AndDES
	PBEWithMD5AndTripleDES
	PBEWithSHA1AndDESede
	PBEWithSHA1AndRC2_40
	PBEWithSHA1AndRC2_128
	PBEWithSHA1AndRC4_40
	PBEWithSHA1AndRC4_128
	PBKDF2WithHmacSHA1
	PBKDF2WithHmacSHA224
	PBKDF2WithHmacSHA256
	PBKDF2WithHmacSHA384
	PBKDF2WithHmacSHA512
	PBEWithHmacSHA1AndAES_128
	PBEWithHmacSHA224AndAES_128
	PBEWithHmacSHA256AndAES_128
	PBEWithHmacSHA384AndAES_128
	PBEWithHmacSHA512AndAES_128
	PBEWithHmacSHA1AndAES_256
	PBEWithHmacSHA224AndAES_256
	PBEWithHmacSHA256AndAES_256
	PBEWithHmacSHA384AndAES_256
	PBEWithHmacSHA512AndAES 256

The following table lists cipher transformations available in the SunJCE provider.

Table 4-14 The SunJCE Provider Cipher Transformations

Algorithm Names	Modes	Paddings
AES	ECB, CBC, PCBC, CTR, CTS, CFB[1], CFB8CFB128, OFB[1], OFB8OFB128	NoPadding, PKCS5Padding, ISO10126Padding
AES	GCM	NoPadding
AESWrap	ECB	NoPadding
AESWrap_128	ECB	NoPadding
AESWrap_192	ECB	NoPadding
AESWrap_256	ECB	NoPadding
ARCFOUR	ECB	NoPadding
Blowfish, DES, DESede, RC2	ECB, CBC, PCBC, CTR, CTS, CFB[1], CFB8CFB64, OFB[1], OFB8OFB64	NoPadding, PKCS5Padding, ISO10126Padding
DESedeWrap	CBC	NoPadding
PBEWithMD5AndDES,	CBC	PKCS5Padding
PBEWithMD5AndTripleDES [2],		
PBEWithSHA1AndDESede,		
PBEWithSHA1AndRC2_40,		
PBEWithSHA1AndRC2_128,		
PBEWithSHA1AndRC4_40,		
PBEWithSHA1AndRC4_128,		
PBEWithHmacSHA1AndAES_ 128,		
PBEWithHmacSHA224AndAE S_128,		
PBEWithHmacSHA256AndAE S_128,		
PBEWithHmacSHA384AndAE S_128,		
PBEWithHmacSHA512AndAE S_128,		
PBEWithHmacSHA1AndAES_ 256,		
PBEWithHmacSHA224AndAE S_256,		
PBEWithHmacSHA256AndAE S_256,		
PBEWithHmacSHA384AndAE S_256,		
PBEWithHmacSHA512AndAE S_256		



Table 4-14 (Cont.) The SunJCE Provider Cipher Transformations

Algorithm Names	Modes	Paddings
RSA	ECB	NoPadding,
		PKCS1Padding,
		OAEPPadding
		OAEPWithMD5AndMGF1Pad ding,
		OAEPWithSHA1AndMGF1Pa dding,
		OAEPWithSHA-1AndMGF1Pa dding,
		OAEPWithSHA-224AndMGF1 Padding,
		OAEPWithSHA-256AndMGF1 Padding,
		OAEPWithSHA-384AndMGF1 Padding,
		OAEPWithSHA-512AndMGF1 Padding

Footnote 1: CFB/OFB with no specified value defaults to the block size of the algorithm. (i.e. AES is 128; Blowfish, DES, DESede, and RC2 are 64.)

Footnote 2: PBEWithMD5AndTripleDES is a proprietary algorithm that has not been standardized.

Keysize Restrictions

The SunJCE provider uses the following default key sizes (in bits) and enforces the following restrictions:

KeyGenerator

Table 4-15 The SunJCE Provider Key Size Restrictions

Algorithm Name	Default Key size	Restrictions/Comments
AES	128	Key size must be equal to 128, 192, or 256.
ARCFOUR (RC4)	128	Key size must range between 40 and 1024 (inclusive).
Blowfish	128	Key size must be a multiple of 8, ranging from 32 to 448 (inclusive).
DES	56	Key size must be equal to 56.



Table 4-15 (Cont.) The SunJCE Provider Key Size Restrictions

Algorithm Name	Default Key size	Restrictions/Comments
DESede (Triple DES)	168	Key size must be equal to 112 or 168.
		A key size of 112 will generate a Triple DES key with 2 intermediate keys, and a key size of 168 will generate a Triple DES key with 3 intermediate keys.
		Due to the "Meet-In-The-Middle" problem, even though 112 or 168 bits of key material are used, the effective key size is 80 or 112 bits respectively.
HmacMD5	512	No key size restriction.
HmacSHA1	512	No key size restriction.
HmacSHA224	224	No key size restriction.
HmacSHA256	256	No key size restriction.
HmacSHA384	384	No key size restriction.
HmacSHA512	512	No key size restriction.
RC2	128	Key size must range between 40 and 1024 (inclusive).



The various Password-Based Encryption (PBE) algorithms use various algorithms to generate key data, and ultimately depends on the targeted Cipher algorithm. For example,

"PBEWithMD5AndDES" will always generate 56-bit keys.

Table 4-16 KeyPairGenerator

Algorithm Name	Default Key size	Restrictions/Comments
Diffie-Hellman (DH)	2048	Key size must be a multiple of 64, ranging from 512 to 1024, plus 1536, 2048, 3072, 4096, 6144, 8192.



Table 4-17 AlgorithmParameterGenerator

Algorithm Name	Default Key size	Restrictions/Comments
Diffie-Hellman (DH)	2048	Key size must be a multiple of 64, ranging from 512 to 1024, plus 2048 and 3072.

The SunJGSS Provider

The following algorithms are available in the SunJGSS provider:

Table 4-18 The SunJGSS Provider

OID	Name
1.2.840.113554.1.2.2	Kerberos v5
1.3.6.1.5.5.2	SPNEGO

The SunSASL Provider

The following algorithms are available in the <code>sunsasl</code> provider:

Table 4-19 The SunSASL Provider Algorithm Names for Engine Classes

Algorithm Names
CRAM-MD5
DIGEST-MD5
EXTERNAL
NTLM
PLAIN
CRAM-MD5
DIGEST-MD5
NTLM

The XMLDSig Provider

The following algorithms are available in the XMLDSig provider:

Table 4-20 The XMLDSig Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
KeyInfoFactory	DOM



Table 4-20 (Cont.) The XMLDSig Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
TransformService	http://www.w3.org/TR/2001/REC-xml-c14n-20010315 - (CanonicalizationMethod.INCLUSIVE)
	http://www.w3.org/TR/2001/REC-xml- c14n-20010315#WithComments - (CanonicalizationMethod.INCLUSIVE_WITH_COMMEN TS)
	http://www.w3.org/2001/10/xml-exc-c14n# - (CanonicalizationMethod.EXCLUSIVE)
	http://www.w3.org/2001/10/xml-exc- c14n#WithComments - (CanonicalizationMethod.EXCLUSIVE_WITH_COMME NTS)
	http://www.w3.org/2000/09/xmldsig#base64 - (Transform.BASE64)
	http://www.w3.org/2000/09/xmldsig#enveloped-signature - (<i>Transform.ENVELOPED</i>)
	http://www.w3.org/TR/1999/REC-xpath-19991116 - (Transform.XPATH)
	http://www.w3.org/2002/06/xmldsig-filter2 - (<i>Transform.XPATH2</i>)
	http://www.w3.org/TR/1999/REC-xslt-19991116 - (Transform.XSLT)
XMLSignatureFactory	DOM

The SunPCSC Provider

The SunPCSC provider enables applications to use the Java Smart Card I/O API to interact with the PC/SC Smart Card stack of the underlying operating system. Consult your operating system documentation for details.

On Solaris and Linux platforms, SunPCSC accesses the PC/SC stack via the <code>libpcsclite.so</code> library. It looks for this library in the directories <code>/usr/\$LIBISA</code> and <code>/usr/local/\$LIBISA</code>, where <code>\$LIBISA</code> is expanded to <code>lib/64</code> on 64-bit Solaris platforms and <code>lib64</code> on 64-bit Linux platforms. The system property <code>sun.security.smartcardio.library</code> may also be set to the full filename of an alternate <code>libpcsclite.so</code> implementation. On Windows platforms, <code>SunPCSC</code> always calls into <code>winscard.dll</code> and no Java-level configuration is necessary or possible.

If PC/SC is available on the host platform, the SunPCSC implementation can be obtained via TerminalFactory.getDefault() and TerminalFactory.getInstance("PC/SC"). If PC/SC is not available or not correctly configured, a getInstance() call will fail with a NoSuchAlgorithmException and getDefault() will return a JRE built-in implementation that does not support any terminals.

The following algorithms are available in the SunPCSC provider:



Table 4-21 The SunPCSC Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
TerminalFactory	PC/SC

The SunMSCAPI Provider

The SunMSCAPI provider enables applications to use the standard JCA/JCE APIs to access the native cryptographic libraries, certificates stores and key containers on the Microsoft Windows platform. The SunMSCAPI provider itself does not contain cryptographic functionality, it is simply a conduit between the Java environment and the native cryptographic services on Windows.

The following algorithms are available in the Sunmscapi provider:

Table 4-22 The SunMSCAPI Algorithm Names for Engine Classes

Algorithm Names
RSA RSA/ECB/PKCS1Padding only
RSA
Windows-MY: The keystore type that identifies the native Microsoft Windows MY keystore. It contains the user's personal certificates and associated private keys.
Windows-ROOT : The keystore type that identifies the native Microsoft Windows ROOT keystore. It contains the certificates of Root certificate authorities and other self-signed trusted certificates.
Windows-PRNG: The name of the native pseudo-random number generation (PRNG) algorithm.
MD5withRSA
MD2withRSA
NONEwithRSA
SHA1withRSA
SHA256withRSA
SHA384withRSA
SHA512withRSA

Keysize Restrictions

The SunMSCAPI provider uses the following default keysizes (in bits) and enforce the following restrictions:

KeyGenerator



Table 4-23 The SunMSCAPI Provider Keysize Restrictions

Alg. Name	Default Keysize	Restrictions/Comments
Alg. Name	Delault Reysize	Restrictions/Comments
RSA	2048	Keysize ranges from 512 bits to 16,384 bits (depending on the underlying Microsoft Windows cryptographic service provider).

The SunEC Provider

The SunEC provider implements Elliptical Curve Cryptography (ECC). Compared to traditional cryptosystems such as RSA, ECC offers equivalent security with smaller key sizes, which results in faster computations, lower power consumption, and memory and bandwidth savings.

Applications can now use the standard JCA/JCE APIs to access ECC functionality without the dependency on external ECC libraries (via SunPKCS11), as was the case in the JDK 6 release.

The following algorithms are available in the SunEC provider:

Table 4-24 The SunEC Provider Names for Engine Classes

Engine	Algorithm Name(s)
AlgorithmParameters	EC
KeyAgreement	ECDH
KeyFactory	EC
KeyPairGenerator	EC
Signature	NONEwithECDSA
	SHA1withECDSA
	SHA224withECDSA
	SHA256withECDSA
	SHA384withECDSA
	SHA512withECDSA
	NONEwithECDSAinP1363Format
	SHA1withECDSAinP1363Format
	SHA224withECDSAinP1363Format
	SHA256withECDSAinP1363Format
	SHA384withECDSAinP1363Format
	SHA512withECDSAinP1363Format

Keysize Restrictions

The SunEC provider uses the following default keysizes (in bits) and enforces the following restrictions:

KeyPairGenerator



Table 4-25 The SunEC Provider Keysize Restrictions

Alg. Name	Default Keysize	Restrictions/Comments
EC	256	Keysize must range from 112 to 571 (inclusive).

The OracleUcrypto Provider

The Solaris-only security provider <code>OracleUcrypto</code> leverages the Solaris Ucrypto library to offload and delegate cryptographic operations supported by the Oracle SPARC T4 based on-core cryptographic instructions. The <code>OracleUcrypto</code> provider itself does not contain cryptographic functionality; it is simply a conduit between the Java environment and the Solaris Ucrypto library.

If the underlying Solaris Ucrypto library does not support a particular algorithm, then the <code>OracleUcrypto</code> provider will not support it either. Consequently, at runtime, the supported algorithms consists of the intersection of those that the Solaris Ucrypto library supports and those that the <code>OracleUcrypto</code> provider recognizes.

The following algorithms are available in the OracleUcrypto provider:

Table 4-26 The OracleUcrypto Provider Algorithm Names for Engine Classes

Engine	Algorithm Name(s)
Cipher	AES
	RSA
	AES/ECB/NoPadding
	AES/ECB/PKCS5Padding
	AES/CBC/NoPadding
	AES/CBC/PKCS5Padding
	AES/CTR/NoPadding
	AES/GCM/NoPadding
	AES/CFB128/NoPadding
	AES/CFB128/PKCS5Padding
	AES_128/ECB/NoPadding
	AES_192/ECB/NoPadding
	AES_256/ECB/NoPadding
	AES_128/CBC/NoPadding
	AES_192/CBC/NoPadding
	AES_256/CBC/NoPadding
	AES_128/GCM/NoPadding
	AES_192/GCM/NoPadding
	AES_256/GCM/NoPadding
	RSA/ECB/PKCS1Padding
	RSA/ECB/NoPadding



Table 4-26 (Cont.) The OracleUcrypto Provider Algorithm Names for Engine Classes

Engine	Algorithm Name(s)
Signature	MD5withRSA
	SHA1withRSA
	SHA256withRSA
	SHA384withRSA
	SHA512withRSA
MessageDigest	MD5
	SHA
	SHA-224
	SHA-256
	SHA-384
	SHA-512
	SHA3-224
	SHA3-256
	SHA3-384
	SHA3-512

Keysize Restrictions

The OracleUcrypto provider does not specify any default keysizes or keysize restrictions; these are specified by the underlying Solaris Ucrypto library.

OracleUcrypto Provider Configuration File

The <code>OracleUcrypto</code> provider has a configuration file named <code>ucrypto-solaris.cfg</code> that resides in the <code>\$JAVA_HOME/conf/security</code> directory. Modify this configuration file to specify which algorithms to disable by default. For example, the following configuration file disables AES with CFB128 mode by default:

```
#
# Configuration file for the OracleUcrypto provider
#
disabledServices = {
   Cipher.AES/CFB128/PKCS5Padding
   Cipher.AES/CFB128/NoPadding
}
```

The Apple Provider

The ${\tt Apple}$ provider implements a ${\tt java.security.KeyStore}$ that provides access to the macOS Keychain.

The following algorithms are available in the Apple provider:

Table 4-27 The Apple Provider Algorithm Name for Engine Classes

Engine	Algorithm Name(s)
KeyStore	KeychainStore

The JdkLDAP Provider

The JdkLDAP provider was introduced in JDK 9 as a replacement for the LDAP CertStore implementation in the SUN provider.

The following algorithms are available in the Jdkldap provider:

Table 4-28 The JdkLDAP Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
CertStore	LDAP

The JdkSASL Provider

The following algorithms are available in the Jdksasl provider:

Table 4-29 The JdkSASL Provider Algorithm Names for Engine Classes

Engine	Algorithm Names
SaslClient	GSSAPI
SaslServer	GSSAPI



5

PKCS#11 Reference Guide

The Java platform defines a set of programming interfaces for performing cryptographic operations. These interfaces are collectively known as the Java Cryptography Architecture (JCA) and the Java Cryptography Extension (JCE). See Java Cryptography Architecture (JCA) Reference Guide.

The cryptographic interfaces are provider-based. Specifically, applications talk to Application Programming Interfaces (APIs), and the actual cryptographic operations are performed in configured providers which adhere to a set of Service Provider Interfaces (SPIs). This architecture supports different provider implementations. Some providers may perform cryptographic operations in software; others may perform the operations on a hardware token (for example, on a smartcard device or on a hardware cryptographic accelerator).

The Cryptographic Token Interface Standard, PKCS#11, is produced by RSA Security and defines native programming interfaces to cryptographic tokens, such as hardware cryptographic accelerators and smartcards. Existing applications that use the JCA and JCE APIs can access native PKCS#11 tokens with the PKCS#11 provider. No modifications to the application are required. The only requirement is to properly configure the provider.

Although an application can make use of most PKCS#11 features using existing APIs, some applications might need more flexibility and capabilities. For example, an application might want to deal with smartcards being removed and inserted dynamically more easily. Or, a PKCS#11 token might require authentication for some non-key-related operations and therefore, the application must be able to log into the token without using keystore. The JCA gives applications greater flexibility in dealing with different providers.

This document describes how native PKCS#11 tokens can be configured into the Java platform for use by Java applications. It also describes how the JCA makes it easier for applications to deal with different types of providers, including PKCS#11 providers.

SunPKCS11 Provider

The SunPKCS11 provider, in contrast to most other providers, does not implement cryptographic algorithms itself. Instead, it acts as a bridge between the Java JCA and JCE APIs and the native PKCS#11 cryptographic API, translating the calls and conventions between the two.

This means that Java applications calling standard JCA and JCE APIs can, without modification, take advantage of algorithms offered by the underlying PKCS#11 implementations, such as, for example,

- Cryptographic smartcards,
- Hardware cryptographic accelerators, and
- High performance software implementations.





Java SE only facilitates accessing native PKCS#11 implementations, it does not itself include a native PKCS#11 implementation. However, cryptographic devices such as Smartcards and hardware accelerators often come with software that includes a PKCS#11 implementation, which you need to install and configure according to manufacturer's instructions.

SunPKCS11 Requirements

The SunPKCS11 provider requires an implementation of PKCS#11 v2.0 or later to be installed on the system. This implementation must take the form of a shared-object library (.so on Solaris and Linux) or dynamic-link library (.dll on Windows). Please consult your vendor documentation to find out if your cryptographic device includes such a PKCS#11 implementation, how to configure it, and what the name of the library file is.

The SunPKCS11 provider supports a number of algorithms, provided that the underlying PKCS#11 implementation offers them. The algorithms and their corresponding PKCS#11 mechanisms are listed in the table in SunPKCS11 Provider Supported Algorithms.

SunPKCS11 Configuration

The SunPKCS11 provider is in the module jdk.crypto.cryptoki. To use the provider, you must first install it statically or programmatically.

To install the provider statically, add the provider to the Java security properties file (java-home/conf/security/java.security).

For example, here's a fragment of the <code>java.security</code> file that installs the SunPKCS11 provider with the configuration file <code>/opt/bar/cfg/pkcs11.cfg</code>.

```
# configuration for security providers 1-12 ommitted
security.provider.13=SunPKCS11 /opt/bar/cfg/pkcs11.cfg
```

To install the provider dynamically, create an instance of the provider with the appropriate configuration filename and then install it. Here is an example.

```
String configName = "/opt/bar/cfg/pkcs11.cfg";
Provider p = Security.getProvider("SunPKCS11");
p = p.configure(configName);
Security.addProvider(p);
```

To use more than one slot per PKCS#11 implementation, or to use more than one PKCS#11 implementation, simply repeat the installation for each with the appropriate configuration file. This will result in a SunPKCS11 provider instance for each slot of each PKCS#11 implementation.

The configuration file is a text file that contains entries in the following format:

```
attribute=value
```

The valid values for *attribute* and *value* are described in the table in this section:

The two mandatory attributes are name and library.



Here is a sample configuration file:

name = FooAccelerator
library = /opt/foo/lib/libpkcs11.so

Comments are denoted by lines starting with the # (number) symbol.

Table 5-1 Attributes in the PKCS#11 Provider Configuration File

Attribute	Value	Description
library	pathname of PKCS#11 implementation	This is the full pathname (including extension) of the PKCS#11 implementation; the format of the pathname is platform dependent. For example, /opt/foo/lib/libpkcs11.so might be the pathname of a PKCS#11 implementation on Solaris and Linux while C:\foo \mypkcs11.dll might be the pathname on Windows.
name	name suffix of this provider instance	This string is concatenated with the prefix SunPKCS11- to produce this provider instance's name (that is, the string returned by its Provider.getName() method). For example, if the name attribute is "FooAccelerator", then the provider instance's name will be "SunPKCS11-FooAccelerator".
description	description of this provider instance	This string will be returned by the provider instance's Provider.getInfo() method. If none is specified, a default description will be returned.
slot	slot id	This is the id of the slot that this provider instance is to be associated with. For example, you would use 1 for the slot with the id 1 under PKCS#11. At most one of slot or slotListIndex may be specified. If neither is specified, the default is a slotListIndex of 0.



Table 5-1 (Cont.) Attributes in the PKCS#11 Provider Configuration File

Attribute	Value	Description
slotListIndex	slot index	This is the slot index that this provider instance is to be associated with. It is the index into the list of all slots returned by the PKCS#11 function C_GetSlotList. For example, 0 indicates the first slot in the list. At most one of slot or slotListIndex may be specified. If neither is specified, the default is a slotListIndex of 0.
enabledMechanisms	brace enclosed, whitespace- separated list of PKCS#11 mechanisms to enable	This is the list PKCS#11 mechanisms that this provider instance should use, provided that they are supported by both the SunPKCS11 provider and PKCS#11 token. All other mechanisms will be ignored. Each entry in the list is the name of a PKCS#11 mechanism. Here is an example that lists two PKCS#11 mechanisms.
		<pre>enabledMechanisms = { CKM_RSA_PKCS CKM_RSA_PKCS_KEY_PAIR_GEN } At most one of enabledMechanisms or disabledMechanisms may be specified. If neither is specified, the mechanisms enabled are those that are supported by both the</pre>
		SunPKCS11 provider (see SunPKCS11 Provider Supported Algorithms) and the PKCS#11 token.



Table 5-1 (Cont.) Attributes in the PKCS#11 Provider Configuration File

Attribute	Value	Description
disabledMechanisms	brace enclosed, whitespace- separated list of PKCS#11 mechanisms to disable	This is the list of PKCS#11 mechanisms that this provider instance should ignore. Any mechanism listed will be ignored by the provider, even if they are supported by the token and the SunPKCS11 provider. The strings SecureRandom and KeyStore may be specified to disable those services. At most one of enabledMechanisms or disabledMechanisms may be specified. If neither is specified, the mechanisms enabled are those that are supported by both the SunPKCS11 provider (see SunPKCS11 Provider Supported Algorithms) and the PKCS#11 token.
attributes	see below	The attributes option can be used to specify additional PKCS#11 that should be set when creating PKCS#11 key objects. This makes it possible to accomodate tokens that require particular attributes. For details, see the section below.

Attributes Configuration

The attributes option allows you to specify additional PKCS#11 attributes that should be set when creating PKCS#11 key objects. By default, the SunPKCS11 provider only specifies mandatory PKCS#11 attributes when creating objects. For example, for RSA public keys it specifies the key type and algorithm (CKA_CLASS and CKA_KEY_TYPE) and the key values for RSA public keys (CKA_MODULUS and CKA_PUBLIC_EXPONENT). The PKCS#11 library you are using will assign implementation specific default values to the other attributes of an RSA public key, for example that the key can be used to encrypt and verify messages (CKA_ENCRYPT and CKA_VERIFY = true).

The attributes option can be used if you do not like the default values your PKCS#11 implementation assigns or if your PKCS#11 implementation does not support defaults and requires a value to be specified explicitly. Note that specifying attributes that your PKCS#11 implementation does not support or that are invalid for the type of key in question may cause the operation to fail at runtime.

The option can be specified zero or more times, the options are processed in the order specified in the configuration file as described below. The attributes option has the format:



```
attributes(operation, keytype, keyalgorithm) = {
  name1 = value1
  [...]
}
```

Valid values for operation are:

- generate, for keys generated via a KeyPairGenerator or KeyGenerator
- import, for keys created via a KeyFactory or SecretKeyFactory. This also applies to Java software keys automatically converted to PKCS#11 key objects when they are passed to the initialization method of a cryptographic operation, for example Signature.initSign().
- *, for keys created using either a generate or a create operation.

Valid values for keytype are CKO_PUBLIC_KEY, CKO_PRIVATE_KEY, and CKO_SECRET_KEY, for public, private, and secret keys, respectively, and * to match any type of key.

Valid values for keyalgorithm are one of the CKK_XXX constants from the PKCS#11 specification, or * to match keys of any algorithm. The algorithms currently supported by the SunPKCS11 provider are CKK_RSA, CKK_DSA, CKK_DH, CKK_AES, CKK_DES, CKK_DES, CKK_DES, CKK_RC4, CKK_BLOWFISH, and CKK_GENERIC.

The attribute names and values are specified as a list of one or more name-value pairs. name must be a CKA_XXX constant from the PKCS#11 specification, for example CKA_SENSITIVE. value can be one of the following:

- A boolean value, true or false
- An integer, in decimal form (default) or in hexadecimal form if it begins with 0x.
- null, indicating that this attribute should not be specified when creating objects.

If the attributes option is specified multiple times, the entries are processed in the order specified with the attributes aggregated and later attributes overriding earlier ones. For example, consider the following configuration file excerpt:

```
attributes(*,CKO_PRIVATE_KEY,*) = {
   CKA_SIGN = true
}

attributes(*,CKO_PRIVATE_KEY,CKK_DH) = {
   CKA_SIGN = null
}

attributes(*,CKO_PRIVATE_KEY,CKK_RSA) = {
   CKA_DECRYPT = true
}
```

The first entry says to specify $cka_sign = true$ for all private keys. The second option overrides that with null for Diffie-Hellman private keys, so the cka_sign attribute will not specified for them at all. Finally, the third option says to also specify $cka_decrypt = true$ for RSA private keys. That means RSA private keys will have both $cka_sign = true$ and $cka_decrypt = true$ set.

There is also a special form of the attributes option. You can write attributes = compatibility in the configuration file. That is a shortcut for a whole set of attribute statements. They are designed to provider maximum compatibility with existing Java applications, which may expect, for example, all key components to be accessible and secret keys to be useable for both encryption and decryption. The compatibility



attributes line can be used together with other attributes lines, in which case the same aggregation and overriding rules apply as described earlier.

Accessing Network Security Services (NSS)

Network Security Services (NSS) is a set of open source security libraries whose crypto APIs are based on PKCS#11 but it includes special features that are outside of the PKCS#11 standard. The SunPKCS11 provider includes code to interact with these NSS specific features, including several NSS specific configuration directives, which are described below.

For best results, we recommend that you use the latest version of NSS available. It should be at least version 3.12.

The SunPKCS11 provider uses NSS specific code when any of the nss configuration directives described below are used. In that case, the regular configuration commands library, slot, and slotListIndex cannot be used.

Table 5-2 NSS Attributes and Values

Attribute	Value	Description
nssLibraryDirectory	directory containing the NSS and NSPR libraries	This is the full pathname of the directory containing the NSS and NSPR libraries. It must be specified unless NSS has already been loaded and initialized by another component running in the same process as the Java VM. Depending on your platform, you may have to set LD_LIBRARY_PATH or PATH (on Windows) to include this directory in order to allow the operating system to locate the dependent libraries.
nssSecmodDirectory	directory containing the NSS DB files	The full pathname of the directory containing the NSS configuration and key information (secmod.db, key3.db, and cert8.db). This directive must be specified unless NSS has already been initialized by another component (see above) or NSS is used without database files as described below.



Table 5-2 (Cont.) NSS Attributes and Values

Attribute	Value	Description
nssDbMode	one of readWrite, readOnly, and noDb	This directives determines how the NSS database is accessed. In read-write mode, full access is possible but only one process at a time should be accessing the databases. Read-only mode disallows modifications to the files. The noDb mode allows NSS to be used without database files purely as a cryptographic provider. It is not possible to create persistent keys using the PKCS11 KeyStore. This mode is useful because NSS includes highly optimized implementations and algorithms not currently available in Oracle's bundled Java-based crypto providers, for example Elliptic Curve Cryptography (ECC).



Table 5-2 (Cont.) NSS Attributes and Values

Attribute	Value	Description
nssModule	one of keystore, crypto, fips, and trustanchors	NSS makes its functionality available using several different libraries and slots. This directive determines which of these modules is accessed by this instance of SunPKCS11. The crypto module is the default in noDb mode. It supports crypto operations without login but no persistent keys.
		The fips module is the default if the NSS secmod.db has been set to FIPS-140 compliant mode. In this mode, NSS restricts the available algorithms and the PKCS#11 attributes with which keys can be created.
		The keystore module is the default in other configurations. It supports persistent keys using the PKCS11 KeyStore, which are stored in the NSS DB files. This module requires login.
		The trustanchors module enables access to NSS trust anchor certificates via the PKCS11 KeyStore, if secmod.db has been configured to include the trust anchor library.

Example 5-1 SunPKCS11 Configuration Files for NSS

NSS as a pure cryptography provider

name = NSScrypto
nssLibraryDirectory = /opt/tests/nss/lib
nssDbMode = noDb
attributes = compatibility

NSS as a FIPS 140 compliant crypto token

name = NSSfips
nssLibraryDirectory = /opt/tests/nss/lib
nssSecmodDirectory = /opt/tests/nss/fipsdb
nssModule = fips



Troubleshooting PKCS#11

There could be issues with PKCS#11 which requires debugging. To show debug info about Library, Slots, Token and Mechanism, add showInfo=true in < java-home > / conf/security/sunpkcs11-solaris.cfg file.

For additional debugging info, users can start or restart the Java processes with one of the following options:

For general SunPKCS11 provider debugging info:

```
-Djava.security.debug=sunpkcs11
```

For PKCS#11 keystore specific debugging info:

-Djava.security.debug=pkcs11keystore

Disabling PKCS#11 Providers and/or Individual PKCS#11 Mechanisms

As part of the troubleshooting process, it could be helpful to temporarily disable a PKCS#11 provider or the specific mechanism of a given provider.

Please note that disabling a PKCS#11 provider, is only a temporary measure. By disabling the PKCS#11 provider, the provider is no longer available which can cause applications to break or have a performance impact. Once the issue has been identified, only that specific mechanism should remain disabled.

Disabling PKCS#11 Providers

A PKCS#11 provider can be disabled by using one of the following methods:

 Disable PKCS#11 for a single Java process. Start or restart the Java process with the following Java command line flag:

```
-Dsun.security.pkcs11.enable-solaris=false
```



This step is only applicable to the SunPKCS11 provider when backed by the default Solaris PKCS#11 provider files (sun.security.pkcs11.SunPKCS11, /conf/security/sunpkcs11-solaris.cfg), and /conf/security/sunpkcs11-solaris.cfg).

2. Disable PKCS#11 for all Java processes run with a particular Java installation: This can be done dynamically by using the API (not shown in this section) or statically by editing the
/conf/security/java.security file and commenting out the SunPKCS11 security provider (do not forget to re-number the order of providers, if necessary) as shown below.

```
# List of providers and their preference orders (see above):
# security.provider.1=SUN
security.provider.2=SunRsaSign
security.provider.3=SunEC
```



```
security.provider.4=SunJSSE
security.provider.5=SunJCE
security.provider.6=SunJGSS
security.provider.7=SunSASL
security.provider.8=XMLDSig
security.provider.9=SunPCSC
security.provider.10=JdkLDAP
security.provider.11=JdkSASL
security.provider.12=SunMSCAPI
#security.provider.13=SunPKCS11
```

Start or restart the Java processes being run on this installation of Java.

Disabling Specific Mechanisms

When an issue occurs in one of the mechanisms of PKCS#11, it can be resolved by disabling only that particular mechanism, rather than the entire PKCS#11 provider (do not forget to re-enable the PKCS#11 provider if it was disabled earlier).

For example, to disable the SecureRandom mechanism only, you can add SecureRandom to the list of disabled mechanisms in the
// security/sunpkcs11-solaris.cfg file:

```
name = Solaris

description = SunPKCS11 accessing Solaris Cryptographic Framework

library = /usr/lib/$ISA/libpkcs11.so

handleStartupErrors = ignoreAll

# Use the X9.63 encoding for EC points (do not wrap in an ASN.1 OctetString).
useEcX963Encoding = true

attributes = compatibility

disabledMechanisms = {
   CKM_DSA_KEY_PAIR_GEN
   SecureRandom
}
```

Application Developers

Java applications can use the existing JCA and JCE APIs to access PKCS#11 tokens through the SunPKCS11 provider.

Token Login

You can login to the keystore using a Personal Identification Number and perform PKCS#11 operations.

Certain PKCS#11 operations, such as accessing private keys, require a login using a Personal Identification Number, or PIN, before the operations can proceed. The most common type of operations that require login are those that deal with keys on the token. In a Java application, such operations often involve first loading the keystore. When accessing the PKCS#11 token as a keystore via the <code>java.security.KeyStore</code> class, you can supply the PIN in the password input parameter to the <code>load</code> method.



The PIN will then be used by the SunPKCS11 provider for logging into the token. Here is an example.

```
char[] pin = ...;
KeyStore ks = KeyStore.getInstance("PKCS11");
ks.load(null, pin);
```

This is fine for an application that treats PKCS#11 tokens as static keystores. For an application that wants to accommodate PKCS#11 tokens more dynamically, such as smartcards being inserted and removed, you can use the new Keystore.Builder class. Here is an example of how to initialize the builder for a PKCS#11 keystore with a callback handler.

```
KeyStore.CallbackHandlerProtection chp =
   new KeyStore.CallbackHandlerProtection(new MyGuiCallbackHandler());
KeyStore.Builder builder =
   KeyStore.Builder.newInstance("PKCS11", null, chp);
```

For the SunPKCS11 provider, the callback handler must be able to satisfy a PasswordCallback, which is used to prompt the user for the PIN. Whenever the application needs access to the keystore, it uses the builder as follows.

```
KeyStore ks = builder.getKeyStore();
Key key = ks.getKey(alias, null);
```

The builder will prompt for a password as needed using the previously configured callback handler. The builder will prompt for a password only for the initial access. If the user of the application continues using the same Smartcard, the user will not be prompted again. If the user removes and inserts a different smartcard, the builder will prompt for a password for the new card.

Depending on the PKCS#11 token, there may be non-key-related operations that also require token login. Applications that use such operations can use the <code>java.security.AuthProvider</code> class. The <code>AuthProvider</code> class extends from <code>java.security.Provider</code> and defines methods to perform login and logout operations on a provider, as well as to set a callback handler for the provider to use.

For the SunPKCS11 provider, the callback handler must be able to satisfy a PasswordCallback, which is used to prompt the user for the PIN.

Here is an example of how an application might use an AuthProvider to log into the token.

```
AuthProvider aprov = (AuthProvider)Security.getProvider("SunPKCS11");
aprov.login(subject, new MyGuiCallbackHandler());
```

Token Keys

Java Key objects may or may not contain actual key material.

- A software Key object does contain the actual key material and allows access to that material.
- An unextractable key on a secure token (such as a smartcard) is represented by a
 Java Key object that does not contain the actual key material. The Key object only
 contains a reference to the actual key.

Applications and providers must use the correct interfaces to represent these different types of Key objects. Software Key objects (or any Key object that has access to the

actual key material) should implement the interfaces in the

java.security.interfaces and javax.crypto.interfaces packages (such as DSAPrivateKey). Key objects representing unextractable token keys should only implement the relevant generic interfaces in the java.security and javax.crypto packages (PrivateKey, PublicKey, Or SecretKey). Identification of the algorithm of a key should be performed using the Key.getAlgorithm() method.

Note that a Key object for an unextractable token key can only be used by the provider associated with that token.

Delayed Provider Selection

Java cryptography <code>getInstance()</code> methods, such as <code>Cipher.getInstance("AES")</code>, return the implementation from the first provider that implemented the requested algorithm. However, the JDK delays the selection of the provider until the relevant initialization method is called. The initialization method accepts a <code>Key</code> object and can determine at that point which provider can accept the specified <code>Key</code> object. This ensures that the selected provider can use the specified <code>Key</code> object. (If an application attempts to use a <code>Key</code> object for an unextractable token key with a provider that only accepts software key objects, then the provider throws an <code>InvalidKeyException</code>. This is an issue for the <code>Cipher</code>, <code>KeyAgreement</code>, <code>Mac</code>, and <code>Signature</code> classes.) The following represents the affected initialization methods.

```
    Cipher.init(..., Key key, ...)
    KeyAgreement.init(Key key, ...)
    Mac.init(Key key, ...)
    Signature.initSign(PrivateKey privateKey)
```

Furthermore, if an application calls the initialization method multiple times (each time with a different key, for example), the proper provider for the given key is selected each time. In other words, a different provider may be selected for each initialization call.

Although this delayed provider selection is hidden from the application, it does affect the behavior of the <code>getProvider()</code> method for <code>Cipher</code>, <code>KeyAgreement</code>, <code>Mac</code>, and <code>Signature</code>. If <code>getProvider()</code> is called before the initialization operation has occurred (and therefore before provider selection has occurred), then the first provider that supports the requested algorithm is returned. This may not be the same provider as the one selected after the initialization method is called. If <code>getProvider()</code> is called after the initialization operation has occurred, then the actual selected provider is returned. It is recommended that applications only call <code>getProvider()</code> after they have called the relevant initialization method.

In addition to getProvider(), the following additional methods are similarly affected.

- Cipher.getBlockSize
- Cipher.getExcemptionMechanism
- Cipher.getIV
- Cipher.getOutputSize
- Cipher.getParameters
- Mac.getMacLength
- Signature.getParameters



Signature.setParameter

JAAS KeyStoreLoginModule

The JDK comes with a JAAS keystore login module, <code>KeyStoreLoginModule</code>, that allows an application to authenticate using its identity in a specified keystore. After authentication, the application would acquire its principal and credentials information (certificate and private key) from the keystore. By using this login module and configuring it to use a PKCS#11 token as a keystore, the application can acquire this information from a PKCS#11 token.

Use the following options to configure the KeyStoreLoginModule to use a PKCS#11 token as the keystore.

- keyStoreURL="NONE"
- keyStoreType="PKCS11"
- keyStorePasswordURL=some_pin_url

where

some pin url

The location of the PIN. If the keyStorePasswordURL option is omitted, then the login module will get the PIN via the application's callback handler, supplying it with a PasswordCallback . Here is an example of a configuration file that uses a PKCS#11 token as a keystore.

```
other {
    com.sun.security.auth.module.KeyStoreLoginModule required
        keyStoreURL="NONE"
    keyStoreType="PKCS11"
    keyStorePasswordURL="file:/home/joe/scpin";
};
```

If more than one SunPKCS11 provider has been configured dynamically or in the <code>java.security</code> security properties file, you can use the <code>keyStoreProvider</code> option to target a specific provider instance. The argument to this option is the name of the provider. For the SunPKCS11 provider, the provider name is of the form <code>SunPKCS11-TokenName</code>, where <code>TokenName</code> is the name suffix that the provider instance has been configured with, as detailed in the <code>Table 5-1</code>. For example, the following configuration file names the PKCS#11 provider instance with name suffix <code>SmartCard</code>.

```
other {
    com.sun.security.auth.module.KeyStoreLoginModule required
        keyStoreURL="NONE"
    keyStoreType="PKCS11"
    keyStorePasswordURL="file:/home/joe/scpin"
    keyStoreProvider="SunPKCS11-SmartCard";
};
```

Some PKCS#11 tokens support login via a protected authentication path. For example, a smartcard may have a dedicated PIN-pad to enter the pin. Biometric devices will also have their own means to obtain authentication information. If the PKCS#11 token has a protected authentication path, then use the protected=true option and omit the keyStorePasswordURL option. Here is an example of a configuration file for such a token.



```
other {
    com.sun.security.auth.module.KeyStoreLoginModule required
        keyStoreURL="NONE"
    keyStoreType="PKCS11"
    protected=true;
};
```

Tokens as JSSE Keystore and Trust Stores

To use PKCS#11 tokens as JSSE keystores or trust stores, the JSSE application can use the APIs described in Token Login to instantiate a KeyStore that is backed by a PKCS#11 token and pass it to its key manager and trust manager. The JSSE application will then have access to the keys on the token.

JSSE also supports configuring the use of keystores and trust stores via system properties, as described in the Java Secure Socket Extension (JSSE) Reference Guide. To use a PKCS#11 token as a keystore or trust store, set the javax.net.ssl.keyStoreType and javax.net.ssl.trustStoreType system properties, respectively, to "PKCS11", and set the javax.net.ssl.keyStore and javax.net.ssl.trustStore system properties, respectively, to NONE. To specify the use of a specific provider instance, use the javax.net.ssl.keyStoreProvider and javax.net.ssl.trustStoreProvider system properties (for example, "SunPKCS11-SmartCard").

Using keytool and jarsigner with PKCS#11 Tokens

If the SunPKCS11 provider has been configured in the <code>java.security</code> security properties file (located in the <code>\$JAVA_HOME/conf/security</code> directory of the Java runtime), then keytool and jarsigner can be used to operate on the PKCS#11 token by specifying the following options.

- -keystore NONE
- -storetype PKCS11

Here an example of a command to list the contents of the configured PKCS#11 token.

```
keytool -keystore NONE -storetype PKCS11 -list
```

The PIN can be specified using the -storepass option. If none has been specified, then keytool and jarsigner will prompt for the token PIN. If the token has a protected authentication path (such as a dedicated PIN-pad or a biometric reader), then the -protected option must be specified, and no password options can be specified.

If more than one SunPKCS11 provider has been configured in the <code>java.security</code> security properties file, you can use the <code>-providerName</code> option to target a specific provider instance. The argument to this option is the name of the provider.

-providerName providerName

For the SunPKCS11 provider, providerName is of the form SunPKCS11-TokenName where:

TokenName

The name suffix that the provider instance has been configured with, as detailed in Table 5-1. For example, the following command lists the contents of the PKCS#11 keystore provider instance with name suffix SmartCard.



If the SunPKCS11 provider has not been configured in the <code>java.security</code> security properties file, you can use the following options to instruct <code>keytool</code> and <code>jarsigner</code> to install the provider dynamically.

- -providerClass sun.security.pkcs11.SunPKCS11
- -providerArg ConfigFilePath

ConfigFilePath

The path to the token configuration file. Here is an example of a command to list a PKCS#11 keystore when the SunPKCS11 provider has not been configured in the java.security file.

```
keytool -keystore NONE -storetype PKCS11 \
    -providerClass sun.security.pkcs11.SunPKCS11 \
    -providerArg /foo/bar/token.config \
    -list
```

Policy Tool



The Policy Tool is deprecated in JDK 9.

The keystore entry in the default policy implementation has the following syntax, which accommodates a PIN and multiple PKCS#11 provider instances:

```
keystore "some_keystore_url", "keystore_type", "keystore_provider";
keystorePasswordURL "some_password_url";
```

Where

keystore_provider

The keystore provider name (for example, "SunPKCS11-SmartCard").

some_password_url

A URL pointing to the location of the token PIN. Both *keystore_provider* and the *keystorePasswordURL* line are optional. If *keystore_provider* has not been specified, then the first configured provider that supports the specified keystore type is used. If the keystorePasswordURL line has not been specified, then no password is used.

Example 5-2 Keystore Policy Entry for a PKCS#11 Token

The following is an example keystore policy entry for a PKCS#11 token:

```
keystore "NONE", "PKCS11", "SunPKCS11-SmartCard";
keystorePasswordURL "file:/foo/bar/passwordFile";
```



Provider Developers

The <code>java.security.Provider</code> class enables provider developers to more easily support PKCS#11 tokens and cryptographic services through provider services and parameter support.

See Example Provider for an example of a simple provider designed to demonstrate provider services and parameter support.

Provider Services

For each service implemented by the provider, there must be a property whose name is the type of service (Cipher, Signature, etc), followed by a period and the name of the algorithm to which the service applies. The property value must specify the fully qualified name of the class implementing the service. Here is an example of a provider setting KeyAgreement.DiffieHellman property to have the value com.sun.crypto.provider.DHKeyAgreement.

```
put("KeyAgreement.DiffieHellman", "com.sun.crypto.provider.DHKeyAgreement")
```

The public static nested class Provider.Service encapsulates the properties of a provider service (including its type, attributes, algorithm name, and algorithm aliases). Providers can instantiate Provider.Service objects and register them by calling the Provider.putService() method. This is equivalent to creating a Property entry and calling the Provider.put() method. Note that legacy Property entries registered via Provider.put are still supported.

Here is an example of a provider creating a Service object with the KeyAgreement type, for the DiffieHellman algorithm, implemented by the class com.sun.crypto.provider.DHKeyAgreement.

Using Provider. Servicee objects instead of legacy Property entries has a couple of major benefits. One benefit is that it allows the provider to have greater flexibility when Instantiating Engine Classes. Another benefit is that it allows the provider to test Parameter Support. These features are discussed in detail next.

Instantiating Engine Classes

putService(s);

By default, the Java Cryptography framework looks up the provider property for a particular service and directly instantiates the engine class registered for that property. A provider can to override this behavior and instantiate the engine class for the requested service itself.

To override the default behavior, the provider overrides the Provider.Service.newInstance() method to add its custom behavior. For example, the provider might call a custom constructor, or might perform initialization using information not accessible outside the provider (or that are only known by the provider).



Parameter Support

The Java Cryptography framework may attempt a fast check to determine whether a provider's service implementation can use an application-specified parameter. To perform this fast check, the framework calls Provider. Service.supportsParameter().

The framework relies on this fast test during delayed provider selection (see Delayed Provider Selection). When an application invokes an initialization method and passes it a <code>Key</code> object, the framework asks an underlying provider whether it supports the object by calling its <code>Service.supportsParameter()</code> method. If <code>supportsParameter()</code> returns <code>false</code>, the framework can immediately remove that provider from consideration. If <code>supportsParameter()</code> returns <code>true</code>, the framework passes the <code>Key</code> object to that provider's initialization engine class implementation. A provider that requires software <code>Key</code> objects should override this method to return <code>false</code> when it is passed non-software <code>keys</code>. Likewise, a provider for a PKCS#11 token that contains unextractable keys should only return <code>true</code> for <code>Key</code> objects that it created, and which therefore correspond to the keys on its respective token.



The default implementation of supportsParameter() returns true. This allows existing providers to work without modification. However, because of this lenient default implementation, the framework must be prepared to catch exceptions thrown by providers that reject the κey object inside their initialization engine class implementations. The framework treats these cases the same as when supportsParameter() returns false.

SunPKCS11 Provider Supported Algorithms

Table 5-3 lists the Java algorithms supported by the SunPKCS11 provider and corresponding PKCS#11 mechanisms needed to support them. When multiple mechanisms are listed, they are given in the order of preference and any one of them is sufficient.

Note:

SunPKCS11 can be instructed to ignore mechanisms by using the disabledMechanisms and enabledMechanisms configuration directives (see SunPKCS11 Configuration).

For Elliptic Curve mechanisms, the SunPKCS11 provider will only use keys that use the namedCurve choice as encoding for the parameters and only allow the uncompressed point format. The SunPKCS11 provider assumes that a token supports all standard named domain parameters.



Table 5-3 Java Agorithms Supported by the SunPKCS11 Provider

Java Algorithm	PKCS#11 Mechanisms
Signature.MD2withRSA	CKM_MD2_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.MD5withRSA	CKM_MD5_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA1withRSA	CKM_SHA1_RSA_PKCS, CKM_RSA_PKCS CKM_RSA_X_509
Signature.SHA224withRSA	CKM_SHA224_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA256withRSA	CKM_SHA256_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA384withRSA	CKM_SHA384_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA512withRSA	CKM_SHA512_RSA_PKCS, CKM_RSA_PKCS, CKM_RSA_X_509
Signature.SHA1withDSA	CKM_DSA_SHA1, CKM_DSA
Signature.NONEwithDSA	CKM_DSA
Signature.SHA1withECDSA	CKM_ECDSA_SHA1, CKM_ECDSA
Signature.SHA224withECDSA	CKM_ECDSA
Signature.SHA256withECDSA	CKM_ECDSA
Signature.SHA384withECDSA	CKM_ECDSA
Signature.SHA512withECDSA	CKM_ECDSA
Signature.NONEwithECDSA	CKM_ECDSA
Cipher.RSA/ECB/PKCS1Padding	CKM_RSA_PKCS
Cipher.ARCFOUR	CKM_RC4
Cipher.DES/CBC/NoPadding	CKM_DES_CBC
Cipher.DESede/CBC/NoPadding	CKM_DES3_CBC
Cipher.AES/CBC/NoPadding	CKM_AES_CBC
Cipher.Blowfish/CBC/NoPadding	CKM_BLOWFISH_CBC
Cipher.RSA/ECB/NoPadding	CKM_RSA_X_509
Cipher.AES/CTR/NoPadding	CKM_AES_CTR
KeyAgreement.ECDH	CKM_ECDH1_DERIVE
KeyAgreement.DiffieHellman	CKM_DH_PKCS_DERIVE
KeyPairGenerator.RSA	CKM_RSA_PKCS_KEY_PAIR_GEN
KeyPairGenerator.DSA	CKM_DSA_KEY_PAIR_GEN
KeyPairGenerator.EC	CKM_EC_KEY_PAIR_GEN
KeyPairGenerator.DiffieHellman	CKM_DH_PKCS_KEY_PAIR_GEN
KeyGenerator.ARCFOUR	CKM_RC4_KEY_GEN
KeyGenerator.DES	CKM_DES_KEY_GEN
KeyGenerator.DESede	CKM_DES3_KEY_GEN
KeyGenerator.AES	CKM_AES_KEY_GEN
KeyGenerator.Blowfish	CKM_BLOWFISH_KEY_GEN
Mac.HmacMD5	CKM_MD5_HMAC
Mac.HmacSHA1	CKM_SHA_1_HMAC
Mac.HmacSHA224	CKM_SHA224_HMAC



Table 5-3 (Cont.) Java Agorithms Supported by the SunPKCS11 Provider

Java Algorithm	PKCS#11 Mechanisms
Mac.HmacSHA256	CKM_SHA256_HMAC
Mac.HmacSHA384	CKM_SHA384_HMAC
Mac.HmacSHA512	CKM_SHA512_HMAC
MessageDigest.MD2	CKM_MD2
MessageDigest.MD5	CKM_MD5
MessageDigest.SHA1	CKM_SHA_1
MessageDigest.SHA-224	CKM_SHA224
MessageDigest.SHA-256	CKM_SHA256
MessageDigest.SHA-384	CKM_SHA384
MessageDigest.SHA-512	CKM_SHA512
KeyFactory.RSA	Any supported RSA mechanism
KeyFactory.DSA	Any supported DSA mechanism
KeyFactory.EC	Any supported EC mechanism
KeyFactory.DiffieHellman	Any supported Diffie-Hellman mechanism
SecretKeyFactory.ARCFOUR	CKM_RC4
SecretKeyFactory.DES	CKM_DES_CBC
SecretKeyFactory.DESede	CKM_DES3_CBC
SecretKeyFactory.AES	CKM_AES_CBC
SecretKeyFactory.Blowfish	CKM_BLOWFISH_CBC
SecureRandom.PKCS11	CK_TOKEN_INFO has the CKF_RNG bit set
KeyStore.PKCS11	Always available
· · · · · · · · · · · · · · · · · · ·	

SunPKCS11 Provider KeyStore Requirements

The following describes the requirements placed by the SunPKCS11 provider's KeyStore implementation on the underlying native PKCS#11 library.



Changes may be made in future releases to maximize interoperability with as many existing PKCS#11 libraries as possible.

Read-Only Access

To map existing objects stored on a PKCS#11 token to KeyStore entries, the SunPKCS11 provider's KeyStore implementation performs the following operations.

- 1. A search for all private key objects on the token is performed by calling C_FindObjects[Init|Final]. The search template includes the following attributes:
 - CKA TOKEN = true
 - CKA_CLASS = CKO_PRIVATE_KEY



- 2. A search for all certificate objects on the token is performed by calling C_FindObjects[Init|Final]. The search template includes the following attributes:
 - CKA TOKEN = true
 - CKA_CLASS = CKO_CERTIFICATE
- 3. Each private key object is matched with its corresponding certificate by retrieving their respective CKA_ID attributes. A matching pair must share the same unique CKA_ID.

For each matching pair, the certificate chain is built by following the issuer->subject path. From the end entity certificate, a call for <code>c_FindObjects[Init|Final]</code> is made with a search template that includes the following attributes:

- CKA TOKEN = true
- CKA_CLASS = CKO_CERTIFICATE
- CKA_SUBJECT = [DN of certificate issuer]

This search is continued until either no certificate for the issuer is found, or until a self-signed certificate is found. If more than one certificate is found the first one is used.

Once a private key and certificate have been matched (and its certificate chain built), the information is stored in a private key entry with the CKA_LABEL value from end entity certificate as the KeyStore alias.

If the end entity certificate has no CKA_LABEL, then the alias is derived from the CKA_ID. If the CKA_ID can be determined to consist exclusively of printable characters, then a String alias is created by decoding the CKA_ID bytes using the UTF-8 charset. Otherwise, a hex String alias is created from the CKA_ID bytes ("0xFFFF...", for example).

If multiple certificates share the same CKA_LABEL, then the alias is derived from the CKA_LABEL plus the end entity certificate issuer and serial number ("MyCert/CN=foobar/1234", for example).

- 4. Each certificate not part of a private key entry (as the end entity certificate) is checked whether it is trusted. If the CKA_TRUSTED attribute is true, then a KeyStore trusted certificate entry is created with the CKA_LABEL value as the KeyStore alias. If the certificate has no CKA_LABEL, or if multiple certificates share the same CKA_LABEL, then the alias is derived as described above. If the CKA_TRUSTED attribute is not supported then no trusted certificate entries are created.
- 5. Any private key or certificate object not part of a private key entry or trusted certificate entry is ignored.
- 6. A search for all secret key objects on the token is performed by calling C_FindObjects[Init|Final]. The search template includes the following attributes:
 - CKA TOKEN = true
 - CKA_CLASS = CKO_SECRET_KEY

A KeyStore secret key entry is created for each secret key object, with the CKA_LABEL value as the KeyStore alias. Each secret key object must have a unique CKA_LABEL.



Write Access

To create new KeyStore entries on a PKCS#11 token to KeyStore entries, the SunPKCS11 provider's KeyStore implementation performs the following operations.

 When creating a KeyStore entry (during KeyStore.setEntry, for example), C_CreateObject is called with CKA_TOKEN=true to create token objects for the respective entry contents.

Private key objects are stored with CKA_PRIVATE=true. The KeyStore alias (UTF8-encoded) is set as the CKA_ID for both the private key and the corresponding end entity certificate. The KeyStore alias is also set as the CKA_LABEL for the end entity certificate object.

Each certificate in a private key entry's chain is also stored. The CKA_LABEL is not set for CA certificates. If a CA certificate is already in the token, a duplicate is not stored.

Secret key objects are stored with $\tt CKA_PRIVATE=true$. The KeyStore alias is set as the CKA LABEL.

- 2. If an attempt is made to convert a session object to a token object (for example, if KeyStore.setEntry is called and the private key object in the specified entry is a session ojbect), then C CopyObject is called with <code>cka_Token=true</code>.
- 3. If multiple certificates in the token are found to share the same CKA_LABEL, then the write capabilities to the token are disabled.
- 4. Since the PKCS#11 specification does not allow regular applications to set CKA_TRUSTED=true (only token initialization applications may do so), trusted certificate entries can not be created.

Miscellaneous

In addition to the searches listed above, the following searches may be used by the SunPKCS11 provider's KeyStore implementation to perform internal functions. Specifically, $C_{\texttt{FindObjects[Init|Final]}}$ may be called with any of the following attribute templates:

```
CKA_TOKEN true
CKA_CLASS CKO_CERTIFICATE
CKA_SUBJECT [subject DN]

CKA_TOKEN true
CKA_CLASS CKO_SECRET_KEY
CKA_LABEL [label]

CKA_TOKEN true
CKA_CLASS CKO_CERTIFICATE or CKO_PRIVATE_KEY
CKA_LABEL [cka_id]
```

Example Provider

The following is an example of a simple provider that demonstrates features of the Provider class.

```
package com.foo;
import java.io.*;
import java.lang.reflect.*;
```



```
import java.security.*;
import javax.crypto.*;
 ^{\star} Example provider that demonstrates some Provider class features.
   . implement multiple different algorithms in a single class.
     Previously each algorithm needed to be implemented in a separate class
      (e.g. one for SHA-256, one for SHA-384, etc.)
   . multiple concurrent instances of the provider frontend class each
     associated with a different backend.
    . it uses "unextractable" keys and lets the framework know which key
     objects it can and cannot support
 * Note that this is only a simple example provider designed to demonstrate
 * several of the new features. It is not explicitly designed for efficiency.
public final class ExampleProvider extends Provider {
    // reference to the crypto backend that implements all the algorithms
    final CryptoBackend cryptoBackend;
    public ExampleProvider(String name, CryptoBackend cryptoBackend) {
        super(name, 1.0, "JCA/JCE provider for " + name);
        this.cryptoBackend = cryptoBackend;
        // register the algorithms we support (SHA-256, SHA-384, DESede, and AES)
       putService(new MyService
            (this, "MessageDigest", "SHA-256",
"com.foo.ExampleProvider$MyMessageDigest"));
        putService(new MyService
            (this, "MessageDigest", "SHA-384",
"com.foo.ExampleProvider$MyMessageDigest"));
       putService(new MyCipherService
            (this, "Cipher", "DES", "com.foo.ExampleProvider$MyCipher"));
       putService(new MyCipherService
            (this, "Cipher", "AES", "com.foo.ExampleProvider$MyCipher"));
    // the API of our fictitious crypto backend
    static abstract class CryptoBackend {
        abstract byte[] digest(String algorithm, byte[] data);
       abstract byte[] encrypt(String algorithm, KeyHandle key, byte[] data);
       abstract byte[] decrypt(String algorithm, KeyHandle key, byte[] data);
       abstract KeyHandle createKey(String algorithm, byte[] keyData);
    // the shell of the representation the crypto backend uses for keys
   private static final class KeyHandle {
        // fill in code
    // we have our own ServiceDescription implementation that overrides newInstance()
    // that calls the (Provider, String) constructor instead of the no-args
constructor
   private static class MyService extends Service {
       private static final Class[] paramTypes = {Provider.class, String.class};
       MyService(Provider provider, String type, String algorithm,
```

```
String className) {
            super(provider, type, algorithm, className, null, null);
        }
        public Object newInstance(Object param) throws NoSuchAlgorithmException {
            try {
                // get the Class object for the implementation class
                Class clazz;
                Provider provider = getProvider();
                ClassLoader loader = provider.getClass().getClassLoader();
                if (loader == null) {
                    clazz = Class.forName(getClassName());
                } else {
                    clazz = loader.loadClass(getClassName());
                // fetch the (Provider, String) constructor
                Constructor cons = clazz.getConstructor(paramTypes);
                // invoke constructor and return the SPI object
                Object obj = cons.newInstance(new Object[] {provider,
getAlgorithm()});
                return obj;
            } catch (Exception e) {
                throw new NoSuchAlgorithmException("Could not instantiate service",
e);
    // custom ServiceDescription class for Cipher objects. See supportsParameter()
below
    private static class MyCipherService extends MyService {
        MyCipherService(Provider provider, String type, String algorithm,
                String className) {
            super(provider, type, algorithm, className);
        // we override supportsParameter() to let the framework know which
        // keys we can support. We support instances of MySecretKey, if they
        // are stored in our provider backend, plus SecretKeys with a RAW encoding.
        public boolean supportsParameter(Object obj) {
            if (obj instanceof SecretKey == false) {
                return false;
            SecretKey key = (SecretKey)obj;
            if (key.getAlgorithm().equals(getAlgorithm()) == false) {
                return false;
            if (key instanceof MySecretKey) {
                MySecretKey myKey = (MySecretKey)key;
                return myKey.provider == getProvider();
            } else {
                return "RAW".equals(key.getFormat());
        }
    }
    // our generic MessageDigest implementation. It implements all digest
    \ensuremath{//} algorithms in a single class. We only implement the bare minimum
    // of MessageDigestSpi methods
    private static final class MyMessageDigest extends MessageDigestSpi {
        private final ExampleProvider provider;
        private final String algorithm;
```

```
private ByteArrayOutputStream buffer;
       MyMessageDigest(Provider provider, String algorithm) {
            super();
            this.provider = (ExampleProvider)provider;
            this.algorithm = algorithm;
            engineReset();
       protected void engineReset() {
            buffer = new ByteArrayOutputStream();
       protected void engineUpdate(byte b) {
            buffer.write(b);
       protected void engineUpdate(byte[] b, int ofs, int len) {
            buffer.write(b, ofs, len);
       protected byte[] engineDigest() {
            byte[] data = buffer.toByteArray();
            byte[] digest = provider.cryptoBackend.digest(algorithm, data);
            engineReset();
            return digest;
    // our generic Cipher implementation, only partially complete. It implements
    // all cipher algorithms in a single class. We implement only as many of the
    // CipherSpi methods as required to show how it could work
    private static abstract class MyCipher extends CipherSpi {
       private final ExampleProvider provider;
       private final String algorithm;
       private int opmode;
       private MySecretKey myKey;
       private ByteArrayOutputStream buffer;
       MyCipher(Provider provider, String algorithm) {
            super();
            this.provider = (ExampleProvider)provider;
            this.algorithm = algorithm;
       protected void engineInit(int opmode, Key key, SecureRandom random)
                throws InvalidKeyException {
            this.opmode = opmode;
            myKey = MySecretKey.getKey(provider, algorithm, key);
            if (myKey == null) {
                throw new InvalidKeyException();
            buffer = new ByteArrayOutputStream();
       protected byte[] engineUpdate(byte[] b, int ofs, int len) {
            buffer.write(b, ofs, len);
            return new byte[0];
       protected int engineUpdate(byte[] b, int ofs, int len, byte[] out, int
outOfs) {
            buffer.write(b, ofs, len);
            return 0;
        protected byte[] engineDoFinal(byte[] b, int ofs, int len) {
            buffer.write(b, ofs, len);
            byte[] in = buffer.toByteArray();
            byte[] out;
            if (opmode == Cipher.ENCRYPT_MODE) {
```

```
out = provider.cryptoBackend.encrypt(algorithm, myKey.handle, in);
            } else {
                out = provider.cryptoBackend.decrypt(algorithm, myKey.handle, in);
            buffer = new ByteArrayOutputStream();
            return out;
        // code for remaining CipherSpi methods goes here
    // our SecretKey implementation. All our keys are stored in our crypto
    // backend, we only have an opaque handle available. There is no
    // encoded form of these keys.
    private static final class MySecretKey implements SecretKey {
        final String algorithm;
        final Provider provider;
        final KeyHandle handle;
        MySecretKey(Provider provider, String algorithm, KeyHandle handle) {
            this.provider = provider;
            this.algorithm = algorithm;
            this.handle = handle;
        public String getAlgorithm() {
            return algorithm;
        public String getFormat() {
            return null; // this key has no encoded form
        public byte[] getEncoded() {
            return null; // this key has no encoded form
        // Convert the given key to a key of the specified provider, if possible
        static MySecretKey getKey(ExampleProvider provider, String algorithm, Key
key) {
            if (key instanceof SecretKey == false) {
                return null;
            // algorithm name must match
            if (!key.getAlgorithm().equals(algorithm)) {
                return null;
            // if key is already an instance of MySecretKey and is stored
            // on this provider, return it right away
            if (key instanceof MySecretKey) {
                MySecretKey myKey = (MySecretKey)key;
                if (myKey.provider == provider) {
                    return myKey;
            // otherwise, if the input key has a RAW encoding, convert it
            if (!"RAW".equals(key.getFormat())) {
                return null;
            byte[] encoded = key.getEncoded();
            KeyHandle handle = provider.cryptoBackend.createKey(algorithm, encoded);
            return new MySecretKey(provider, algorithm, handle);
```

}



6

Java Authentication and Authorization Service (JAAS)

JAAS Reference Guide in the JDK 8 documentation describes Java Authentication and Authorization Service (JAAS), which enables you to authenticate users and securely determine who is currently executing Java code, and authorize users to ensure that they have the access control rights, or permissions, required to do the actions performed.

JAAS Tutorials in the JDK 8 documentation provides tutorials about Java Authentication and Authorization Service (JAAS) authentication and authorization.

Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide shows you how to implement the LoginModule interface, which you plug into an application to provide a particular type of authentication.

JAAS Reference Guide

See JAAS Reference Guide in the JDK 8 documentation for information about Java Authentication and Authorization Service (JAAS), which enables you to do the following:

- Authenticate users and securely determine who is currently executing Java code, regardless of whether the code is running as an application, an applet, a bean, or a servlet.
- Authorize users to ensure that they have the access control rights (permissions) required to do the actions performed.

JAAS Tutorials

See JAAS Tutorials in the JDK 8 documentation for tutorials about Java Authentication and Authorization Service (JAAS) authentication and authorization.

Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide

The Java Authentication and Authorization Service (JAAS) was introduced as an optional package to the Java 2 SDK, Standard Edition (J2SDK), v 1.3. JAAS was integrated into the Java Standard Edition Development Kit starting with J2SDK 1.4.

JAAS provides subject-based authorization on authenticated identities. This document focuses on the authentication aspect of JAAS, specifically the Interface LoginModule .

Who Should Read This Document

This document is intended for experienced programmers who require the ability to write a Interface LoginModule implementing an authentication technology.

Related Documentation

This document assumes you have already read the following:

Java Authentication and Authorization Service (JAAS) Reference Guide

It also discusses various classes and interfaces in the JAAS API. See the Javadoc API documentation for the JAAS API specification for more detailed information:

- javax.security.auth
- com.sun.security.auth.callback
- javax.security.auth.kerberos
- com.sun.security.auth.login
- javax.security.auth.spi
- javax.security.auth.x500
- com.sun.security.auth
- com.sun.security.auth.callback
- com.sun.security.auth.login
- com.sun.security.auth.module

The following **tutorials** for JAAS authentication and authorization can be run by everyone:

- JAAS Authentication Tutorial
- JAAS Authorization Tutorial

Similar tutorials for JAAS authentication and authorization, but which demonstrate the use of a Kerberos LoginModule and thus which require a Kerberos installation, can be found at

- JAAS Authentication
- JAAS Authorization

These two tutorials are a part of the JAAS and Java GSS-API Tutorial that utilize Kerberos as the underlying technology for authentication and secure communication.

Introduction to LoginModule

LoginModules are plugged in under applications to provide a particular type of authentication.

The Interface LoginModule documentation describes the interface that must be implemented by authentication technology providers.

While applications write to the LoginContext Application Programming Interface (API), authentication technology providers implement the LoginModule interface. A Configuration Specifies the LoginModule(s) to be used with a particular login application.



Different LoginModules can be plugged in under the application without requiring any modifications to the application itself.

The LoginContext is responsible for reading the Configuration and instantiating the specified LoginModules. Each LoginModule is initialized with a Subject, a Interface CallbackHandler, Shared LoginModule State, and LoginModule-specific options.

The subject represents the user or service currently being authenticated and is updated by a LoginModule with relevant Interface Principal and credentials if authentication succeeds. LoginModules use the CallbackHandler to communicate with users (to prompt for user names and passwords, for example), as described in the login method description. Note that the CallbackHandler may be null. A LoginModule that requires a CallbackHandler to authenticate the Subject may throw a LoginException if it was initialized with a null CallbackHandler. LoginModules optionally use the **shared state** to share information or data among themselves.

The LoginModule-specific options represent the options configured for this LoginModule in the login Configuration. The options are defined by the LoginModule itself and control the behavior within it. For example, a LoginModule may define options to support debugging/testing capabilities. Options are defined using a key-value syntax, such as debug=true. The LoginModule stores the options as a Map so that the values may be retrieved using the key. Note that there is no limit to the number of options a LoginModule chooses to define.

The calling application sees the authentication process as a single operation invoked via a call to the LoginContext's login method. However, the authentication process within each LoginModule proceeds in two distinct phases. In the first phase of authentication, the LoginContext's login method invokes the login method of each LoginModule specified in the Configuration. The login method for a LoginModule performs the actual authentication (prompting for and verifying a password for example) and saves its authentication status as private state information. Once finished, the LoginModule's login method returns true (if it succeeded) or false (if it should be ignored), or it throws a LoginException to specify a failure. In the failure case, the LoginModule must not retry the authentication or introduce delays. The responsibility of such tasks belongs to the application. If the application attempts to retry the authentication, each LoginModule's login method will be called again.

In the second phase, if the LoginContext's overall authentication succeeded (calls to the relevant required, requisite, sufficient and optional LoginModules' login methods succeeded), then the commit method for each LoginModule gets invoked. (For an explanation of the LoginModule flags required, requisite, sufficient and optional, please consult the Configuration documentation and Appendix B: Example Login Configurations in the JAAS Reference Guide.) The commit method for a LoginModule checks its privately saved state to see if its own authentication succeeded. If the overall LoginContext authentication succeeded and the LoginModule's own authentication succeeded, then the commit method associates the relevant Principals (authenticated identities) and credentials (authentication data such as cryptographic keys) with the Subject.

If the LoginContext's overall authentication failed (the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL LoginModules' login methods did not succeed), then the abort method for each LoginModule gets invoked. In this case, the LoginModule removes/destroys any authentication state originally saved.

Logging out a subject involves only one phase. The LoginContext invokes the LoginModule's logout method. The logout method for the LoginModule then performs the



Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide

logout procedures, such as removing Principals or credentials from the Subject, or logging session information.

Steps to Implement a LoginModule

The steps required in order to implement and test a LoginModule:

Step 1: Understand the Authentication Technology

The first thing you need to do is understand the authentication technology to be implemented by your new LoginModule provider, and determine its requirements.

 Determine whether or not your LoginModule will require some form of user interaction (retrieving a user name and password, for example). If so, you will need to become familiar with the Interface CallbackHandler and the javax.security.auth.callback.

In that package you will find several possible <code>callback</code> implementations to use. (Alternatively, you can create your own <code>Callback</code> implementations.) The <code>LoginModule</code> will invoke the <code>CallbackHandler</code> specified by the application itself and passed to the <code>LoginModule</code>'s <code>initialize</code> method. The <code>LoginModule</code> passes the <code>CallbackHandler</code> an array of appropriate <code>CallbackS</code>. See the login method in Step 3.



It is possible for ${\tt LoginModule}$ implementations not to have any end-user interactions. Such ${\tt LoginModules}$ would not need to access the callback package.

Determine what configuration options you want to make available to the user, who
specifies configuration information in whatever form the current Configuration
implementation expects (for example, in files). For each option, decide the option
name and possible values.

For example, if a LoginModule may be configured to consult a particular authentication server host, decide on the option's key name ("auth_server", for example), as well as the possible server hostnames valid for that option ("server_one.example.com" and "server_two.example.com", for example).

Step 2: Name the LoginModule Implementation

Decide on the proper package and class name for your LoginModule.

For example, a LoginModule developed by IBM might be called com.ibm.auth.Module where com.ibm.auth is the package name and Module is the name of the LoginModule class implementation.

Step 3: Implement the Abstract LoginModule Methods

The LoginModule interface specifies five abstract methods that require implementations.

LoginModule.initialize Method

public void initialize (



```
Subject subject,
CallbackHandler handler,
Map<java.lang.String, ?> sharedState,
Map<java.lang.String, ?> options) { ... }
```

The initialize method is called to initialize the LoginModule with the relevant authentication and state information.

This method is called by a LoginContext immediately after this LoginModule has been instantiated, and prior to any calls to its other public methods. The method implementation should store away the provided arguments for future use.

The initialize method may additionally peruse the provided sharedState to determine what additional authentication state it was provided by other LoginModules, and may also traverse through the provided options to determine what configuration options were specified to affect the LoginModule's behavior. It may save option values in variables for future use.

Note: JAAS LoginModules may use the options defined in PAM (use_first_pass, try_first_pass, use_mapped_pass, and try_mapped_pass) to achieve single-sign on. See Making Login Services Independent from Authentication Technologies for further information.

Below is a list of options commonly supported by LoginModules. Note that the following is simply a guideline. Modules are free to support a subset (or none) of the following options.

- try_first_pass If true, the first LoginModule in the stack saves the password entered, and subsequent LoginModules also try to use it. If authentication fails, the LoginModules prompt for a new password and retry the authentication.
- use_first_pass If true, the first LoginModule in the stack saves the password entered, and subsequent LoginModules also try to use it. LoginModules do not prompt for a new password if authentication fails (authentication simply fails).
- try_mapped_pass If true, the first LoginModule in the stack saves the password entered, and subsequent LoginModules attempt to map it into their service-specific password. If authentication fails, the LoginModules prompt for a new password and retry the authentication.
- use_mapped_pass If true, the first LoginModule in the stack saves the password entered, and subsequent LoginModules attempt to map it into their service-specific password. LoginModules do not prompt for a new password if authentication fails (authentication simply fails).
- moduleBanner If true, then when invoking the CallbackHandler, the LoginModule provides a TextOutputCallback as the first Callback, which describes the LoginModule performing the authentication.
- debug If true, instructs a LoginModule to output debugging information.

The initialize method may freely ignore state or options it does not understand, although it would be wise to log such an event if it does occur.

Note that the LoginContext invoking this LoginModule (and the other configured LoginModules, as well), all share the same references to the provided Subject and sharedState. Modifications to the Subject and sharedState will, therefore, be seen by all.



LoginModule.login Method

boolean login() throws LoginException;

The login method is called to authenticate a Subject. This is phase 1 of authentication.

This method implementation should perform the actual authentication. For example, it may cause prompting for a user name and password, and then attempt to verify the password against a password database. Another example implementation may inform the user to insert their finger into a fingerprint reader, and then match the input fingerprint against a fingerprint database.

If your LoginModule requires some form of user interaction (retrieving a user name and password, for example), it should not do so directly. That is because there are various ways of communicating with a user, and it is desirable for LoginModules to remain independent of the different types of user interaction. Rather, the LoginModule's login method should invoke the handle method of the Interface CallbackHandler passed to the initialize method to perform the user interaction and set appropriate results, such as the user name and password. The LoginModule passes the CallbackHandler an array of appropriate Callbacks, for example a NameCallbackfor the user name and a PasswordCallbackfor the password, and the CallbackHandler performs the requested user interaction and sets appropriate values in the Callbacks. For example, to process a NameCallback, the CallbackHandler may prompt for a name, retrieve the value from the user, and call the NameCallback's setName method to store the name.

The authentication process may also involve communication over a network. For example, if this method implementation performs the equivalent of a *kinit* in Kerberos, then it would need to contact the KDC. If a password database entry itself resides in a remote naming service, then that naming service needs to be contacted, perhaps via the Java Naming and Directory Interface (JNDI). Implementations might also interact with an underlying operating system. For example, if a user has already logged into an operating system like Solaris, Linux, macOS, or Windows NT, this method might simply import the underlying operating system's identity information.

The login method should

- 1. Determine whether or not this LoginModule should be ignored. One example of when it should be ignored is when a user attempts to authenticate under an identity irrelevant to this LoginModule (if a user attempts to authenticate as root using NIS, for example). If this LoginModule should be ignored, login should return false. Otherwise, it should do the following:
- 2. Call the CallbackHandler handle method if user interaction is required.
- 3. Perform the authentication.
- 4. Store the authentication result (success or failure).
- 5. If authentication succeeded, save any relevant state information that may be needed by the commit method.
- 6. Return true if authentication succeeds, or throw a LoginException such as FailedLoginException if authentication fails.

Note that the login method implementation should not associate any new Principal or credential information with the saved Subject object. This method merely performs the authentication, and then stores away the authentication result and corresponding authentication state. This result and state will later be accessed by the commit or abort



method. Note that the result and state should typically not be saved in the *sharedState* Map, as they are not intended to be shared with other LoginModules.

An example of where this method might find it useful to store state information in the *sharedState* Map is when LoginModules are configured to share passwords. In this case, the entered password would be saved as shared state. By sharing passwords, the user only enters the password once, and can still be authenticated to multiple LoginModules. The standard conventions for saving and retrieving names and passwords from the *sharedState* Map are the following:

- javax.security.auth.login.name Use this as the shared state map key for saving/ retrieving a name.
- javax.security.auth.login.password Use this as the shared state map key for saving/retrieving a password.

If authentication fails, the <code>login</code> method should not retry the authentication. This is the responsibility of the application. Multiple <code>LoginContext login</code> method calls by an application are preferred over multiple login attempts from within <code>LoginModule.login()</code>.

LoginModule.commit Method

boolean commit() throws LoginException;

The commit method is called to commit the authentication process. This is phase 2 of authentication when phase 1 succeeds. It is called if the LoginContext's overall authentication succeeded (that is, if the relevant REQUIRED, REQUISITE, SUFFICIENT and OPTIONAL LoginModules succeeded.)

This method should access the authentication result and corresponding authentication state saved by the <code>login</code> method.

If the authentication result denotes that the login method failed, then this commit method should remove/destroy any corresponding state that was originally saved.

If the saved result instead denotes that this LoginModule's login method succeeded, then the corresponding state information should be accessed to build any relevant Principal and credential information. Such Principals and credentials should then be added to the Subject stored away by the initialize method.

After adding Principals and credentials, dispensable state fields should be destroyed expeditiously. Likely fields to destroy would be user names and passwords stored during the authentication process.

The commit method should save private state indicating whether the commit succeeded or failed.

The following chart depicts what a LoginModule's commit method should return. The different boxes represent the different situations that may occur. For example, the top-left corner box depicts what the commit method should return if both the previous call to login succeeded and the commit method itself succeeded.

Table 6-1 LoginModule.commit method return values

Login Status	COMMIT: SUCCESS	COMMIT: FAILURE
LOGIN: SUCCESS	return TRUE	throw EXCEPTION
LOGIN: FAILURE	return FALSE	return FALSE



LoginModule.abort Method

boolean abort() throws LoginException;

The abort method is called to abort the authentication process. This is phase 2 of authentication when phase 1 fails. It is called if the LoginContext's overall authentication failed.

This method first accesses this LoginModule's authentication result and corresponding authentication state saved by the login (and possibly commit) methods, and then clears out and destroys the information. Sample state to destroy would be user names and passwords.

If this LoginModule's authentication attempt failed, then there shouldn't be any private state to clean up.

The following charts depict what a LoginModule's abort method should return. This first chart assumes that the previous call to login succeeded. For instance, the top-left corner box depicts what the abort method should return if both the previous call to login and commit succeeded, and the abort method itself also succeeded.

Table 6-2 LoginModule.abort method return values; login succeeded

Login Status	ABORT: SUCCESS	ABORT: FAILURE
COMMIT: SUCCESS	return TRUE	throw EXCEPTION
COMMIT: FAILURE	return TRUE	throw EXCEPTION

LoginModule.logout Method

boolean logout() throws LoginException;

The logout method is called to log out a Subject.

This method removes Principals, and removes/destroys credentials associated with the Subject during the commit operation. This method should not touch those Principals or credentials previously existing in the Subject, or those added by other LoginModuleS.

If the Subject has been marked read-only (the Subject's isReadOnly method returns true), then this method should only destroy credentials associated with the Subject during the commit operation (removing the credentials is not possible). If the Subject has been marked as read-only and the credentials associated with the Subject during the commit operation are not destroyable (they do not implement the Destroyable interface), then this method may throw a LoginException.

The logout method should return true if logout succeeds, or otherwise throw a LoginException.

Step 4: Choose or Write a Sample Application

Either choose an existing sample application for your testing, or write a new one.



See Java Authentication and Authorization Service (JAAS) Reference Guide for information about application requirements and a sample application you can use for your testing.

Step 5: Compile the LoginModule and Application

Compile your new LoginModule and the application you will use for testing.

Step 6: Prepare for Testing

Prepare for testing the LoginModule.

Step 6a: Place Your LoginModule and Application Code in JAR Files

Place your LoginModule and application code in separate JAR files, in preparation for referencing the JAR files in the policy in Step 6c: Set LoginModule and Application JAR File Permissions. Here is a sample command for creating a JAR file:

jar cvf <JAR file name> <list of classes, separated by spaces>

This command creates a JAR file with the specified name containing the specified classes.

For more information on the jar tool, see jar.

Step 6b: Decide Where to Store the JAR Files

The application can be stored essentially anywhere you like.

Your LoginModule can also be placed anywhere you (and other clients) like. If the LoginModule is fully trusted, it can be placed in the JRE's lib/ext (standard extension) directory.

You will need to test the LoginModule being located both in the lib/ext directory and elsewhere because in one situation your LoginModule will need to explicitly be granted Permissions in the Java Development Kit (JDK) required for any security-sensitive operations it does, while in the other case such permissions are not needed.

If your LoginModule is placed in the JRE's lib/ext directory, it will be treated as an installed extension and no permissions need to be granted, since the default system Policy File Syntax grants all permissions to installed extensions.

If your LoginModule is placed anywhere else, the permissions need to be granted, for example by grant statements in a policy file.

Decide where you will store the LoginModule JAR file for testing the case where it is not an installed extension. In the next step, you grant permissions to the JAR file, in the specified location.

Step 6c: Set LoginModule and Application JAR File Permissions

If your LoginModule and/or application performs security-sensitive tasks that will trigger security checks (making network connections, reading or writing files on a local disk, etc), it will need to be granted the required Permissions in the Java Development Kit (JDK) if it is not an installed extension (see Step 6b: Decide Where to Store the JAR Files) and it is run while a security manager is installed.



Since LoginModules usually associate Principals and credentials with an authenticated Subject, some types of permissions a LoginModule will typically require are AuthPermissions with target names "modifyPrincipals", "modifyPublicCredentials", and "modifyPrivateCredentials".

A sample statement granting permissions to a LoginModule whose code is in MyLM. jar appears below. Such a statement could appear in a policy file. In this example, the MyLM. jar file is assumed to be in the /localWork directory.

```
grant codeBase "file:/localWork/MyLM.jar" {
  permission javax.security.auth.AuthPermission "modifyPrincipals";
  permission javax.security.auth.AuthPermission "modifyPublicCredentials";
  permission javax.security.auth.AuthPermission "modifyPrivateCredentials";
};
```

Note:

Since a LoginModule is always invoked within an AccessController.doPrivileged call, it should not have to call doPrivileged itself. If it does, it may inadvertently open up a security hole. For example, a LoginModule that invokes the application-provided CallbackHandler inside a doPrivileged call opens up a security hole by permitting the application's CallbackHandler to gain access to resources it would otherwise not have been able to access.

Step 6d: Create a Configuration Referencing the LoginModule

Because JAAS supports a pluggable authentication architecture, your new LoginModule can be used without requiring modifications to existing applications. Only the login Configuration needs to be updated in order to indicate use of a new LoginModule.

The default <code>configuration</code> implementation from Oracle reads configuration information from configuration files, as described in <code>configFile</code>.

Create a configuration file to be used for testing. For example, to configure the previously-mentioned hypothetical IBM LoginModule for an application, the configuration file might look like this:

```
AppName {
    com.ibm.auth.Module REQUIRED debug=true;
};
```

where AppName should be whatever name the application uses to refer to this entry in the login configuration file. The application specifies this name as the first argument to the LoginContext constructor.

Step 7: Test Use of the LoginModule

Test your application and its use of the LoginModule. When you run the application, specify the login configuration file to be used. For example, suppose your application is named MyApp, it is located in MyApp, jar, and your configuration file is test.conf.

You could run the application and specify the configuration file via the following:



```
java -classpath MyApp.jar
-Djava.security.auth.login.config=test.conf MyApp
```

Type all that on one line. Multiple lines are used here for legibility.

To specify a policy file named my.policy and run the application with a security manager installed, do the following:

```
java -classpath MyApp.jar -Djava.security.manager
-Djava.security.policy=my.policy
-Djava.security.auth.login.config=test.conf MyApp
```

Again, type all that on one line.

You may want to configure the LoginModule with a *debug* option to help ensure that it is working correctly.

Debug your code and continue testing as needed. If you have problems, review the steps above and ensure they are all completed.

Be sure to vary user input and the LoginModule options specified in the configuration file.

Be sure to also include testing using different installation options (e.g., making the LoginModule an installed extension or placing it on the class path) and execution environments (with or without a security manager running). Installation options are discussed in Step 6b: Decide Where to Store the JAR Files. In particular, in order to ensure your LoginModule works when a security manager is installed and the LoginModule and applications are not installed extensions, you need to test such an installation and execution environment, after granting required permissions, as described in Step 6c: Set LoginModule and Application JAR File Permissions.

- 1. If you find during testing that your LoginModule or application needs modifications, make the modifications, recompile (Step 5: Compile the LoginModule and Application).
- 2. Place the updated code in a JAR file (Step 6a: Place Your LoginModule and Application Code in JAR Files).
- 3. Re-install the JAR file (Step 6b: Decide Where to Store the JAR Files).
- 4. If needed fix or add to the permissions (Step 6c: Set LoginModule and Application JAR File Permissions).
- 5. If needed modify the login configuration file (Step 6d: Create a Configuration Referencing the LoginModule).
- 6. Re-run the application and repeat these steps as needed.

Step 8: Document Your LoginModule Implementation

Write documentation for clients of your LoginModule.

Example documentation you may want to include is:

- A README or User Guide describing
 - 1. The authentication process employed by your LoginModule implementation.
 - 2. Information on how to install the LoginModule.



- 3. Configuration options accepted by the LoginModule. For each option, specify the option name and possible values (or types of values), as well as the behavior the option controls.
- 4. The permissions required by your LoginModule when it is run with a security manager (and it is not an installed extension).
- An example Configuration file that references your new LoginModule.
- An example policy file granting your LoginModule the required permissions.
- API documentation. Putting javadoc comments into your source code as you
 write it will make the API javadocs easy to generate.

Step 9: Make LoginModule JAR File and Documents Available

Make your LoginModule JAR file and documentation available to clients.



7

Java Generic Security Services (Java GSS-API)

Java Generic Security Services (Java GSS-API) is used for securely exchanging messages between communicating applications.

Java GSS-API and JAAS Tutorials for Use with Kerberos

See Java GSS-API and JAAS Tutorials for Use with Kerberos in the JDK 8 documentation for a series of tutorials demonstrating various aspects of Java Authentication and Authorization Service (JAAS) and Java Generic Security Services Application Program Interface (GSS-API).

Single Sign-on Using Kerberos in Java

See Single Sign-on Using Kerberos in Java in the JDK 8 documentation for more information about how to use single sign-on based on the Kerberos V5 protocol.

Java GSS Advanced Security Programming

See Java GSS Advanced Security Programming in the JDK 8 documentation for exercises that show you to use the Java GSS-API to build applications that authenticate users, communicate securely with other applications and services, and configure applications in a Kerberos environment to achieve Single Sign-On. In addition, these exercises show you how to use stronger encryption algorithms in a Kerberos environment and Java GSS mechanisms, such as SPNEGO, to secure the association.

The Kerberos 5 GSS-API Mechanism

See The Kerberos 5 GSS-API Mechanism in the JDK 8 documentation for information about Java Generic Security Services (Java GSS) for Kerberos 5.

8

Java Secure Socket Extension (JSSE) Reference Guide

The Java Secure Socket Extension (JSSE) enables secure Internet communications. It provides a framework and an implementation for a Java version of the SSL, TLS, and DTLS protocols and includes functionality for data encryption, server authentication, message integrity, and optional client authentication.

Introduction to JSSE

Data that travels across a network can easily be accessed by someone who is not the intended recipient. When the data includes private information, such as passwords and credit card numbers, steps must be taken to make the data unintelligible to unauthorized parties. It is also important to ensure that the data has not been modified, either intentionally or unintentionally, during transport. The Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols were designed to help protect the privacy and integrity of data while it is being transferred across a network.

The Java Secure Socket Extension (JSSE) enables secure Internet communications. It provides a framework and an implementation for a Java version of the SSL and TLS protocols and includes functionality for data encryption, server authentication, message integrity, and optional client authentication. Using JSSE, developers can provide for the secure passage of data between a client and a server running any application protocol (such as HTTP, Telnet, or FTP) over TCP/IP. For an introduction to SSL, see Secure Sockets Layer (SSL) Protocol Overview.

By abstracting the complex underlying security algorithms and handshaking mechanisms, JSSE minimizes the risk of creating subtle but dangerous security vulnerabilities. Furthermore, it simplifies application development by serving as a building block that developers can integrate directly into their applications.

JSSE provides both an application programming interface (API) framework and an implementation of that API. The JSSE API supplements the core network and cryptographic services defined by the <code>java.security</code> and <code>java.net</code> packages by providing extended networking socket classes, trust managers, key managers, SSL contexts, and a socket factory framework for encapsulating socket creation behavior. Because the <code>SSLSocket</code> class is based on a blocking I/O model, the Java Development Kit (JDK) includes a nonblocking <code>SSLEngine</code> class to enable implementations to choose their own I/O methods.

The JSSE API supports the following security protocols:

SSL: version 3.0

TLS: version 1.0, 1.1, and 1.2

DTLS: versions 1.0 and 1.2

These security protocols encapsulate a normal bidirectional stream socket, and the JSSE API adds transparent support for authentication, encryption, and integrity

protection.sions. Note that the JSSE implementation that is shipped with the JDK does not implement SSL 2.0.

JSSE is a security component of the Java SE platform, and is based on the same design principles found elsewhere in the Java Cryptography Architecture (JCA) Reference Guide framework. This framework for cryptography-related security components allows them to have implementation independence and, whenever possible, algorithm independence. JSSE uses the Cryptographic Service Providers defined by the JCA framework.

Other security components in the Java SE platform include the Java Authentication and Authorization Service (JAAS) Reference Guide and the Java Security Tools. JSSE encompasses many of the same concepts and algorithms as those in JCA but automatically applies them underneath a simple stream socket API.

The JSSE API was designed to allow other SSL/TLS/DTLS protocol and Public Key Infrastructure (PKI) implementations to be plugged in seamlessly. Developers can also provide alternative logic to determine if remote hosts should be trusted or what authentication key material should be sent to a remote host.

JSSE Features and Benefits

JSSE includes the following important benefits and features:

- Included as a standard component of the JDK
- Extensible, provider-based architecture
- Implemented in 100% pure Java
- Provides API support for SSL/TLS/DTLS
- Provides implementations of SSL 3.0, TLS (versions 1.0, 1.1, and 1.2), and DTLS (versions 1.0 and 1.2)
- Includes classes that can be instantiated to create secure channels (SSLSocket, SSLServerSocket, and SSLEngine)
- Provides support for cipher suite negotiation, which is part of the SSL/TLS/DTLS handshaking used to initiate or verify secure communications
- Provides support for client and server authentication, which is part of the normal SSL/TLS/DTLS handshaking
- Provides support for HTTP encapsulated in the SSL/TLS protocol, which allows access to data such as web pages using HTTPS
- Provides server session management APIs to manage memory-resident SSL sessions
- Provides support for server name indication extension, which facilitates secure connections to virtual servers.
- Provides support for certificate status request extension (OCSP stapling), which saves client certificate validation round-trips and resources.
- Provides support for Server Name Indication (SNI) Extension, which extends the SSL/TLS/DTLS protocols to indicate what server name the client is attempting to connect to during handshaking.
- Provides support for endpoint identification during handshaking, which prevents man-in-the-middle attacks.



 Provides support for cryptographic algorithm constraints, which provides finegrained control over algorithms negotiated by JSSE.

JSSE uses the following cryptographic algorithms:

Table 8-1 Cryptographic Algorithms Used by JSSE

Cryptographic Functionality	Cryptographic Algorithm ¹	Key Lengths (Bits) ²
Bulk encryption	Advanced Encryption Standard (AES)	256 ³ 128
Bulk encryption	Data Encryption Standard (DES)	64 (56 effective) 64 (40 effective)
Bulk encryption	Rivest Cipher 4 (RC4)	128 128 (40 effective)
Bulk encryption	Triple DES (3DES)	192 (112 effective)
Hash algorithm	Message Digest Algorithm (MD5)	128
Hash algorithm	Secure Hash Algorithm 1 (SHA1)	160
Hash algorithm	Secure Hash Algorithm 224 (SHA224)	224
Hash algorithm	Secure Hash Algorithm 256 (SHA256)	256
Hash algorithm	Secure Hash Algorithm 384 (SHA384)	384
Hash algorithm	Secure Hash Algorithm 512 (SHA512)	512
Authentication	Digital Signature Algorithm (DSA)	1024, 2048, 3072
Authentication	Elliptic Curve Digital Signature Algorithm (ECDSA)	160 through 512
Authentication and key exchange	Rivest-Shamir-Adleman (RSA)	512 and larger
Key exchange	Static Elliptic Curve Diffie- Hellman (ECDH)	160 through 512
Key exchange	Ephemeral Elliptic Curve Diffie-Hellman (ECDHE)	160 through 512
Key agreement	Diffie-Hellman (DH)	512, 768, 1024, 2048, 3072 4096, 6144, 8192

¹ The SunJSSE implementation uses the Java Cryptography Architecture (JCA) for all its cryptographic algorithms.

JSSE Standard API

The JSSE standard API, available in the <code>javax.net</code> and <code>javax.net.ssl</code> packages, provides:

Secure sockets tailored to client and server-side applications.



 $^{^{\,2}\,\,}$ A JSSE provider may disable or deactivate weak algorithms and weak keys.

³ Cipher suites that use AES_256 require the appropriate Java Cryptography Extension (JCE) unlimited strength jurisdiction policy file set, which is included in the JDK. By default, the active cryptography policy is unlimited. See Cryptographic Strength Configuration.

- A non-blocking engine for producing and consuming streams of SSL/TLS/DTLS data (SSLEngine).
- Factories for creating sockets, server sockets, SSL sockets, and SSL server sockets. By using socket factories, you can encapsulate socket creation and configuration behavior.
- A class representing a secure socket context that acts as a factory for secure socket factories and engines.
- Key and trust manager interfaces (including X.509-specific key and trust managers), and factories that can be used for creating them.
- A class for secure HTTP URL connections (HTTPS).

SunJSSE Provider

Oracle's implementation of Java SE includes a JSSE provider named *SunJSSE*, which comes preinstalled and preregistered with the JCA. This provider supplies the following cryptographic services:

- An implementation of the SSL 3.0, TLS (versions 1.0, 1.1, and 1.2), and DTLS (versions 1.0 and 1.2) security protocols.
- An implementation of the most common SSL, TLS, and DTLS cipher suites. This
 implementation encompasses a combination of authentication, key agreement,
 encryption, and integrity protection.
- An implementation of an X.509-based key manager that chooses appropriate authentication keys from a standard JCA keystore.
- An implementation of an X.509-based trust manager that implements rules for certificate chain path validation.

See The SunJSSE Provider.

JSSE Related Documentation

The following list contains links to online documentation and names of books about related subjects:

JSSE API Documentation

- javax.net package
- javax.net.ssl package

Java SE Security

- The Java SE Security home page
- The Security Features in Java SE trail of the Java Tutorial
- Java PKI Programmers Guide
- Inside Java 2 Platform Security, Second Edition: Architecture, API Design and Implementation



Cryptography

- The Cryptography and Security page by Dr. Ronald L. Rivest (no longer maintained)
- Applied Cryptography, Second Edition by Bruce Schneier. John Wiley and Sons, Inc., 1996.
- Cryptography Theory and Practice by Doug Stinson. CRC Press, Inc., 1995. Third edition published in 2005.
- *Cryptography & Network Security: Principles & Practice* by William Stallings. Prentice Hall, 1998. Fifth edition published in 2010.

Secure Sockets Layer (SSL)

- The Secure Sockets Layer (SSL) Protocol Version 3.0 RFC
- The TLS Protocol Version 1.0 RFC
- HTTP Over TLS RFC
- SSL and TLS: Designing and Building Secure Systems by Eric Rescorla. Addison Wesley Professional, 2000.
- SSL and TLS Essentials: Securing the Web by Stephen Thomas. John Wiley and Sons, Inc., 2000.
- Java 2 Network Security, Second Edition, by Marco Pistoia, Duane F Reller, Deepak Gupta, Milind Nagnur, and Ashok K Ramani. Prentice Hall, 1999.

Transport Layer Security (TLS)

- The TLS Protocol Version 1.0 RFC
- The TLS Protocol Version 1.1 RFC
- The TLS Protocol Version 1.2 RFC
- Transport Layer Security (TLS) Extensions
- HTTP Over TLS RFC

Datagram Transport Layer Security (DTLS)

- The DTLS Protocol Version 1.0 RFC
- The DTLS Protocol Version 1.2 RFC

U.S. Encryption Policies

- U.S. Department of Commerce
- Technology CEO Council
- Current export policies: Encryption and Export Administration Regulations (EAR)
- NIST Computer Security Publications

Terms and Definitions

The following are commonly used cryptography terms and their definitions.



authentication

The process of confirming the identity of a party with whom one is communicating.

certificate

A digitally signed statement vouching for the identity and public key of an entity (person, company, and so on). Certificates can either be self-signed or issued by a Certificate Authority (CA) an entity that is trusted to issue valid certificates for other entities. Well-known CAs include Comodo, Entrust, and GoDaddy. X509 is a common certificate format that can be managed by the JDK's keytool.

cipher suite

A combination of cryptographic parameters that define the security algorithms and key sizes used for authentication, key agreement, encryption, and integrity protection.

cryptographic hash function

An algorithm that is used to produce a relatively small fixed-size string of bits (called a hash) from an arbitrary block of data. A cryptographic hash function is similar to a checksum and has three primary characteristics: it's a one-way function, meaning that it is not possible to produce the original data from the hash; a small change in the original data produces a large change in the resulting hash; and it doesn't require a cryptographic key.

Cryptographic Service Provider (CSP)

Sometimes referred to simply as providers for short, the Java Cryptography Architecture (JCA) defines it as a package (or set of packages) that implements one or more engine classes for specific cryptographic algorithms. An engine class defines a cryptographic service in an abstract fashion without a concrete implementation.

Datagram Transport Layer Security (DTLS) Protocol

A protocol that manages client and server authentication, data integrity, and encrypted communication between the client and server based on an unreliable transport channel such as UDP.

decryption

See encryption/decryption.

digital signature

A digital equivalent of a handwritten signature. It is used to ensure that data transmitted over a network was sent by whoever claims to have sent it and that the data has not been modified in transit. For example, an RSA-based digital signature is calculated by first computing a cryptographic hash of the data and then encrypting the hash with the sender's private key.

encryption/decryption

Encryption is the process of using a complex algorithm to convert an original message (cleartext) to an encoded message (ciphertext) that is unintelligible unless it is decrypted. Decryption is the inverse process of producing cleartext from ciphertext. The algorithms used to encrypt and decrypt data typically come in two categories: secret key (symmetric) cryptography and public key (asymmetric) cryptography.

endpoint identification

An IPv4 or IPv6 address used to identify an endpoint on the network. Endpoint identification procedures are handled during SSL/TLS handshake.



handshake protocol

The negotiation phase during which the two socket peers agree to use a new or existing session. The handshake protocol is a series of messages exchanged over the record protocol. At the end of the handshake, new connection-specific encryption and integrity protection keys are generated based on the key agreement secrets in the session.

java-home

Variable placeholder used throughout this document to refer to the directory where the Java Development Kit (JDK) is installed.

key agreement

A method by which two parties cooperate to establish a common key. Each side generates some data, which is exchanged. These two pieces of data are then combined to generate a key. Only those holding the proper private initialization data can obtain the final key. Diffie-Hellman (DH) is the most common example of a key agreement algorithm.

key exchange

A method by which keys are exchanged. One side generates a private key and encrypts it using the peer's public key (typically RSA). The data is transmitted to the peer, who decrypts the key using the corresponding private key.

key manager/trust manager

Key managers and trust managers use keystores for their key material. A key manager manages a keystore and supplies public keys to others as needed (for example, for use in authenticating the user to others). A trust manager decides who to trust based on information in the truststore it manages.

keystore/truststore

A keystore is a database of key material. Key material is used for a variety of purposes, including authentication and data integrity. Various types of keystores are available, including PKCS12 and Oracle's JKS.

Generally speaking, keystore information can be grouped into two categories: key entries and trusted certificate entries. A key entry consists of an entity's identity and its private key, and can be used for a variety of cryptographic purposes. In contrast, a trusted certificate entry contains only a public key in addition to the entity's identity. Thus, a trusted certificate entry can't be used where a private key is required, such as in a <code>javax.net.ssl.KeyManager</code>. In the JDK implementation of JKS, a keystore may contain both key entries and trusted certificate entries.

A truststore is a keystore that is used when making decisions about what to trust. If you receive data from an entity that you already trust, and if you can verify that the entity is the one that it claims to be, then you can assume that the data really came from that entity.

An entry should only be added to a truststore if the user trusts that entity. By either generating a key pair or by importing a certificate, the user gives trust to that entry. Any entry in the truststore is considered a trusted entry.

It may be useful to have two different keystore files: one containing just your key entries, and the other containing your trusted certificate entries, including CA certificates. The former contains private information, whereas the latter does not. Using two files instead of a single keystore file provides a cleaner separation of the logical distinction between your own certificates (and corresponding private keys) and others' certificates. To provide more protection for your private keys, store them in a keystore with restricted access, and provide the trusted certificates in a more publicly accessible keystore if needed.



message authentication code (MAC)

Provides a way to check the integrity of information transmitted over or stored in an unreliable medium, based on a secret key. Typically, MACs are used between two parties that share a secret key in order to validate information transmitted between these parties.

A MAC mechanism that is based on cryptographic hash functions is referred to as HMAC. HMAC can be used with any cryptographic hash function, such as Message Digest 5 (MD5) and the Secure Hash Algorithm (SHA-256), in combination with a secret shared key. HMAC is specified in RFC 2104.

public-key cryptography

A cryptographic system that uses an encryption algorithm in which two keys are produced. One key is made public, whereas the other is kept private. The public key and the private key are cryptographic inverses; what one key encrypts only the other key can decrypt. Public-key cryptography is also called asymmetric cryptography.

Record Protocol

A protocol that packages all data (whether application-level or as part of the handshake process) into discrete records of data much like a TCP stream socket converts an application byte stream into network packets. The individual records are then protected by the current encryption and integrity protection keys.

secret-key cryptography

A cryptographic system that uses an encryption algorithm in which the same key is used both to encrypt and decrypt the data. Secret-key cryptography is also called symmetric cryptography.

Secure Sockets Layer (SSL) Protocol

A protocol that manages client and server authentication, data integrity, and encrypted communication between the client and server.

session

A named collection of state information including authenticated peer identity, cipher suite, and key agreement secrets that are negotiated through a secure socket handshake and that can be shared among multiple secure socket instances.

Transport Layer Security (TLS) Protocol

A protocol that manages client and server authentication, data integrity, and encrypted communication between the client and server based on a reliable transport channel such as TCP.

TLS 1 is the successor of the SSL 3.0 protocol.

trust manager

See key manager/trust manager.

truststore

See keystore/truststore.

Secure Sockets Layer (SSL) Protocol Overview

Secure Sockets Layer (SSL) is the most widely used protocol for implementing cryptography on the web. SSL uses a combination of cryptographic processes to provide secure communication over a network. This section provides an introduction to SSL and the cryptographic processes it uses.



SSL provides a secure enhancement to the standard TCP/IP sockets protocol used for Internet communications. As shown in Table 8-2, the secure sockets layer is added between the transport layer and the application layer in the standard TCP/IP protocol stack. The application most commonly used with SSL is Hypertext Transfer Protocol (HTTP), the protocol for Internet web pages. Other applications, such as Net News Transfer Protocol (NNTP), Telnet, Lightweight Directory Access Protocol (LDAP), Interactive Message Access Protocol (IMAP), and File Transfer Protocol (FTP), can be used with SSL as well.

Table 8-2 TCP/IP Protocol Stack with SSL

TCP/IP Layer	Protocol
Application Layer	HTTP, NNTP, Telnet, FTP, and so on
Secure Sockets Layer	SSL
Transport Layer	TCP
Internet Layer	IP

SSL was developed by Netscape in 1994, and with input from the Internet community, has evolved to become a standard. It is now under the control of the international standards organization, the Internet Engineering Task Force (IETF). The IETF renamed SSL to Transport Layer Security (TLS), and released the first specification, version 1.0, in January 1999. TLS 1.0 is a modest upgrade to the most recent version of SSL, version 3.0. This upgrade corrected defects in previous versions and prohibited the use of known weak algorithms. TLS 1.1 was released in April 2006, and TLS 1.2 in August 2008.

Why Use SSL?

Transferring sensitive information over a network can be risky due to the following issues:

- You can't always be sure that the entity with whom you are communicating is really who you think it is.
- Network data can be intercepted, so it's possible that it can be read by an unauthorized third party, sometimes known as an attacker.
- An attacker who intercepts data may be able to modify it before sending it on to the receiver.

SSL addresses each of these issues. It addresses the first issue by optionally allowing each of two communicating parties to ensure the identity of the other party in a process called authentication. After the parties are authenticated, SSL provides an encrypted connection between the two parties for secure message transmission. Encrypting the communication between the two parties provides privacy and therefore addresses the second issue. The encryption algorithms used with SSL include a secure hash function, which is similar to a checksum. This ensures that data isn't modified in transit. The secure hash function addresses the third issue of data integrity.



Note:

Both authentication and encryption are optional and depend on the negotiated cipher suites between the two entities.

An e-commerce transaction is an obvious example of when to use SSL. In an e-commerce transaction, it would be foolish to assume that you can guarantee the identity of the server with whom you are communicating. It would be easy enough for someone to create a phony website promising great services if only you enter your credit card number. SSL allows you, the client, to authenticate the identity of the server. It also allows the server to authenticate the identity of the client, although in Internet transactions, this is seldom done.

After the client and the server are comfortable with each other's identity, SSL provides privacy and data integrity through the encryption algorithms that it uses. This allows sensitive information, such as credit card numbers, to be transmitted securely over the Internet.

Although SSL provides authentication, privacy, and data integrity, it doesn't provide nonrepudiation services. Nonrepudiation means that an entity that sends a message can't later deny sending it. When the digital equivalent of a signature is associated with a message, the communication can later be proved. SSL alone does not provide nonrepudiation.

How SSL Works

One of the reasons that SSL is effective is that it uses several different cryptographic processes. SSL uses public-key cryptography to provide authentication, and secret-key cryptography with hash functions to provide for privacy and data integrity. Before you can understand SSL, it's helpful to understand these cryptographic processes.

Cryptographic Processes

The primary purpose of cryptography is to make it difficult for an unauthorized third party to access and understand private communication between two parties. It is not always possible to restrict all unauthorized access to data, but private data can be made unintelligible to unauthorized parties through the process of encryption. Encryption uses complex algorithms to convert the original message (cleartext) to an encoded message (ciphertext). The algorithms used to encrypt and decrypt data that is transferred over a network typically come in two categories: secret-key cryptography and public-key cryptography.

Both secret-key cryptography and public-key cryptography depend on the use of an agreed-upon cryptographic key or pair of keys. A key is a string of bits that is used by the cryptographic algorithm or algorithms during the process of encrypting and decrypting the data. A cryptographic key is like a key for a lock; only with the right key can you open the lock.

Safely transmitting a key between two communicating parties is not a trivial matter. A public key certificate enables a party to safely transmit its public key, while providing assurance to the receiver of the authenticity of the public key. See Public Key Certificates.



The descriptions of the cryptographic processes in secret-key cryptography and public-key cryptography follow conventions widely used by the security community: the two communicating parties are labeled with the names Alice and Bob. The unauthorized third party, also known as the attacker, is named Charlie.

Secret-Key Cryptography

With secret-key cryptography, both communicating parties, Alice and Bob, use the same key to encrypt and decrypt the messages. Before any encrypted data can be sent over the network, both Alice and Bob must have the key and must agree on the cryptographic algorithm that they will use for encryption and decryption

One of the major problems with secret-key cryptography is the logistical issue of how to get the key from one party to the other without allowing access to an attacker. If Alice and Bob are securing their data with secret-key cryptography, and if Charlie gains access to their key, then Charlie can understand any secret messages he intercepts between Alice and Bob. Not only can Charlie decrypt Alice's and Bob's messages, but he can also pretend that he is Alice and send encrypted data to Bob. Bob won't know that the message came from Charlie, not Alice.

After the problem of secret key distribution is solved, secret-key cryptography can be a valuable tool. The algorithms provide excellent security and encrypt data relatively quickly. The majority of the sensitive data sent in an SSL session is sent using secret-key cryptography.

Secret-key cryptography is also called symmetric cryptography because the same key is used to both encrypt and decrypt the data. Well-known secret-key cryptographic algorithms include Advanced Encryption Standard (AES), Triple Data Encryption Standard (3DES), and Rivest Cipher 4 (RC4).

Public-Key Cryptography

Public-key cryptography solves the logistical problem of key distribution by using both a public key and a private key. The public key can be sent openly through the network while the private key is kept private by one of the communicating parties. The public and the private keys are cryptographic inverses of each other; what one key encrypts, the other key will decrypt.

Assume that Bob wants to send a secret message to Alice using public-key cryptography. Alice has both a public key and a private key, so she keeps her private key in a safe place and sends her public key to Bob. Bob encrypts the secret message to Alice using Alice's public key. Alice can later decrypt the message with her private key.

If Alice encrypts a message using her private key and sends the encrypted message to Bob, then Bob can be sure that the data he receives comes from Alice; if Bob can decrypt the data with Alice's public key, the message must have been encrypted by Alice with her private key, and only Alice has Alice's private key. The problem is that anybody else can read the message as well because Alice's public key is public. Although this scenario does not allow for secure data communication, it does provide the basis for digital signatures. A digital signature is one of the components of a public key certificate, and is used in SSL to authenticate a client or a server. See Public Key Certificates and Digital Signatures.

Public-key cryptography is also called *asymmetric cryptography* because different keys are used to encrypt and decrypt the data. A well-known public key cryptographic algorithm often used with SSL is the Rivest Shamir Adleman (RSA) algorithm. Another



public key algorithm used with SSL that is designed specifically for secret key exchange is the Diffie-Hellman (DH) algorithm. Public-key cryptography requires extensive computations, making it very slow. It is therefore typically used only for encrypting small pieces of data, such as secret keys, rather than for the bulk of encrypted data communications.

Comparison Between Secret-Key and Public-Key Cryptography

Both secret-key cryptography and public-key cryptography have strengths and weaknesses. With secret-key cryptography, data can be encrypted and decrypted quickly, but because both communicating parties must share the same secret key information, the logistics of exchanging the key can be a problem. With public-key cryptography, key exchange is not a problem because the public key does not need to be kept secret, but the algorithms used to encrypt and decrypt data require extensive computations, and are therefore very slow

Public Key Certificates

A public key certificate provides a safe way for an entity to pass on its public key to be used in asymmetric cryptography. The public key certificate avoids the following situation: if Charlie creates his own public key and private key, he can claim that he is Alice and send his public key to Bob. Bob will be able to communicate with Charlie, but Bob will think that he is sending his data to Alice.

A public key certificate can be thought of as the digital equivalent of a passport. It is issued by a trusted organization and provides identification for the bearer. A trusted organization that issues public key certificates is known as a Certificate Authority (CA). The CA can be likened to a notary public. To obtain a certificate from a CA, one must provide proof of identity. Once the CA is confident that the applicant represents the organization it says it represents, the CA signs the certificate attesting to the validity of the information contained within the certificate.

A public key certificate contains the following fields:

Issuer

The CA that issued the certificate. If a user trusts the CA that issued the certificate, and if the certificate is valid, then the user can trust the certificate.

Period of validity

A certificate has an expiration date. This date should be checked when verifying the validity of a certificate.

Subject

Includes information about the entity that the certificate represents.

Subject's public key

The primary piece of information that the certificate provides is the subject's public key. All the other fields are provided to ensure the validity of this key.

Signature

The certificate is digitally signed by the CA that issued the certificate. The signature is created using the CA's private key and ensures the validity of the certificate. Because only the certificate is signed, not the data sent in the SSL transaction, SSL does not provide for nonrepudiation.



If Bob only accepts Alice's public key as valid when she sends it in a public key certificate, then Bob won't be fooled into sending secret information to Charlie when Charlie masquerades as Alice.

Multiple certificates may be linked in a certificate chain. When a certificate chain is used, the first certificate is always that of the sender. The next is the certificate of the entity that issued the sender's certificate. If more certificates are in the chain, then each is that of the authority that issued the previous certificate. The final certificate in the chain is the certificate for a root CA. A root CA is a public Certificate Authority that is widely trusted. Information for several root CAs is typically stored in the client's Internet browser. This information includes the CA's public key. Well-known CAs include VeriSign, Entrust, and GTE CyberTrust.

Cryptographic Hash Functions

When sending encrypted data, SSL typically uses a cryptographic hash function to ensure data integrity. The hash function prevents Charlie from tampering with data that Alice sends to Bob.

A cryptographic hash function is similar to a checksum. The main difference is that whereas a checksum is designed to detect accidental alterations in data, a cryptographic hash function is designed to detect deliberate alterations. When data is processed by a cryptographic hash function, a small string of bits, known as a hash, is generated. The slightest change to the message typically makes a large change in the resulting hash. A cryptographic hash function does not require a cryptographic key. A hash function often used with SSL is Secure Hash Algorithm (SHA). SHA was proposed by the U.S. National Institute of Standards and Technology (NIST).

Message Authentication Code

A message authentication code (MAC) is similar to a cryptographic hash, except that it is based on a secret key. When secret key information is included with the data that is processed by a cryptographic hash function, then the resulting hash is known as an HMAC.

If Alice wants to be sure that Charlie does not tamper with her message to Bob, then she can calculate an HMAC for her message and append the HMAC to her original message. She can then encrypt the message plus the HMAC using a secret key that she shares with Bob. When Bob decrypts the message and calculates the HMAC, he will be able to tell if the message was modified in transit. With SSL, an HMAC is used with the transmission of secure data.

Digital Signatures

Once a cryptographic hash is created for a message, the hash is encrypted with the sender's private key. This encrypted hash is called a digital signature.

The SSL Handshake

Communication using SSL begins with an exchange of information between the client and the server. This exchange of information is called the SSL handshake. The SSL handshake includes the following stages:

Negotiating the cipher suite



The SSL session begins with a negotiation between the client and the server as to which cipher suite they will use. A cipher suite is a set of cryptographic algorithms and key sizes that a computer can use to encrypt data. The cipher suite includes information about the public key exchange algorithms or key agreement algorithms, and cryptographic hash functions. The client tells the server which cipher suites it has available, and the server chooses the best mutually acceptable cipher suite.

2. Authenticating the server's identity (optional)

In SSL, the authentication step is optional, but in the example of an e-commerce transaction over the web, the client will generally want to authenticate the server. Authenticating the server allows the client to be sure that the server represents the entity that the client believes the server represents.

To prove that a server belongs to the organization that it claims to represent, the server presents its public key certificate to the client. If this certificate is valid, then the client can be sure of the identity of the server.

The client and server exchange information that allows them to agree on the same secret key. For example, with RSA, the client uses the server's public key, obtained from the public key certificate, to encrypt the secret key information. The client sends the encrypted secret key information to the server. Only the server can decrypt this message because the server's private key is required for this decryption.

3. Agreeing on encryption mechanisms

Both the client and the server now have access to the same secret key. With each message, they use the cryptographic hash function, chosen in the first step of the handshake, and shared secret information, to compute an HMAC that they append to the message. They then use the secret key and the secret key algorithm negotiated in the first step of the handshake to encrypt the secure data and the HMAC. The client and server can now communicate securely using their encrypted and hashed data.

The SSL Protocol

The SSL Handshake provides a high-level description of the SSL handshake, which is the exchange of information between the client and the server prior to sending the encrypted message. Figure 8-1 provides more detail. It shows the sequence of messages that are exchanged in the SSL handshake. Messages that are sent only in certain situations are noted as optional. Each of the SSL messages is described in detail afterward.



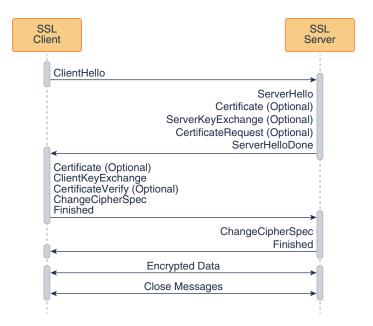


Figure 8-1 The SSL/TLS Handshake

The SSL messages are sent in the following order:

- 1. Client hello: The client sends the server information including the highest version of SSL that it supports and a list of the cipher suites that it supports (TLS 1.0 is indicated as SSL 3.1). The cipher suite information includes cryptographic algorithms and key sizes.
- 2. **Server hello:** The server chooses the highest version of SSL and the best cipher suite that both the client and server support and sends this information to the client.
- 3. (Optional) **Certificate:** The server sends the client a certificate or a certificate chain. A certificate chain typically begins with the server's public key certificate and ends with the certificate authority's root certificate. This message is optional, but is used whenever server authentication is required.
- 4. (Optional) Certificate request: If the server must authenticate the client, then it sends the client a certificate request. In Internet applications, this message is rarely sent.
- 5. (Optional) **Server key exchange:** The server sends the client a server key exchange message if the public key information from the Certificate is not sufficient for key exchange. For example, in cipher suites based on Diffie-Hellman (DH), this message contains the server's DH public key.
- **6. Server hello done:** The server tells the client that it is finished with its initial negotiation messages.
- 7. (Optional) **Certificate:** If the server Certificate request from the client, the client sends its certificate chain, just as the server did previously.



Only a few Internet server applications ask for a certificate from the client.

- 8. Client key exchange: The client generates information used to create a key to use for symmetric encryption. For RSA, the client then encrypts this key information with the server's public key and sends it to the server. For cipher suites based on DH, this message contains the client's DH public key.
- 9. (Optional) Certificate verify: This message is sent by the client when the client presents a certificate as previously explained. Its purpose is to allow the server to complete the process of authenticating the client. When this message is used, the client sends information that it digitally signs using a cryptographic hash function. When the server decrypts this information with the client's public key, the server is able to authenticate the client.
- **10.** Change cipher spec: The client sends a message telling the server to change to encrypted mode.
- **11. Finished** The client tells the server that it is ready for secure data communication to begin.
- **12. Change cipher spec:** The server sends a message telling the client to change to encrypted mode.
- **13. Finished:** The server tells the client that it is ready for secure data communication to begin. This is the end of the SSL handshake.
- 14. Encrypted data: The client and the server communicate using the symmetric encryption algorithm and the cryptographic hash function negotiated during the client hello and server hello, and using the secret key that the client sent to the server during the client key exchange. The handshake can be renegotiated at this time. See Handshaking Again (Renegotiation).
- **15.** Close Messages:At the end of the connection, each side sends a close_notify alert to inform the peer that the connection is closed.

If the parameters generated during an SSL session are saved, then these parameters can sometimes be reused for future SSL sessions. Saving SSL session parameters allows encrypted communication to begin much more quickly.

Handshaking Again (Renegotiation)

Once the initial handshake is finished and application data is flowing, either side is free to initiate a new handshake at any time. An application might like to use a stronger cipher suite for especially critical operations, or a server application might want to require client authentication.

Regardless of the reason, the new handshake takes place over the existing encrypted session, and application data and handshake messages are interleaved until a new session is established.

Your application can initiate a new handshake by using one of the following methods:

- SSLSocket.startHandshake()
- SSLEngine.beginHandshake()



✓ Note:

a protocol flaw related to renegotiation was found in 2009. The protocol and the Java SE implementation have both been fixed. See Transport Layer Security (TLS) Renegotiation Issue.

Cipher Suite Choice and Remote Entity Verification

The SSL/TLS protocols define a specific series of steps to ensure a *protected* connection. However, the choice of cipher suite directly affects the type of security that the connection enjoys. For example, if an anonymous cipher suite is selected, then the application has no way to verify the remote peer's identity. If a suite with no encryption is selected, then the privacy of the data cannot be protected. Additionally, the SSL/TLS protocols do not specify that the credentials received must match those that peer might be expected to send. If the connection were somehow redirected to a rogue peer, but the rogue's credentials were acceptable based on the current trust material, then the connection would be considered valid.

When using raw <code>sslsocket</code> and <code>sslengine</code> classes, you should always check the peer's credentials before sending any data. The <code>sslsocket</code> and <code>sslengine</code> classes do not automatically verify that the host name in a URL matches the host name in the peer's credentials. An application could be exploited with URL spoofing if the host name is not verified. Since JDK 7, endpoint identification/verification procedures can be handled during <code>SSL/TLS</code> handshaking. See the

SSLParameters.getEndpointIdentificationAlgorithm method.

Protocols such as HTTPS (HTTP Over TLS) do require host name verification. Since JDK 7, the HTTPS endpoint identification is enforced during handshaking for HttpsURLConnection by default. See the

SSLParameters.getEndpointIdentificationAlgorithm method. Alternatively, applications can use the HostnameVerifier interface to override the default HTTPS host name rules. See The HostnameVerifier Interface and HttpsURLConnection Class.

Client-Driven OCSP and OCSP Stapling

Use the Online Certificate Status Protocol (OCSP) to determine the X.509 certificate revocation status during the Transport Layer Security (TLS) handshake.

X.509 certificates used in TLS can be revoked by the issuing Certificate Authority (CA) if there is reason to believe that a certificate is compromised. You can check the revocation status of certificates during the TLS handshake by using one of the following approaches.

Certificate Revocation List (CRL)

A CRL is a simple list of revoked certificates. The application receiving a certificate gets the CRL from a CRL server and checks if the certificate received is on the list. There are two disadvantages to using CRLs that mean a certificate could be revoked, but the revoked certificate is not listed in the CRL:

 CRLs can become very large so there can be a substantial increase in network traffic.



Many CRLs are created with longer validity periods, which increases the
possibility of a certificate being revoked within that validity period and not
showing up until the next CRL refresh.

See Certificate/CRL Storage Classes topic of the Java PKI Programmer's Guide.

Client-driven OCSP

In client-driven OCSP, the client uses OCSP to contact an OCSP responder to check the certificate's revocation status. The amount of data required is small, and the OCSP responder is likely to be more up-to-date with the revocation status than a CRL. Each client connecting to a server requires an OCSP response for each certificate being checked. If the server is a popular one, and many of the clients are using client-driven OCSP, these OCSP requests can have a negative effect on the performance of the OCSP responder.

OCSP stapling

OCSP stapling enables the server, rather than the client, to make the request to the OCSP responder. The server staples the OCSP response to the certificate and returns it to the client during the TLS handshake. This approach enables the presenter of the certificate, rather than the issuing CA, to bear the resource cost of providing OCSP responses. It also enables the server to cache the OCSP responses and supply them to all clients. This significantly reduces the load on the OCSP responder because the response can be cached and periodically refreshed by the server rather than by each client.

Client-Driven OCSP and Certificate Revocation

Client-driven Online Certificate Status Protocol (OCSP) enables the client to check the certificate revocation status by connecting to an OCSP responder during the Transport Layer Security (TLS) handshake.

The client-driven OCSP request occurs during the TLS handshake just after the client receives the certificate from the server and validates it. See SSL Handshake.

TLS Handshake with Client-Driven OCSP

Client-driven OCSP is used during the TLS handshake between the client and the server to check the server certificate revocation status. After the client receives the certificate it performs certificate validation. If the validation is successful, then the client verifies that the certificate was not revoked by the issuer. This is done by sending an OCSP request to an OCSP responder. After receiving the OCSP response, the client checks this response before to completing the TLS handshake.



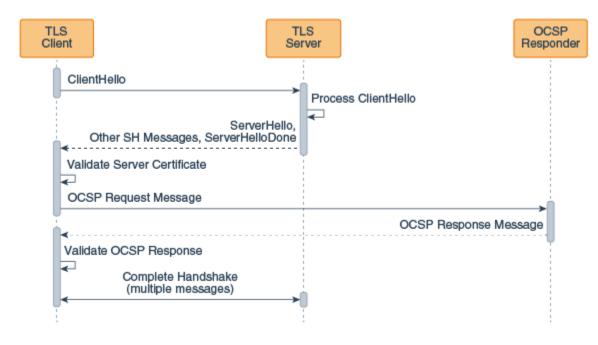


Figure 8-2 TLS Handshake with Client-Driven OCSP

Usually the client finds the OCSP responder's URL by looking in the Authority Information Access (AIA) extension of the certificate, but it can be set to a static URL through the use of a system property.

Setting up a Java Client to use Client-Driven OCSP

Client-driven OCSP is enabled by enabling revocation checking and enabling OCSP.

To configure a Java client to use client-driven OCSP, the Java client must already be set up to connect to a server using TLS.

- 1. Enable revocation checking. You can do this in two different ways.
 - Set the system property com.sun.net.ssl.checkRevocation to true.
 - Use the setRevocationEnabled method on PKIXParameters. See The PKIXParameters Class.
- 2. Enable client-driven OCSP:

Set the Security Property ocsp.enable to true.

Both steps are necessary. The ocsp.enable setting has no effect unless revocation checking is enabled.

OCSP Stapling and Certificate Revocation

Online Certificate Status Protocol (OCSP) stapling enables the presenter of a certificate, rather than the issuing Certificate Authority (CA), to bear the resource cost of providing the OCSP responses that contain the certificate's revocation status.

TLS Handshake with OCSP Stapling

OCSP stapling is used during the Transport Layer Security (TLS) handshake between the client and the server to check the server certificate revocation status. The server makes the OCSP request to the OCSP responder and staples the OCSP responses to the certificates returned to the client. By having the server make the request to the OCSP responder, the responses can be cached, and then used multiple times for many clients.

The TLS handshake begins with the TLS <code>clientHello</code> message. When OCSP stapling is used, this message is sent to the server with the <code>status_request</code> extension that indicates that the server should perform an OCSP request. After processing the <code>clientHello</code> message, the server sends an OCSP request to the appropriate OCSP responder for each certificate. When the server receives the OCSP responses from the OCSP responders, it sends a <code>serverHello</code> message with its <code>status_request</code> extension, indicating that OCSP responses will be provided in the TLS handshake. The server will then present the server certificate chain, followed by a message that consists of one or more OCSP responses for those certificates. The client receiving the certificates with stapled OCSP responses validates each certificate, and then checks the OCSP responses before continuing with the handshake.

If, from the client's perspective, the stapled OCSP response from the server for a certificate is missing, the client will attempt to use client-driven OCSP or CRLs to get revocation information, if either of these are enabled and revocation checking is set to true.



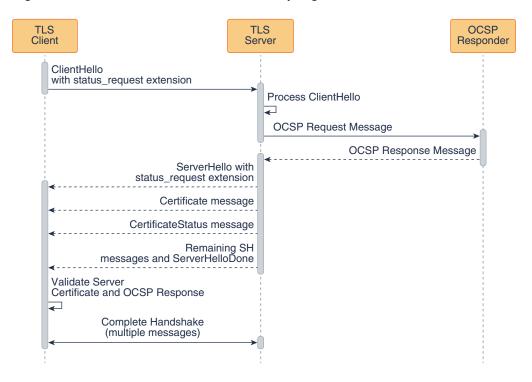


Figure 8-3 TLS Handshake with OCSP Stapling

For more information about TLS handshake messages, see The SSL Handshake.

Status Request Versus Multiple Status Request

The OCSP stapling feature implements the TLS Certificate Status Request extension (section 8 of RFC 6066) and the Multiple Certificate Status Request Extension (RFC 6961).

The TLS Certificate Status Request extension requests revocation information for only the server certificate in the certificate chain while the Multiple Certificate Status Request Extension requests revocation information for all certificates in the certificate chain. In the case where only the server certificate's revocation information is sent to the client, other certificates in the chain may be verified using using the Certificate Revocation Lists (CRLs) or client-driven OCSP (but the client will need to be set up to do this).

Although TLS allows the server to also request the client's certificate, there is no provision in OCSP stapling that enables the client to contact the appropriate OCSP responder and staple the response to the certificate sent to the server.

The OCSP Request and Response

OCSP request and response messages are usually sent over unencrypted HTTP. The response is signed by the CA.

If necessary, the stapled responses can be obtained in the client code by calling the getStatusResponses method on the ExtendedSSLSession object. The method signature is:

public List<byte[]> getStatusResponses();



The OCSP response is encoded using the Distinguished Encoding Rules (DER) in a format described by the ASN.1 found in RFC 6960.

Setting Up a Java Client to Use OCSP Stapling

Online Certificate Status Protocol (OCSP) stapling is enabled on the client side by setting the system property jdk.tls.client.enableStatusRequestExtension to true (its default value).

To configure a Java client to make use of the OCSP response stapled to the certificate returned by a server, the Java client must already be set up to connect to a server using TLS, and the server must be set up to staple an OCSP response to the certificate it returns part of the TLS handshake.

1. Enable OCSP stapling on the client:

If necessary, set the system property jdk.tls.client.enableStatusRequestExtension to true.

- 2. Enable revocation checking. You can do this in two different ways.
 - Set the system property com.sun.net.ssl.checkRevocation to true. You can do this from the command line or in the code.
 - Use the setRevocationEnabled method on the PKIXParameters class. See The PKIXParameters Class.

For the client to include the stapled responses received from the server in the certificate validation, revocation checking must be set to \mathtt{true} . If revocation checking is not set to \mathtt{true} , then the connection will be allowed to proceed regardless of the presence or status of the revocation information

Setting Up a Java Server to Use OCSP Stapling

Online Certificate Status Protocol (OCSP) stapling is enabled on the server by setting the system property jdk.tls.server.enableStatusRequestExtension to true. (It is set to false by default.)

The following steps can be used to configure a Java server to connect to an OCSP responder and staple the OCSP response to the certificate to be returned to the client. The Java server must already be set up to respond to clients using TLS.

- Enable OCSP stapling on the server:
 - Set the system property jdk.tls.server.enableStatusRequestExtension to true.
- Optional: Set other properties as required. See OCSP Stapling Configuration Properties for a list of the valid properties.

OCSP Stapling Configuration Properties

This topic lists the effects of setting various properties when using the Online Certificate Status Protocol (OCSP). It shows the properties used in both client-driven OCSP and OCSP stapling.

Server-side Properties

Most of the properties are read at SSLContext instantiation time. This means that if you set a property, you must obtain a new SSLContext object so that an SSLSocket or



SSLEngine object you obtain from that SSLContext object will reflect the property setting. The one exception is the jdk.tls.stapling.responseTimeout property. That property is evaluated when the ServerHandshaker object is created (essentially at the same time that an SSLSocket or SSLEngine object gets created).

Table 8-3 Server-Side OCSP stapling Properties

Property	Description	Default Value
jdk.tls.server.enableStatusRequestEx tension	Enables the server-side support for OCSP stapling.	False
jdk.tls.stapling.responseTimeout	Controls the maximum amount of time the server will use to obtain OCSP responses, whether from the cache or by contacting an OCSP responder.	5000 (integer value in milliseconds)
	The responses that are already received will be sent in a CertificateStatus message, if applicable based on the type of stapling being done.	
jdk.tls.stapling.cacheSize	Controls the maximum cache size in entries.	256 objects
	If the cache is full and a new response needs to be cached, then the least recently used cache entry will be replaced with the new one. A value of zero or less for this property means that the cache will have no upper bound on the number of responses it can contain.	
jdk.tls.stapling.cacheLifetime	stapling.cacheLifetime Controls the maximum life of a cached response.	
	It is possible for responses to have shorter lifetimes than the value set with this property if the response has a nextUpdate field that expires sooner than the cache lifetime. A value of zero or less for this property disables the cache lifetime. If an object has no nextUpdate value and cache lifetimes are disabled, then the response will not be cached.	hour)
jdk.tls.stapling.responderURI	Enables the administrator to set a default URI in the event that certificates used for TLS do not have the Authority Info Access (AIA) extension.	Not set
	It will not override the Authority Info Access extension value unless the jdk.tls.stapling.responder0verr ide property is set.	



Table 8-3 (Cont.) Server-Side OCSP stapling Properties

Property	Description	Default Value
jdk.tls.stapling.responderOverride	Enables a URI provided through the jdk.tls.stapling.responderURI property to override any AIA extension value.	False
jdk.tls.stapling.ignoreExtensions	Disables the forwarding of OCSP extensions specified in the status_request or status_request_v2 TLS extensions.	False

Client-Side Settings

Table 8-4 Client-Side Settings Used in OCSP Stapling

PKIXBuilderParamet ers	checkRevocation Property	PKIXRevocationChe cker	Result
Default	Default	Default	Revocation checking is disabled.
Default	True	Default	Revocation checking is enabled.[1]
Instantiated	Default	Default	Revocation checking is enabled.[1]
Instantiated	Default	Instantiated, added to PKIXBuilderParamete rs object.	Revocation checking is enabled and[1]will behave according to the PKIXRevocationCheck er settings.

Footnote 1 Note that client-side OCSP fallback will occur only if the ocsp.enable Security Property is set to true.

Developers have some flexibility in how to handle the responses provided through OCSP stapling. OCSP stapling makes no changes to the current methodologies involved in certificate path checking and revocation checking. This means that it is possible to have both client and server assert the status_request extensions, obtain OCSP responses through the CertificateStatus message, and provide user flexibility in how to react to revocation information, or the lack thereof.

If no PKIXBuilderParameters is provided by the caller, then revocation checking is disabled. If the caller creates a PKIXBuilderParameters object and uses the setRevocationEnabled method to enable revocation checking, then stapled OCSP responses will be evaluated. This is also the case if the com.sun.net.ssl.checkRevocation property is set to true.



JSSE Classes and Interfaces

To communicate securely, both sides of the connection must be SSL-enabled. In the JSSE API, the endpoint classes of the connection are <code>SSLSocket</code> and <code>SSLEngine</code>. In Figure 8-4, the major classes used to create <code>SSLSocket</code> and <code>SSLEngine</code> are laid out in a logical ordering.

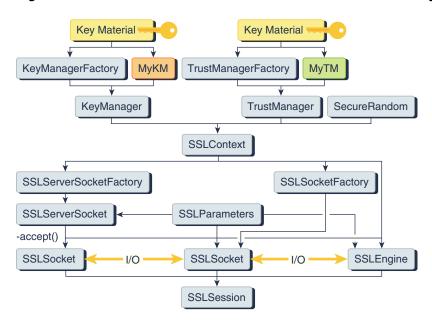


Figure 8-4 JSSE Classes Used to Create SSLSocket and SSLEngine

An SSLSocket is created either by an SSLSocketFactory or by an SSLServerSocket accepting an inbound connection. In turn, an SSLServerSocket is created by an SSLServerSocketFactory. Both SSLSocketFactory and SSLServerSocketFactory objects are created by an SSLContext. An SSLEngine is created directly by an SSLContext, and relies on the application to handle all I/O.

Note:

When using raw <code>sslsocket</code> or <code>sslengine</code> classes, you should always check the peer's credentials before sending any data. Since JDK 7, endpoint identification/verification procedures can be handled during <code>SSL/TLS</code> handshaking. See the method

SSLParameters.setEndpointIdentificationAlgorithm.

For example, the host name in a URL matches the host name in the peer's credentials. An application could be exploited with URL spoofing if the host name is not verified.

JSSE Core Classes and Interfaces



The core JSSE classes are part of the javax.net and javax.net.ssl packages.

SocketFactory and ServerSocketFactory Classes

The abstract <code>javax.net.SocketFactory</code> class is used to create sockets. Subclasses of this class are factories that create particular subclasses of sockets and thus provide a general framework for the addition of public socket-level functionality. For example, see <code>SSLSocketFactory</code> and <code>SSLServerSocketFactory</code> Classes.

The abstract <code>javax.net.ServerSocketFactory</code> class is analogous to the <code>SocketFactory</code> class, but is used specifically for creating server sockets.

Socket factories are a simple way to capture a variety of policies related to the sockets being constructed, producing such sockets in a way that does not require special configuration of the code that asks for the sockets:

- Due to polymorphism of both factories and sockets, different kinds of sockets can be used by the same application code just by passing different kinds of factories.
- Factories can themselves be customized with parameters used in socket construction. For example, factories could be customized to return sockets with different networking timeouts or security parameters already configured.
- The sockets returned to the application can be subclasses of java.net.socket (or javax.net.ssl.SSLSocket), so that they can directly expose new APIs for features such as compression, security, record marking, statistics collection, or firewall tunneling.

SSLSocketFactory and SSLServerSocketFactory Classes

The <code>javax.net.ssl.SSLSocketFactory</code> class acts as a factory for creating secure sockets. This class is an abstract subclass of <code>javax.net.SocketFactory</code>.

Secure socket factories encapsulate the details of creating and initially configuring secure sockets. This includes authentication keys, peer certificate validation, enabled cipher suites, and the like.

The javax.net.ssl.SSLServerSocketFactory class is analogous to the SSLSocketFactory class, but is used specifically for creating server sockets.

Obtaining an SSLSocketFactory

The following ways can be used to obtain an SSLSocketFactory:

- Get the default factory by calling the SSLSocketFactory.getDefault() Static method.
- Receive a factory as an API parameter. That is, code that must create sockets but
 does not care about the details of how the sockets are configured can include a
 method with an SSLSocketFactory parameter that can be called by clients to specify
 which SSLSocketFactory to use when creating sockets (for example,
 javax.net.ssl.HttpsURLConnection).
- Construct a new factory with specifically configured behavior.

The default factory is typically configured to support server authentication only so that sockets created by the default factory do not leak any more information about the client than a normal TCP socket would.



Many classes that create and use sockets do not need to know the details of socket creation behavior. Creating sockets through a socket factory passed in as a parameter is a good way of isolating the details of socket configuration, and increases the reusability of classes that create and use sockets.

You can create new socket factory instances either by implementing your own socket factory subclass or by using another class which acts as a factory for socket factories. One example of such a class is <code>SSLContext</code>, which is provided with the <code>JSSE</code> implementation as a provider-based configuration class.

SSLSocket and SSLServerSocket Classes

The <code>javax.net.ssl.SSLSocket</code> class is a subclass of the standard <code>Java.net.Socket</code> class. It supports all of the standard socket methods and adds methods specific to secure sockets. It supports all of the standard socket methods and adds methods specific to secure sockets. Instances of this class encapsulate the <code>SSLContext</code> under which they were created. See <code>The SSLContext</code> Class. There are APIs to control the creation of secure socket sessions for a socket instance, but trust and key management are not directly exposed.

The javax.net.ssl.SSLServerSocket class is analogous to the SSLSocket class, but is used specifically for creating server sockets.

To prevent peer spoofing, you should always verify the credentials presented to an SSLSocket. See Cipher Suite Choice and Remote Entity Verification.



Due to the complexity of the SSL and TLS protocols, it is difficult to predict whether incoming bytes on a connection are handshake or application data, and how that data might affect the current connection state (even causing the process to block). In the Oracle JSSE implementation, the <code>available()</code> method on the object obtained by <code>SSLSocket.getInputStream()</code> returns a count of the number of application data bytes successfully decrypted from the SSL connection but not yet read by the application.

Obtaining an SSLSocket

Instances of SSLSocket can be obtained in one of the following ways:

- An SSLSocket can be created by an instance of SSLSocketFactory via one of the several createSocket methods of that class.
- An SSLSocket can be created through the accept method of the SSLServerSocket class.

SSLEngine Class

SSL/TLS/DTLS is becoming increasingly popular. It is being used in a wide variety of applications across a wide range of computing platforms and devices. Along with this popularity come demands to use SSL/TLS/DTLS with different I/O and threading models to satisfy the applications' performance, scalability, footprint, and other



requirements. There are demands to use SSL/TLS/DTLS with blocking and nonblocking I/O channels, asynchronous I/O, arbitrary input and output streams, and byte buffers. There are demands to use it in highly scalable, performance-critical environments, requiring management of thousands of network connections.

Abstraction of the I/O transport mechanism using the SSLEngine class in Java SE allows applications to use the SSL/TLS/DTLS protocols in a transport-independent way, and thus frees application developers to choose transport and computing models that best meet their needs. Not only does this abstraction allow applications to use nonblocking I/O channels and other I/O models, it also accommodates different threading models. This effectively leaves the I/O and threading decisions up to the application developer. Because of this flexibility, the application developer must manage I/O and threading (complex topics in and of themselves), as well as have some understanding of the SSL/TLS/DTLS protocols. The abstraction is therefore an advanced API: beginners should use SSLSocket.

Users of other Java programming language APIs such as the Java Generic Security Services (Java GSS-API) and the Java Simple Authentication Security Layer (Java SASL) will notice similarities in that the application is also responsible for transporting data.

The core class is <code>javax.net.ssl.SSLEngine</code>. It encapsulates an SSL/TLS/DTLS state machine and operates on inbound and outbound byte buffers supplied by the user of the <code>SSLEngine</code> class. Figure 8-5 illustrates the flow of data from the application, through <code>SSLEngine</code>, to the transport mechanism, and back.

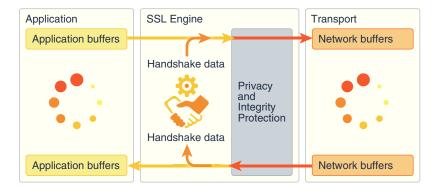


Figure 8-5 Flow of Data Through SSLEngine

The application, shown on the left, supplies application (plaintext) data in an application buffer and passes it to SSLEngine. The SSLEngine object processes the data contained in the buffer, or any handshaking data, to produce SSL/TLS/DTLS encoded data and places it to the network buffer supplied by the application. The application is then responsible for using an appropriate transport (shown on the right) to send the contents of the network buffer to its peer. Upon receiving SSL/TLS/DTLS encoded data from its peer (via the transport), the application places the data into a network buffer and passes it to SSLEngine. The SSLEngine object processes the network buffer's contents to produce handshaking data or application data.

An instance of the SSLEngine class can be in one of the following states:



Creation

The sslengine has been created and initialized, but has not yet been used. During this phase, an application may set any SSLEngine-specific settings (enabled cipher suites, whether the SSLEngine should handshake in client or server mode, and so on). Once handshaking has begun, though, any new settings (except client/server mode) will be used for the next handshake.

Initial handshaking

The initial handshake is a procedure by which the two peers exchange communication parameters until an SSLSession is established. Application data can't be sent during this phase.

Application data

After the communication parameters have been established and the handshake is complete, application data can flow through the sslengine. Outbound application messages are encrypted and integrity protected, and inbound messages reverse the process.

Rehandshaking

Either side can request a renegotiation of the session at any time during the Application Data phase. New handshaking data can be intermixed among the application data. Before starting the rehandshake phase, the application may reset the SSL/TLS/DTLS communication parameters such as the list of enabled ciphersuites and whether to use client authentication, but can not change between client/server modes. As before, after handshaking has begun, any new SSLEngine configuration settings won't be used until the next handshake.

Closure

When the connection is no longer needed, the application should close the SSLEngine and should send/receive any remaining messages to the peer before closing the underlying transport mechanism. Once an engine is closed, it is not reusable: a new SSLEngine must be created.

Creating an SSLEngine Object

Use the SSLContext.createSSLEngine() method to create an SSLEngine object.

Before you create an SSLEngine object, you must configure the engine to act as a client or a server, and set other configuration parameters, such as which cipher suites to use and whether client authentication is required. The SSLContext.createSSLEngine method creates an javax.net.ssl.SSLEngine object.



Note:

The server name and port number are not used for communicating with the server (all transport is the responsibility of the application). They are hints to the JSSE provider to use for SSL session caching, and for Kerberos-based cipher suite implementations to determine which server credentials should be obtained.



Example 8-1 Sample Code for Creating an SSLEngine Client for TLS with JKS as Keystore

The following sample code creates an SSLEngine client for TLS that uses JKS as keystore:

```
import javax.net.ssl.*;
import java.security.*;
// Create and initialize the SSLContext with key material
char[] passphrase = "passphrase".toCharArray();
// First initialize the key and trust material
KeyStore ksKeys = KeyStore.getInstance("JKS");
ksKeys.load(new FileInputStream("testKeys"), passphrase);
KeyStore ksTrust = KeyStore.getInstance("JKS");
ksTrust.load(new FileInputStream("testTrust"), passphrase);
// KeyManagers decide which key material to use
KeyManagerFactory kmf = KeyManagerFactory.getInstance("PKIX");
kmf.init(ksKeys, passphrase);
// TrustManagers decide whether to allow connections
TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
tmf.init(ksTrust);
// Get an instance of SSLContext for SSL/TLS protocols
sslContext = SSLContext.getInstance("TLS");
sslContext.init(kmf.getKeyManagers(), tmf.getTrustManagers(), null);
// Create the engine
SSLEngine engine = sslContext.createSSLengine(hostname, port);
// Use as client
engine.setUseClientMode(true);
```

Generating and Processing SSL/TLS Data

The two main SSLEngine methods are wrap() and unwrap(). They are responsible for generating and consuming network data respectively. Depending on the state of the SSLEngine object, this data might be handshake or application data.

Each SSLEngine object has several phases during its lifetime. Before application data can be sent or received, the SSL/TLS protocol requires a handshake to establish cryptographic parameters. This handshake requires a series of back-and-forth steps by the SSLEngine object. For more details about the handshake itself, see The SSL Handshake.

During the initial handshaking, the wrap() and unwrap() methods generate and consume handshake data, and the application is responsible for transporting the data. The wrap() and unwrap() method sequence is repeated until the handshake is finished. Each SSLEngine operation generates an instance of the SSLEngineResult class, in which the SSLEngineResult.HandshakeStatus field is used to determine what operation must occur next to move the handshake along.

Table 8-5 shows the sequence of methods called during a typical handshake, with corresponding messages and statuses.

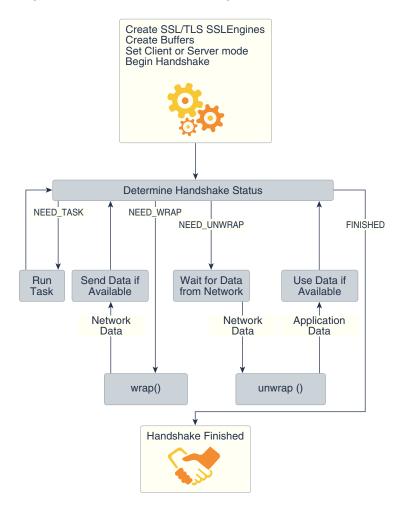


Table 8-5 Typical Handshake

Client	SSL/TLS Message	HandshakeStatus	
wrap()	ClientHello	NEED_UNWRAP	
unwrap()	ServerHello/Cert/ ServerHelloDone	NEED_WRAP	
wrap()	ClientKeyExchange	NEED_WRAP	
wrap()	ChangeCipherSpec	NEED_WRAP	
wrap()	Finished	NEED_UNWRAP	
unwrap()	ChangeCipherSpec	NEED_UNWRAP	
unwrap()	Finished	FINISHED	

Figure 8-6 shows the state machine during a typical SSL/TLS handshake, with corresponding messages and statuses:

Figure 8-6 State Machine during SSL/TLS Handshake



When handshaking is complete, further calls to wrap() will attempt to consume application data and package it for transport. The unwrap() method will attempt the opposite.

To send data to the peer, the application first supplies the data that it wants to send via <code>sslengine.wrap()</code> to obtain the corresponding SSL/TLS encoded data. The application then sends the encoded data to the peer using its chosen transport mechanism. When the application receives the SSL/TLS encoded data from the peer via the transport mechanism, it supplies this data to the <code>sslengine via sslengine.unwrap()</code> to obtain the plaintext data sent by the peer.

Example 8-2 Sample Code for Creating a Nonblocking SocketChannel

The following example is an SSL application that uses a non-blocking <code>socketChannel</code> to communicate with its peer. It sends the string "hello" to the peer by encoding it using the <code>sslengine</code> created in <code>Example 8-1</code>. It uses information from the <code>sslsession</code> to determine how large to make the byte buffers.



The example can be made more robust and scalable by using a selector with the nonblocking socketChannel.

```
// Create a nonblocking socket channel
SocketChannel socketChannel = SocketChannel.open();
socketChannel.configureBlocking(false);
socketChannel.connect(new InetSocketAddress(hostname, port));
// Complete connection
while (!socketChannel.finishedConnect()) {
// do something until connect completed
//Create byte buffers for holding application and encoded data
SSLSession session = engine.getSession();
ByteBuffer myAppData = ByteBuffer.allocate(session.getApplicationBufferSize());
ByteBuffer myNetData = ByteBuffer.allocate(session.getPacketBufferSize());
ByteBuffer peerAppData = ByteBuffer.allocate(session.getApplicationBufferSize());
ByteBuffer peerNetData = ByteBuffer.allocate(session.getPacketBufferSize());
// Do initial handshake
doHandshake(socketChannel, engine, myNetData, peerNetData);
myAppData.put("hello".getBytes());
myAppData.flip();
while (myAppData.hasRemaining()) {
// Generate SSL/TLS/DTLS encoded data (handshake or application data)
SSLEngineResult res = engine.wrap(myAppData, myNetData);
// Process status of call
if (res.getStatus() == SSLEngineResult.Status.OK) {
    myAppData.compact();
    // Send SSL/TLS/DTLS encoded data to peer
```



Example 8-3 Sample Code for Reading Data From Nonblocking SocketChannel

SocketChannelSSLEngineExample 8-1

```
// Read SSL/TLS/DTLS encoded data from peer
int num = socketChannel.read(peerNetData);
if (num == -1) {
   // The channel has reached end-of-stream
} else if (num == 0) {
   // No bytes read; try again ...
} else {
   // Process incoming data
   peerNetData.flip();
   res = engine.unwrap(peerNetData, peerAppData);
    if (res.getStatus() == SSLEngineResult.Status.OK) {
       peerNetData.compact();
    if (peerAppData.hasRemaining()) {
        // Use peerAppData
// Handle other status: BUFFER_OVERFLOW, BUFFER_UNDERFLOW, CLOSED
}
```

Datagram Transport Layer Security (DTLS) Protocol

Datagram Transport Layer Security (DTLS) protocol is designed to construct "TLS over datagram" traffic that doesn't require or provide reliable or in-order delivery of data. Java Secure Socket Extension (JSSE) API and the SunJSSE security provider support the DTLS protocol.

Because the TLS requires a transparent reliable transport channel such as TCP it can't be used to secure unreliable datagram traffic. DTLS is a datagram-compatible variant of TLS.

The JSSE API now supports DTLS Version 1.0 and DTLS Version 1.2 along with Secure Socket Layer (SSL) and Transport Layer Security (TLS) protocols.

The javax.net.ssl.SSLEngine programming model is used by the JSSE API for DTLS.



The DTLS Handshake

Before application data can be sent or received, the DTLS protocol requires a handshake to establish cryptographic parameters. This handshake requires a series of back-and-forth messages between the client and server by the SSLEngine Object.

DTLS handshake requires all messages be received properly. Thus, in unreliable datagram traffic, missing or delayed packets must be retransmitted. Since <code>javax.net.ssl.SSLEngine</code> is not responsible for I/O operations, it is up to the application to provide timers and signal the SSLEngine when a retransmission is needed. It is important that you implement a timer and retransmission strategy for your application. See Handling Retransmissions in DTLS Connections.

The DTLS handshake includes the following stages:

1. Negotiating the cipher suite

The DTLS session begins with a negotiation between the client and the server as to which cipher suite they will use. A cipher suite is a set of cryptographic algorithms and key sizes that a computer can use to encrypt data. The cipher suite includes information about the public key exchange algorithms or key agreement algorithms, and cryptographic hash functions. The client tells the server which cipher suites it has available, and the server chooses the best mutually acceptable cipher suite.

A cookie is exchanged between the client and server along with the cipher suite in order to prevent denial of service attacks (DoS).

2. (Optional) Authenticating the server's identity (optional)

The authentication step is optional, but in the example of an e-commerce transaction over the web, the client chooses to authenticate the server. Authenticating the server allows the client to be sure that the server represents the entity that the client believes the server represents.

To prove that a server belongs to the organization that it claims to represent, the server presents its public key certificate to the client. If this certificate is valid, then the client can be sure of the identity of the server.

The client and server exchange information that allows them to agree on the same secret key. For example, with RSA, the client uses the server's public key, obtained from the public key certificate, to encrypt the secret key information. The client sends the encrypted secret key information to the server. Only the server can decrypt this message because the server's private key is required for this decryption.

3. Agreeing on encryption mechanisms

Both the client and the server now have access to the same secret key. With each message, they use the cryptographic hash function, chosen in the first step of the handshake, and shared secret information, to compute an HMAC that they append to the message. They then use the secret key and the secret key algorithm negotiated in the first step of the handshake to encrypt the secure data and the HMAC. The client and server can now communicate securely using their encrypted and hashed data.

The DTLS Handshake Message Exchange

In a DTLS handshake, series of back-and-forth messages are exchanged between the client and server by the SSLEngine object.



Figure 8-7 shows the sequence of messages that are exchanged in the DTLS handshake. Messages that are sent only in certain situations are noted as optional. Each message is described following the figure.

To know more about DTLS handshake messages, see DTLS Version 1.0 and DTLS Version 1.2.

DTLS DTLS Client Server ClientHello HelloVerifyRequest (contains cookie) ClientHello (with cookie) ServerHello Certificate (Optional) ServerKeyExchange (Optional) CertificateRequest (Optional) ServerHelloDone Certificate (Optional) ClientKeyExchange CertificateVerify (Optional) ChangeCipherSpec Finished ChangeCipherSpec Finished

Figure 8-7 DTLS Handshake

The following handshake messages are exchanged between the client and server during DTLS handshake:

1. ClientHello:

The client sends the server information including the highest version of DTLS that it supports and a list of the cipher suites that it supports. The cipher suite information includes cryptographic algorithms and key sizes.

2. HelloVerifyRequest:

The server responds to the ClientHello message from the client with a cookie.

3. ClientHello:

The client sends a second ClientHello message to the server with highest version of DTLS that it supports and a list of the cipher suites that it supports. The cookie received in the HelloVerifyRequest is sent back to the server.

4. ServerHello:

The server chooses the highest version of DTLS and the best cipher suite that both the client and server support and sends this information to the client.

5. (Optional) Certificate:

The server sends the client a certificate or a certificate chain. A certificate chain typically begins with the server's public key certificate and ends with the certificate authority's root certificate. This message is optional, but is used whenever server authentication is required



6. (Optional) CertificateRequest:

If the server must authenticate the client, then it sends the client a certificate request. In Internet applications, this message is rarely sent.

7. (Optional) ServerKeyExchange:

The server sends the client a server key exchange message if the public key information from the Certificate is not sufficient for key exchange. For example, in cipher suites based on Diffie-Hellman (DH), this message contains the server's DH public key.

8. ServerHelloDone:

The server tells the client that it is finished with its initial negotiation messages.

9. (Optional) Certificate:

If the server Certificate request from the client, the client sends its certificate chain, just as the server did previously.



Only a few Internet server applications ask for a certificate from the client.

10. ClientKeyExchange:

The client generates information used to create a key to use for symmetric encryption. For RSA, the client then encrypts this key information with the server's public key and sends it to the server. For cipher suites based on DH, this message contains the client's DH public key.

11. (Optional) CertificateVerify:

This message is sent by the client when the client presents a certificate as previously explained. Its purpose is to allow the server to complete the process of authenticating the client. When this message is used, the client sends information that it digitally signs using a cryptographic hash function. When the server decrypts this information with the client's public key, the server is able to authenticate the client.

12. ChangeCipherSpec:

The client sends a message telling the server that subsequent data will be protected under the newly negotiated CipherSpec and keys and the data is encrypted

13. Finished:

The client tells the server that it is ready for secure data communication to begin.

14. ChangeCipherSpec:

The server sends a message telling the client that subsequent data will be protected under the newly negotiated CipherSpec and keys and the data is encrypted.

15. Finished:

The server tells the client that it is ready for secure data communication to begin. This is the end of the DTLS handshake.

Handshaking Again (Renegotiation)

Once the initial handshake is finished and application data is flowing, either side is free to initiate a new handshake at any time. An application might like to use a stronger



cipher suite for especially critical operations, or a server application might want to require client authentication.

Regardless of the reason, the new handshake takes place over the existing encrypted session, and application data and handshake messages are interleaved until a new session is established.

Your application can initiate a new handshake by using the ${\tt SSLEngine.beginHandshake()}$ method.



A protocol flaw related to renegotiation was found in 2009. The protocol and the Java SE implementation have both been fixed. See Transport Layer Security (TLS) Renegotiation Issue.

Example 8-4 Sample Code for Handling DTLS handshake Status and Overall Status

```
void handshake(SSLEngine engine, DatagramSocket socket, SocketAddress peerAddr)
throws Exception {
    boolean endLoops = false;
    // private static int MAX HANDSHAKE LOOPS = 60;
    int loops = MAX HANDSHAKE LOOPS;
    engine.beginHandshake();
    while (!endLoops && (serverException == null) && (clientException == null)) {
        if (--loops < 0) {
            throw new RuntimeException("Too many loops to produce handshake
packets");
       SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();
        if (hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP ||
               hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP_AGAIN) {
            ByteBuffer iNet;
            ByteBuffer iApp;
            if (hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP) {
                // Receive ClientHello request and other SSL/TLS/DTLS records
                byte[] buf = new byte[1024];
                DatagramPacket packet = new DatagramPacket(buf, buf.length);
                try {
                    socket.receive(packet);
                } catch (SocketTimeoutException ste) {
                    // Retransmit the packet if timeout
                    List <Datagrampacket> packets = onReceiveTimeout(engine,
peerAddr);
                    for (DatagramPacket p : packets) {
                        socket.send(p);
                    continue;
                iNet = ByteBuffer.wrap(buf, 0, packet.getLength());
                iApp = ByteBuffer.allocate(1024);
                iNet = ByteBuffer.allocate(0);
                iApp = ByteBuffer.allocate(1024);
            SSLEngineResult r = engine.unwrap(iNet, iApp);
            SSLEngineResult.Status rs = r.getStatus();
```



```
hs = r.getHandshakeStatus();
            if (rs == SSLEngineResult.Status.BUFFER_OVERFLOW) {
                // The client maximum fragment size config does not work?
                throw new Exception("Buffer overflow: " +
                                    "incorrect client maximum fragment size");
            } else if (rs == SSLEngineResult.Status.BUFFER_UNDERFLOW) {
                // Bad packet, or the client maximum fragment size
                // config does not work?
                if (hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
                    throw new Exception("Buffer underflow: " +
                                        "incorrect client maximum fragment size");
                } // Otherwise, ignore this packet
            } else if (rs == SSLEngineResult.Status.CLOSED) {
                endLoops = true;
            } // Otherwise, SSLEngineResult.Status.OK
            if (rs != SSLEngineResult.Status.OK) {
                continue;
        } else if (hs == SSLEngineResult.HandshakeStatus.NEED_WRAP) {
            // Call a function to produce handshake packets
            List <DatagramPacket> packets = produceHandshakePackets(engine,
peerAddr);
            for (DatagramPacket p : packets) {
                socket.send(p);
        } else if (hs == SSLEngineResult.HandshakeStatus.NEED_TASK) {
            runDelegatedTasks(engine);
        } else if (hs == SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
            // OK, time to do application data exchange
            endLoops = true;
        } else if (hs == SSLEngineResult.HandshakeStatus.FINISHED) {
            endLoops = true;
    SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();
    if (hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
        throw new Exception("Not ready for application data yet");
```

Handling Retransmissions in DTLS Connections

In SSL/TLS over a reliable connection, data is guaranteed to arrive in the proper order, and retransmission is unnecessary. However, for DTLS, which often works over unreliable media, missing or delayed handshake messages must be retransmitted.

The SSLEngine class operates in a completely transport-neutral manner, and the application layer performs all I/O. Because the SSLEngine class isn't responsible for I/O, the application instead is responsible for providing timers and signalling the SSLEngine class when a retransmission is needed. The application layer must determine the right timeout value and when to trigger the timeout event. During handshaking, if an SSLEngine Object is in HandshakeStatus.NEED_UNWRAP state, a call to SSLEngine.wrap() means that the previous packets were lost, and must be retransmitted. For such cases, the DTLS implementation of the SSLEngine class takes the responsibility to wrap the previous necessary handshaking messages again if necessary.

Note:

In a DTLS engine, only handshake messages must be properly exchanged. Application data can handle packet loss without the need for timers.

Handling Retransmission in an Application

 ${\tt SSLEngine.unwrap()}$ and ${\tt SSLEngine.wrap()}$ can be used together to handle retransmission in an application.

Figure 8-8 shows a typical scenario for handling DTLS handshaking retransmission:

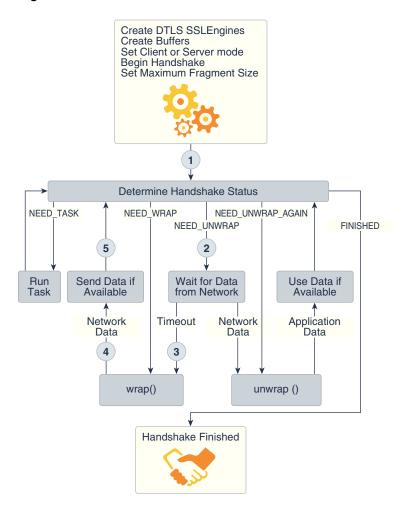


Figure 8-8 DTLS Handshake Retransmission State Flow

- Create and initialize an instance of DTLS SSLEngine.
 - See Creating an SSLEngine Object. The DTLS handshake process begins. See The DTLS Handshake.
- 2. If the handshake status is HandshakeStatus.NEED_UNWRAP, wait for data from network.
- If the timer times out, it indicates that the previous delivered handshake messages may have been lost.

Note:

In DTLS handshaking retransmission, the determined handshake status isn't necessarily <code>HandshakeStatus.NEED_WRAP</code> for the call to <code>SSLEngine.wrap()</code>.

- 4. Call SSLEngine.wrap().
- **5.** The wrapped packets are delivered.

Handling a Buffered Handshake Message in an Application

Datagram transport doesn't require or provide reliable or in-order delivery of data. Handshake messages may be lost or need to be reordered. In the DTLS implementation, a handshake message may need to be buffered for future handling before all previous messages have been received.

The DTLS implementation of SSLEngine takes the responsibility to reorder handshake messages. Handshake message buffering and reordering are transparent to applications.

However, applications must manage ${\tt HandshakeStatus.NEED_UNWRAP_AGAIN}$ status. This status indicates that for the next ${\tt SSLEngine.unwrap}()$ operation no additional data from the remote side is required.

Figure 8-9 shows a typical scenario for using the HandshakeStatus.NEED_UNWRAP_AGAIN.



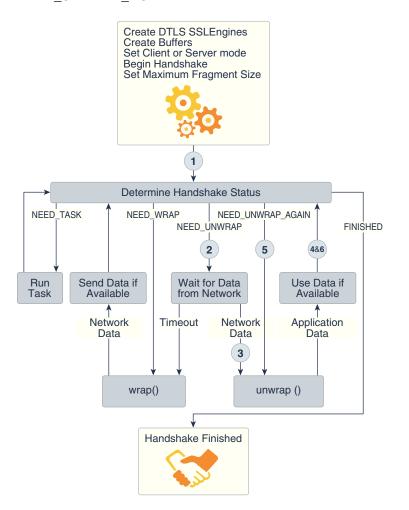


Figure 8-9 State Machine of DTLS Buffered Handshake with NEED UNWRAP AGAIN

- 1. Create and initialize an instance of DTLS sslengine.
 - See Creating an SSLEngine Object. The DTLS handshake process begins, see The DTLS Handshake.
- Optional: If the handshake status is HandshakeStatus.NEED_UNWRAP, wait for data from network.
- 3. Optional: If you received the network data, call SSLEngine.unwrap().
- 4. Determine the handshake status for next processing. The handshake status can be HandshakeStatus.NEED_UNWRAP_AGAIN, HandshakeStatus.NEED_UNWRAP, Or HandshakeStatus.NEED_WRAP.
 - If the handshake status is HandshakeStatus.NEED_UNWRAP_AGAIN, call SSLEngine.unwrap().

Note:

For HandshakeStatus.NEED_UNWRAP_AGAIN status, no additional data from the network is required for an SSLEngine.unwrap() operation.

5. Determine the handshake status for further processing. The handshake status can be HandshakeStatus.NEED_UNWRAP_AGAIN, HandshakeStatus.NEED_UNWRAP, Or HandshakeStatus.NEED_WRAP.

Creating an SSLEngine Object for DTLS

The following examples illustrate how to create an SSLEngine object for DTLS.



The server name and port number are not used for communicating with the server (all transport is the responsibility of the application). They are hints to the JSSE provider to use for DTLS session caching, and for Kerberos-based cipher suite implementations to determine which server credentials should be obtained.

Example 8-5 Sample Code for Creating an SSLEngine Client for DTLS with PKCS12 as Keystore

The following sample code creates an SSLEngine client for DTLS that uses PKCS12 as keystore:

```
import javax.net.ssl.*;
import java.security.*;
// Create and initialize the SSLContext with key material
char[] passphrase = "passphrase".toCharArray();
// First initialize the key and trust material
KeyStore ksKeys = KeyStore.getInstance("PKCS12");
ksKeys.load(new FileInputStream("testKeys"), passphrase);
KeyStore ksTrust = KeyStore.getInstance("PKCS12");
ksTrust.load(new FileInputStream("testTrust"), passphrase);
// KeyManagers decide which key material to use
KeyManagerFactory kmf = KeyManagerFactory.getInstance("PKIX");
kmf.init(ksKeys, passphrase);
// TrustManagers decide whether to allow connections
TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
tmf.init(ksTrust);
// Get an instance of SSLContext for DTLS protocols
sslContext = SSLContext.getInstance("DTLS");
sslContext.init(kmf.getKeyManagers(), tmf.getTrustManagers(), null);
// Create the engine
SSLEngine engine = sslContext.createSSLengine(hostname, port);
```



```
// Use engine as client
engine.setUseClientMode(true);
```

Example 8-6 Sample Code for Creating an SSLEngine Server for DTLS with PKCS12 as Keystore

```
SSLEngine
    import javax.net.ssl.*;
    import java.security.*;
    // Create and initialize the SSLContext with key material
    char[] passphrase = "passphrase".toCharArray();
    // First initialize the key and trust material
    KeyStore ksKeys = KeyStore.getInstance("PKCS12");
    ksKeys.load(new FileInputStream("testKeys"), passphrase);
    KeyStore ksTrust = KeyStore.getInstance("PKCS12");
    ksTrust.load(new FileInputStream("testTrust"), passphrase);
    // KeyManagers decide which key material to use
    KeyManagerFactory kmf = KeyManagerFactory.getInstance("PKIX");
    kmf.init(ksKeys, passphrase);
    // TrustManagers decide whether to allow connections
    TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
    tmf.init(ksTrust);
    // Get an SSLContext for DTLS Protocol without authentication
    sslContext = SSLContext.getInstance("DTLS");
    sslContext.init(null, null, null);
    // Create the engine
    SSLEngine engine = sslContext.createSSLeEngine(hostname, port);
    // Use the engine as server
    engine.setUseClientMode(false);
    // Require client authentication
    engine.setNeedClientAuth(true);
```

Generating and Processing DTLS Data

A DTLS handshake and a SSL/TLS handshake generate and process data similarly. (See Generating and Processing SSL/TLS Data.) They both use the SSLEngine.wrap() and SSLEngine.wrap() methods to generate and consume network data, respectively.

The following diagram shows the state machine during a typical DTLS handshake, with corresponding messages and statuses:



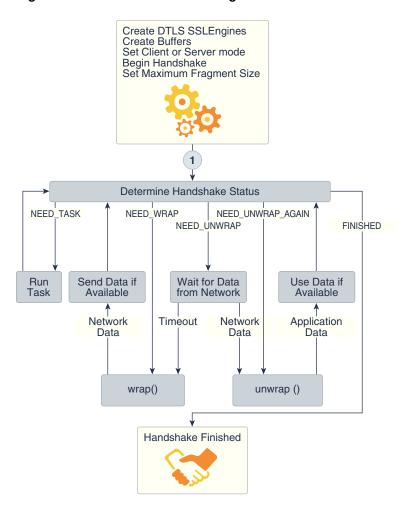


Figure 8-10 State Machine during DTLS Handshake

Difference Between the SSL/TLS and DTLS SSLEngine.wrap() Methods

The <code>SSLEngine.wrap()</code> method for DTLS is different from SSL/TLS as follows:

- In the SSL/TLS implementation of SSLEngine, the output buffer of SSLEngine.wrap() contains one or more TLS records (due to the TLSv1 BEAST Cipher Block Chaining vulnerability).
- In the DTLS implementation of SSLEngine, the output buffer of SSLEngine.wrap() contains at most one record, so that every DTLS record can be marshaled and delivered to the datagram layer individually.



Each record produced by SSLEngine.wrap() should comply to the maximum packet size limitation as specified by SSLParameters.getMaximumPacketSize().



Understanding SSLEngine Operation Statuses

The status of the SSLEngine is represented by SSLEngineResult.Status.

To indicate the status of the engine and what actions the application should take, the <code>SSLEngine.wrap()</code> and <code>SSLEngine.unwrap()</code> methods return an <code>SSLEngineResult</code> instance, as shown in <code>Example 8-2</code>. This <code>SSLEngineResult</code> object contains two pieces of status information: the overall status of the engine and the handshaking status.

The possible overall statuses are represented by the SSLEngineResult.Status enum. The following statuses are available:

OK

There was no error.

CLOSED

The operation closed the SSLEngine or the operation could not be completed because it was already closed.

BUFFER UNDERFLOW

The input buffer had insufficient data, indicating that the application must obtain more data from the peer (for example, by reading more data from the network).

BUFFER OVERFLOW

The output buffer had insufficient space to hold the result, indicating that the application must clear or enlarge the destination buffer.

Example 8-7 illustrates how to handle the <code>BUFFER_UNDERFLOW</code> and <code>BUFFER_OVERFLOW</code> statuses of the <code>SSLEngine.unwrap()</code> method. It uses <code>SSLSession.getApplicationBufferSize()</code> and <code>SSLSession.getPacketBufferSize()</code> to determine how large to make the byte buffers.

The possible handshaking statuses are represented by the

SSLEngineResult.HandshakeStatus enum. They represent whether handshaking has completed, whether the caller must obtain more handshaking data from the peer or send more handshaking data to the peer, and so on. The following handshake statuses are available:

FINISHED

The SSLEngine has just finished handshaking.

NEED TASK

The SSLEngine needs the results of one (or more) delegated tasks before handshaking can continue.

NEED UNWRAP

The SSLEngine needs to receive data from the remote side before handshaking can continue.

NEED UNWRAP AGAIN

The SSLEngine needs to unwrap before handshaking can continue. This value indicates that not-yet-interpreted data has been previously received from the remote side and does not need to be received again; the data has been brought into the JSSE framework but has not been processed yet.



NEED WRAP

The SSLEngine must send data to the remote side before handshaking can continue, so SSLEngine.wrap() should be called.

NOT HANDSHAKING

The SSLEngine is not currently handshaking.

Having two statuses per result allows the SSLEngine to indicate that the application must take two actions: one in response to the handshaking and one representing the overall status of the <code>wrap()</code> and <code>unwrap()</code> methods. For example, the engine might, as the result of a single <code>SSLEngine.unwrap()</code> call, return <code>SSLEngineResult.Status.OK</code> to indicate that the input data was processed successfully and <code>SSLEngineResult.HandshakeStatus.NEED_UNWRAP</code> to indicate that the application should obtain more <code>SSL/TLS/DTLS</code> encoded data from the peer and supply it to <code>SSLEngine.unwrap()</code> again so that handshaking can continue. As you can see, the previous examples were greatly simplified; they would need to be expanded significantly to properly handle all of these statuses.

Example 8-9 and Example 8-8 illustrate how to process handshaking data by checking handshaking status and the overall status of the wrap() and unwrap() methods.

Example 8-7 Sample Code for Handling BUFFER_UNDERFLOW and BUFFER OVERFLOW

The following code sample illustrates how to handle BUFFER_UNDERFLOW and BUFFER_OVERFLOW status:

```
SSLEngineResult res = engine.unwrap(peerNetData, peerAppData);
   switch (res.getStatus()) {
   case BUFFER OVERFLOW:
           // Maybe need to enlarge the peer application data buffer.
       if (engine.getSession().getApplicationBufferSize() > peerAppData.capacity())
{
           // enlarge the peer application data buffer
       } else {
           // compact or clear the buffer
       // retry the operation
   break;
   case BUFFER_UNDERFLOW:
       // Maybe need to enlarge the peer network packet buffer
       if (engine.getSession().getPacketBufferSize() > peerNetData.capacity()) {
       // enlarge the peer network packet buffer
       } else {
       // compact or clear the buffer
       // obtain more inbound network data and then retry the operation
      break;
      // Handle other status: CLOSED, OK
      // ...
   }
```

Example 8-8 Sample Code for Checking and Processing Handshaking Statuses and Overall Statuses

The following code sample illustrates how to process handshaking data by checking handshaking status and the overall status of the wrap() and unwrap() methods:

```
void doHandshake(SocketChannel socketChannel, SSLEngine engine,
    ByteBuffer myNetData, ByteBuffer peerNetData) throws Exception {
    // Create byte buffers to use for holding application data
    int appBufferSize = engine.getSession().getApplicationBufferSize();
    ByteBuffer myAppData = ByteBuffer.allocate(appBufferSize);
    ByteBuffer peerAppData = ByteBuffer.allocate(appBufferSize);
    // Begin handshake
    engine.beginHandshake();
    SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();
    // Process handshaking message
    while (hs != SSLEngineResult.HandshakeStatus.FINISHED &&
       hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
        switch (hs) {
        case NEED_UNWRAP:
            // Receive handshaking data from peer
            if (socketChannel.read(peerNetData) < 0) {</pre>
                // The channel has reached end-of-stream
            // Process incoming handshaking data
            peerNetData.flip();
            SSLEngineResult res = engine.unwrap(peerNetData, peerAppData);
            peerNetData.compact();
            hs = res.getHandshakeStatus();
            // Check status
            switch (res.getStatus()) {
            case OK :
                // Handle OK status
            // Handle other status: BUFFER_UNDERFLOW, BUFFER_OVERFLOW, CLOSED
            break;
        case NEED_WRAP :
            // Empty the local network packet buffer.
            myNetData.clear();
            // Generate handshaking data
            res = engine.wrap(myAppData, myNetData);
            hs = res.getHandshakeStatus();
            // Check status
            switch (res.getStatus()) {
            case OK :
               myNetData.flip();
                // Send the handshaking data to peer
                while (myNetData.hasRemaining()) {
                    socketChannel.write(myNetData);
                break;
            // Handle other status: BUFFER_OVERFLOW, BUFFER_UNDERFLOW, CLOSED
```

```
// ...
}
break;

case NEED_TASK :
    // Handle blocking tasks
    break;

// Handle other status: // FINISHED or NOT_HANDSHAKING
    // ...
}

// Processes after handshaking
// ...
}
```

Example 8-9 Sample Code for Handling DTLS handshake Status and Overall Status

The following code sample illustrates how to handle DTLS handshake status:

```
void handshake(SSLEngine engine, DatagramSocket socket,
               SocketAddress peerAddr) throws Exception {
    boolean endLoops = false;
    // private static int MAX_HANDSHAKE_LOOPS = 60;
    int loops = MAX_HANDSHAKE_LOOPS;
    engine.beginHandshake();
    while (!endLoops && (serverException == null) && (clientException == null)) {
        if (--loops < 0) {
            throw new RuntimeException("Too many loops to produce handshake
packets");
        SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();
        if (hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP | |
                hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP_AGAIN) {
            ByteBuffer iNet;
            ByteBuffer iApp;
            if (hs == SSLEngineResult.HandshakeStatus.NEED_UNWRAP) {
                // receive ClientHello request and other SSL/TLS/DTLS records
                byte[] buf = new byte[1024];
                DatagramPacket packet = new DatagramPacket(buf, buf.length);
                try {
                    socket.receive(packet);
                } catch (SocketTimeoutException ste) {
                    // retransmit the packet if timeout
                    List <Datagrampacket> packets =
                        onReceiveTimeout(engine, peerAddr);
                    for (DatagramPacket p : packets) {
                        socket.send(p);
                    }
                    continue;
                iNet = ByteBuffer.wrap(buf, 0, packet.getLength());
                iApp = ByteBuffer.allocate(1024);
            } else {
                iNet = ByteBuffer.allocate(0);
                iApp = ByteBuffer.allocate(1024);
            SSLEngineResult r = engine.unwrap(iNet, iApp);
            SSLEngineResult.Status rs = r.getStatus();
```

```
hs = r.getHandshakeStatus();
        if (rs == SSLEngineResult.Status.BUFFER_OVERFLOW) {
            // the client maximum fragment size config does not work?
            throw new Exception("Buffer overflow: " +
                                "incorrect client maximum fragment size");
        } else if (rs == SSLEngineResult.Status.BUFFER_UNDERFLOW) {
            // bad packet, or the client maximum fragment size
            // config does not work?
            if (hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
                throw new Exception("Buffer underflow: "
                                    "incorrect client maximum fragment size");
            } // otherwise, ignore this packet
        } else if (rs == SSLEngineResult.Status.CLOSED) {
            endLoops = true;
            // otherwise, SSLEngineResult.Status.OK:
        if (rs != SSLEngineResult.Status.OK) {
            continue;
    } else if (hs == SSLEngineResult.HandshakeStatus.NEED_WRAP) {
       List <DatagramPacket> packets =
            // Call a function to produce handshake packets
            produceHandshakePackets(engine, peerAddr);
        for (DatagramPacket p : packets) {
            socket.send(p);
    } else if (hs == SSLEngineResult.HandshakeStatus.NEED_TASK) {
        runDelegatedTasks(engine);
    } else if (hs == SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
        // OK, time to do application data exchange.
        endLoops = true;
    } else if (hs == SSLEngineResult.HandshakeStatus.FINISHED) {
        endLoops = true;
SSLEngineResult.HandshakeStatus hs = engine.getHandshakeStatus();
if (hs != SSLEngineResult.HandshakeStatus.NOT_HANDSHAKING) {
    throw new Exception("Not ready for application data yet");
```

Dealing With Blocking Tasks

During handshaking, an SSLEngine might encounter tasks that can block or take a long time. For example, a TrustManager may need to connect to a remote certificate validation service, or a KeyManager might need to prompt a user to determine which certificate to use as part of client authentication. To preserve the nonblocking nature of SSLEngine, when the engine encounters such a task, it will return SSLEngineResult.HandshakeStatus.NEED_TASK. Upon receiving this status, the application should invoke SSLEngine.getDelegatedTask() to get the task, and then, using the threading model appropriate for its requirements, process the task. The application might, for example, obtain threads from a thread pool to process the tasks, while the main thread handles other I/O.

The following code executes each task in a newly created thread:

```
if (res.getHandshakeStatus() == SSLEngineResult.HandshakeStatus.NEED_TASK) {
   Runnable task;
   while ((task = engine.getDelegatedTask()) != null) {
       new Thread(task).start();
   }
}
```

```
}
```

The SSLEngine will block future wrap() and unwrap() calls until all of the outstanding tasks are completed.

Shutting Down a SSL/TLS/DTLS Connection

For an orderly shutdown of an SSL/TLS/DTLS connection, the SSL/TLS/DTLS protocols require transmission of close messages. Therefore, when an application is done with the SSL/TLS/DTLS connection, it should first obtain the close messages from the <code>SSLEngine</code>, then transmit them to the peer using its transport mechanism, and finally shut down the transport mechanism. Example 8-10 illustrates this.

In addition to an application explicitly closing the SSLEngine, the SSLEngine might be closed by the peer (via receipt of a close message while it is processing handshake data), or by the SSLEngine encountering an error while processing application or handshake data, indicated by throwing an SSLException. In such cases, the application should invoke SSLEngine.wrap() to get the close message and send it to the peer until SSLEngine.isOutboundDone() returns true (as shown in Example 8-10), or until the SSLEngineResult.getStatus() returns CLOSED.

In addition to orderly shutdowns, there can also be unexpected shutdowns when the transport link is severed before close messages are exchanged. In the previous examples, the application might get -1 or IOException when trying to read from the nonblocking SocketChannel, or get IOException when trying to write to the non-blocking SocketChannel. When you get to the end of your input data, you should call engine.closeInbound(), which will verify with the SSLEngine that the remote peer has closed cleanly from the SSL/TLS/DTLS perspective. Then the application should still try to shut down cleanly by using the procedure in Example 8-10. Obviously, unlike SSLSocket, the application using SSLEngine must deal with more state transitions, statuses, and programming. See Sample Code Illustrating the Use of an SSLEngine.

Example 8-10 Sample Code for Shutting Down a SSL/TLS/DTLS Connection

The following code sample illustrates how to shut down a SSL/TLS/DTLS connection:



SSLSession and ExtendedSSLSession

The <code>javax.net.ssl.SSLSession</code> interface represents a security context negotiated between the two peers of an <code>SSLSocket</code> or <code>SSLEngine</code> connection. After a session has been arranged, it can be shared by future <code>SSLSocket</code> or <code>SSLEngine</code> objects connected between the same two peers.

In some cases, parameters negotiated during the handshake are needed later in the handshake to make decisions about trust. For example, the list of valid signature algorithms might restrict the certificate types that can be used for authentication. The SSLSession can be retrieved during the handshake by calling getHandshakeSession() on an SSLSocket or SSLEngine. Implementations of TrustManager or KeyManager can use the getHandshakeSession() method to get information about session parameters to help them make decisions.

A fully initialized SSLSession contains the cipher suite that will be used for communications over a secure socket as well as a nonauthoritative hint as to the network address of the remote peer, and management information such as the time of creation and last use. A session also contains a shared master secret negotiated between the peers that is used to create cryptographic keys for encrypting and guaranteeing the integrity of the communications over an SSLSocket or SSLEngine connection. The value of this master secret is known only to the underlying secure socket implementation and is not exposed through the SSLSession API.

ExtendedSSLSession extends the SSLSession interface to support additional session attributes. The ExtendedSSLSession class adds methods that describe the signature algorithms that are supported by the local implementation and the peer. The getRequestedServerNames() method called on an ExtendedSSLSession instance is used to obtain a list of sniserverName objects in the requested Server Name Indication (SNI) Extension. The server should use the requested server names to guide its selection of an appropriate authentication certificate, and/or other aspects of the security policy. The client should use the requested server names to guide its endpoint identification of the peer's identity, and/or other aspects of the security policy.

Calls to the <code>getPacketBufferSize()</code> and <code>getApplicationBufferSize()</code> methods on <code>SSLSession</code> are used to determine the appropriate buffer sizes used by <code>SSLEngine</code>.

Note:

The SSL/TLS protocols specify that implementations are to produce packets containing at most 16 kilobytes (KB) of plain text. However, some implementations violate the specification and generate large records up to 32 KB. If the <code>sslengine.unwrap()</code> code detects large inbound packets, then the buffer sizes returned by <code>sslsession</code> will be updated dynamically. Applications should always check the <code>BUFFER_OVERFLOW</code> and <code>BUFFER_UNDERFLOW</code> statuses and enlarge the corresponding buffers if necessary. See <code>Understanding SSlengine Operation Statuses</code>. SunJSSE will always send standard compliant 16 KB records and allow incoming 32 KB records. For a workaround, see the System property <code>jsse.SSlengine.acceptlargeFragments</code> in <code>Customizing JSSE</code>.



HttpsURLConnection Class

The <code>javax.net.ssl.HttpsURLConnection</code> class extends the <code>java.net.HttpURLConnection</code> class and adds support for HTTPS-specific features.

The HTTPS protocol is similar to HTTP, but HTTPS first establishes a secure channel via SSL/TLS sockets and then verifies the identity of the peer (see Cipher Suite Choice and Remote Entity Verification) before requesting or receiving data. The javax.net.ssl.HttpsURLConnection class extends the java.net.HttpURLConnection class and adds support for HTTPS-specific features. To know more about how HTTPS URLs are constructed and used, see the java.net.URL,

java.net.URLConnection, java.net.HttpURLConnection, and javax.net.ssl.HttpsURLConnection classes.

Upon obtaining an httpsurlconnection instance, you can configure a number of HTTP and HTTPS parameters before actually initiating the network connection via the URLConnection.connect() method. Of particular interest are:

- Setting the Assigned SSLSocketFactory
- Setting the Assigned HostnameVerifier

Setting the Assigned SSLSocketFactory

In some situations, it is desirable to specify the <code>SSLSocketFactory</code> that an <code>HttpsURLConnection</code> instance uses. For example, you might want to tunnel through a proxy type that is not supported by the default implementation. The new <code>SSLSocketFactory</code> could return sockets that have already performed all necessary tunneling, thus allowing <code>HttpsURLConnection</code> to use additional proxies.

The HttpsURLConnection class has a default SSLSocketFactory that is assigned when the class is loaded (this is the factory returned by the SSLSocketFactory.getDefault() method). Future instances of HttpsURLConnection will inherit the current default SSLSocketFactory until a new default SSLSocketFactory is assigned to the class via the static HttpsURLConnection.setDefaultSSLSocketFactory() method. Once an instance of HttpsURLConnection has been created, the inherited SSLSocketFactory on this instance can be overridden with a call to the setSSLSocketFactory() method.



Changing the default static ${\tt SSLSocketFactory}$ has no effect on existing instances of ${\tt HttpsURLConnection}$. A call to the ${\tt setSSLSocketFactory}()$ method is necessary to change the existing instances.

You can obtain the per-instance or per-class <code>SSLSocketFactory</code> by making a call to the <code>getSSLSocketFactory()</code> or <code>getDefaultSSLSocketFactory()</code> method, respectively.

Setting the Assigned HostnameVerifier

If the host name of the URL does not match the host name in the credentials received as part of the SSL/TLS handshake, then it is possible that URL spoofing has occurred.



If the implementation cannot determine a host name match with reasonable certainty, then the SSL implementation performs a callback to the instance's assigned <code>HostnameVerifier</code> for further checking. The host name verifier can take whatever steps are necessary to make the determination, such as performing host name pattern matching or perhaps opening an interactive dialog box. An unsuccessful verification by the host name verifier closes the connection. For more information regarding host name verification, see RFC 2818.

The setHostnameVerifier() and setDefaultHostnameVerifier() methods operate in a similar manner to the setSSLSocketFactory() and setDefaultSSLSocketFactory() methods, in that HostnameVerifier objects are assigned on a per-instance and per-class basis, and the current values can be obtained by a call to the getHostnameVerifier() or getDefaultHostnameVerifier() method.

Support Classes and Interfaces

The classes and interfaces in this section are provided to support the creation and initialization of SSLContext objects, which are used to create SSLSocketFactory, SSLServerSocketFactory, and SSLEngine objects. The support classes and interfaces are part of the javax.net.ssl package.

Three of the classes described in this section (The SSLContext Class, The KeyManagerFactory Class, and The TrustManagerFactory Class) are engine classes. An engine class is an API class for specific algorithms (or protocols, in the case of SSLContext), for which implementations may be provided in one or more Cryptographic Service Provider (provider) packages. See JCA Design Principles and Engine Classes and Algorithms.

The SunJSSE provider that comes standard with JSSE provides SSLContext, KeyManagerFactory, and TrustManagerFactory implementations, as well as implementations for engine classes in the standard java.security API. Table 8-6 lists implementations supplied by SunJSSE.

Table 8-6	Implement	tations	Supplied	by Sui	nJSSE
-----------	-----------	---------	----------	--------	-------

Engine Class Implemented	Algorithm or Protocol
KeyStore	PKCS12
KeyManagerFactory	PKIX, SunX509
TrustManagerFactory	PKIX (X509 or SunPKIX), SunX509
SSLContext SSLv3[1], TLSv1, TLSv1. DTLSv1.0, DTLSv1.2	

Footnote 1: Starting with JDK 8u31, the SSLv3 protocol (Secure Socket Layer) has been deactivated and is not available by default. See the <code>java.security.Security</code> property <code>jdk.tls.disabledAlgorithms</code> in the <code><java_home>/conf/security/java.security</code> file. If SSLv3 is absolutely required, the protocol can be reactivated by removing <code>SSLv3</code> from the <code>jdk.tls.disabledAlgorithms</code> property in the <code>java.security</code> file or by dynamically setting this Security Property before JSSE is initialized. To enable SSLv3 protocol at deploy level, after following the previous steps, add the line <code>deployment.security.SSLv3=true</code> to the <code>deployment.properties</code> file.



The SSLContext Class

The <code>javax.net.ssl.SSLContext</code> class is an engine class for an implementation of a secure socket protocol. An instance of this class acts as a factory for <code>sslSocket</code>, <code>sslServerSocket</code>, and <code>sslEngine</code>. An <code>sslContext</code> object holds all of the state information shared across all objects created under that context. For example, session state is associated with the <code>sslContext</code> when it is negotiated through the handshake protocol by sockets created by socket factories provided by the context. These cached sessions can be reused and shared by other sockets created under the same context.

Each instance is configured through its <code>init</code> method with the keys, certificate chains, and trusted root CA certificates that it needs to perform authentication. This configuration is provided in the form of key and trust managers. These managers provide support for the authentication and key agreement aspects of the cipher suites supported by the context.

Currently, only X.509-based managers are supported.

Obtaining and Initializing the SSLContext Class

The SSLContext class is used to create the SSLSocketFactory or SSLServerSocketFactory class.

There are two ways to obtain and initialize an SSLContext:

- The simplest way is to call the static SSLContext.getDefault method on either the SSLSocketFactory Or SSLServerSocketFactory Class. This method creates a default SSLContext with a default KeyManager, TrustManager, and SecureRandom (a secure random number generator). A default KeyManagerFactory and TrustManagerFactory are used to create the KeyManager and TrustManager, respectively. The key material used is found in the default keystore and truststore, as determined by system properties described in Customizing the Default Keystores and Truststores, Store Types, and Store Passwords.
- The approach that gives the caller the most control over the behavior of the created context is to call the static method SSLContext.getDefault
 sSLContext class, and then initialize the context by calling the instance's proper init() method. One variant of the init() method takes three arguments: an array of KeyManager objects, and a SecureRandom object. The KeyManager and TrustManager objects are created by either implementing the appropriate interfaces or using the KeyManagerFactory and TrustManagerFactory classes to generate implementations. The KeyManagerFactory and TrustManagerFactory can then each be initialized with key material contained in the KeyStore passed as an argument to the init() method of the TrustManagerFactory or KeyManagerFactory classes. Finally, the getTrustManagers() method (in KeyManagerFactory) can be called to obtain the array of trust managers or key managers, one for each type of trust or key material.

Once an SSL connection is established, an SSLSession is created which contains various information, such as identities established and cipher suite used. The SSLSession is then used to describe an ongoing relationship and state information between two entities. Each SSL connection involves one session at a time, but that session may be used on many connections between those entities, simultaneously or sequentially.



Creating an SSLContext Object

Like other JCA provider-based engine classes, SSLContext objects are created using the getInstance() factory methods of the SSLContext class. These static methods each return an instance that implements at least the requested secure socket protocol. The returned instance may implement other protocols, too. For example, getInstance("TLSv1") may return an instance that implements TLSv1, TLSv1.1, and TLSv1.2. The getSupportedProtocols() method returns a list of supported protocols when an SSLSocket, SSLServerSocket, or SSLEngine is created from this context. You can control which protocols are actually enabled for an SSL connection by using the setEnabledProtocols(String[] protocols) method.

Note:

An SSLContext object is automatically created, initialized, and statically assigned to the SSLSocketFactory class when you call the SSLSocketFactory.getDefault() method. Therefore, you do not have to directly create and initialize an SSLContext object (unless you want to override the default behavior).

To create an SSLContext object by calling the <code>getInstance()</code> factory method, you must specify the protocol name. You may also specify which provider you want to supply the implementation of the requested protocol:

- public static SSLContext getInstance(String protocol);
- public static SSLContext getInstance(String protocol, String provider);
- public static SSLContext getInstance(String protocol, Provider provider);

If just a protocol name is specified, then the system will determine whether an implementation of the requested protocol is available in the environment. If there is more than one implementation, then it will determine whether there is a preferred one.

If both a protocol name and a provider are specified, then the system will determine whether an implementation of the requested protocol is in the provider requested. If there is no implementation, an exception will be thrown.

A protocol is a string (such as "TLS") that describes the secure socket protocol desired. Common protocol names for SSLContext objects are defined in Java Security Standard Algorithm Names Specification.

An SSLContext can be obtained as follows:

```
SSLContext sc = SSLContext.getInstance("TLS");
```

A newly created SSLContext should be initialized by calling the init method:

```
public void init(KeyManager[] km, TrustManager[] tm, SecureRandom random);
```

If the <code>KeyManager[]</code> parameter is null, then an empty <code>KeyManager</code> will be defined for this context. If the <code>TrustManager[]</code> parameter is null, then the installed security providers will be searched for the highest-priority implementation of the <code>TrustManagerFactory</code> class (see The <code>TrustManagerFactory Class</code>), from which an appropriate <code>TrustManager</code> will be obtained. Likewise, the <code>SecureRandom</code> parameter may be null, in which case a default implementation will be used.



If the internal default context is used, (for example, an <code>SSLContext</code> is created by <code>SSLSocketFactory.getDefault()</code> or <code>SSLServerSocketFactory.getDefault()</code>), then a default <code>KeyManager</code> and <code>TrustManager</code> are created. The default <code>SecureRandom</code> implementation is also chosen.

The TrustManager Interface

The primary responsibility of the TrustManager is to determine whether the presented authentication credentials should be trusted. If the credentials are not trusted, then the connection will be terminated. To authenticate the remote identity of a secure socket peer, you must initialize an SSLContext object with one or more TrustManager objects. You must pass one TrustManager for each authentication mechanism that is supported. If null is passed into the SSLContext initialization, then a trust manager will be created for you. Typically, a single trust manager supports authentication based on X.509 public key certificates (for example, X509TrustManager). Some secure socket implementations may also support authentication based on shared secret keys, Kerberos, or other mechanisms.

TrustManager objects are created either by a TrustManagerFactory, or by providing a concrete implementation of the interface.

The TrustManagerFactory Class

The <code>javax.net.ssl.TrustManagerFactory</code> is an engine class for a provider-based service that acts as a factory for one or more types of <code>TrustManager</code> objects. Because it is provider-based, additional factories can be implemented and configured to provide additional or alternative trust managers that provide more sophisticated services or that implement installation-specific authentication policies.

Creating a TrustManagerFactory

You create an instance of this class in a similar manner to <code>sslContext</code>, except for passing an algorithm name string instead of a protocol name to the <code>getInstance()</code> method:

```
TrustManagerFactory tmf = TrustManagerFactory.getInstance(String algorithm);
TrustManagerFactory tmf = TrustManagerFactory.getInstance(String algorithm, String provider);
TrustManagerFactory tmf = TrustManagerFactory.getInstance(String algorithm, Provider provider);
```

A sample call is as follows:

```
TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX", "SunJSSE");
```

The preceding call creates an instance of the SunJSSE provider's PKIX trust manager factory. This factory can be used to create trust managers that provide X.509 PKIX-based certification path validity checking.

When initializing an SSLContext, you can use trust managers created from a trust manager factory, or you can write your own trust manager, for example, using the CertPath API. See Java PKI Programmer's Guide. You do not need to use a trust manager factory if you implement a trust manager using the X509TrustManager interface.

A newly created factory should be initialized by calling one of the <code>init()</code> methods:

```
public void init(KeyStore ks);
public void init(ManagerFactoryParameters spec);
```

Call whichever init() method is appropriate for the TrustManagerFactory you are using. If you are not sure, then ask the provider vendor.

For many factories, such as the SunX509 TrustManagerFactory from the SunJSSE provider, the KeyStore is the only information required to initialize the TrustManagerFactory and thus the first init method is the appropriate one to call. The TrustManagerFactory will query the KeyStore for information about which remote certificates should be trusted during authorization checks.

Sometimes, initialization parameters other than a Keystore are needed by a provider. Users of that provider are expected to pass an implementation of the appropriate ManagerFactoryParameters as defined by the provider. The provider can then call the specified methods in the ManagerFactoryParameters implementation to obtain the needed information.

For example, suppose the TrustManagerFactory provider requires initialization parameters B, R, and S from any application that wants to use that provider. Like all providers that require initialization parameters other than a KeyStore, the provider requires the application to provide an instance of a class that implements a particular ManagerFactoryParameters subinterface. In the example, suppose that the provider requires the calling application to implement and create an instance of MyTrustManagerFactoryParams and pass it to the second init() method. The following example illustrates what MyTrustManagerFactoryParams can look like:

```
public interface MyTrustManagerFactoryParams extends ManagerFactoryParameters {
  public boolean getBValue();
  public float getRValue();
  public String getSValue();
}
```

Some trust managers can make trust decisions without being explicitly initialized with a KeyStore object or any other parameters. For example, they may access trust material from a local directory service via LDAP, use a remote online certificate status checking server, or access default trust material from a standard local location.

PKIX TrustManager Support

The default trust manager algorithm is PKIX. It can be changed by editing the ${\tt ssl.TrustManagerFactory.algorithm}$ property in the ${\tt java.security}$ file.

The PKIX trust manager factory uses the CertPath PKIX implementation (see PKI Programmers Guide Overview) from an installed security provider. The trust manager factory can be initialized using the normal <code>init(KeyStores)</code> method, or by passing <code>CertPath</code> parameters to the PKIX trust manager using the <code>CertPathTrustManagerParameters</code> class.

Example 8-11 illustrates how to get the trust manager to use a particular LDAP certificate store and enable revocation checking.

If the <code>TrustManagerFactory.init(KeyStore)</code> method is used, then default PKIX parameters are used with the exception that revocation checking is disabled. It can be enabled by setting the <code>com.sun.net.ssl.checkRevocation</code> system property to <code>true</code>. This setting requires that the <code>CertPath</code> implementation can locate revocation information by itself. The PKIX implementation in the provider can do this in many cases but requires that the <code>system property com.sun.security.enableCRLDP</code> be set to <code>true</code>. Note that the



TrustManagerFactory.init(ManagerFactoryParameters) method has revocation checking enabled by default.

See PKIX Classes and The CertPath Class.

Example 8-11 Sample Code for Using a LDAP Certificate to Enable Revocation Checking

The following example illustrates how to get the trust manager to use a particular LDAP certificate store and enable revocation checking:

```
import javax.net.ssl.*;
    import java.security.cert.*;
    import java.security.KeyStore;
    import java.io.FileInputStream;
    // Obtain Keystore password
    char[] pass = System.console().readPassword("Password: ");
    // Create PKIX parameters
    KeyStore anchors = KeyStore.getInstance("JKS");
    anchors.load(new FileInputStream(anchorsFile, pass));
    PKIXBuilderParameters pkixParams = new PKIXBuilderParameters(anchors, new
X509CertSelector());
    // Specify LDAP certificate store to use
    LDAPCertStoreParameters lcsp = new LDAPCertStoreParameters("ldap.imc.org", 389);
    pkixParams.addCertStore(CertStore.getInstance("LDAP", lcsp));
    // Specify that revocation checking is to be enabled
    pkixParams.setRevocationEnabled(true);
    // Wrap PKIX parameters as trust manager parameters
    ManagerFactoryParameters trustParams = new
CertPathTrustManagerParameters(pkixParams);
    // Create TrustManagerFactory for PKIX-compliant trust managers
    TrustManagerFactory factory = TrustManagerFactory.getInstance("PKIX");
    // Pass parameters to factory to be passed to CertPath implementation
    factory.init(trustParams);
    // Use factory
    SSLContext ctx = SSLContext.getInstance("TLS");
    ctx.init(null, factory.getTrustManagers(), null);
```

The X509TrustManager Interface

The <code>javax.net.ssl.X509TrustManager</code> interface extends the general <code>TrustManager</code> interface. It must be implemented by a trust manager when using X.509-based authentication.

To support X.509 authentication of remote socket peers through JSSE, an instance of this interface must be passed to the <code>init</code> method of an <code>SSLContext</code> object.

Creating an X509TrustManager

You can either implement this interface directly yourself or obtain one from a provider-based TrustManagerFactory (such as that supplied by the SunJSSE provider). You could also implement your own interface that delegates to a factory-generated trust manager. For example, you might do this to filter the resulting trust decisions and query an end-user through a graphical user interface.

If a null KeyStore parameter is passed to the SunJSSE PKIX or SunX509 TrustManagerFactory, then the factory uses the following process to try to find trust material:

1. If the <code>javax.net.ssl.trustStore</code> property is defined, then the <code>TrustManagerFactory</code> attempts to find a file using the file name specified by that system property, and uses that file for the <code>KeyStore</code> parameter. If the <code>javax.net.ssl.trustStorePassword</code> system property is also defined, then its value is used to check the integrity of the data in the truststore before opening it.

If the javax.net.ssl.trustStore property is defined but the specified file does not exist, then a default TrustManager using an empty keystore is created.

- 2. If the javax.net.ssl.trustStore system property was not specified, then:
 - if the file java-home/lib/security/jssecacerts exists, that file is used;
 - if the file *java-home*/lib/security/cacerts exists, that file is used;
 - if neither of these files exists, then the SSL cipher suite is anonymous, does not perform any authentication, and thus does not need a truststore.

To know more about what *java-home* refers to, see Terms and Definitions.

The factory looks for a file specified via the <code>javax.net.ssl.trustStore</code> Security Property or for the <code>jssecacerts</code> file before checking for a <code>cacerts</code> file. Therefore, you can provide a JSSE-specific set of trusted root certificates separate from ones that might be present in <code>cacerts</code> for code-signing purposes.

Creating Your Own X509TrustManager

If the supplied X509TrustManager behavior is not suitable for your situation, then you can create your own X509TrustManager by either creating and registering your own TrustManagerFactory or by implementing the X509TrustManager interface directly.

Example 8-12 illustrates a MyX509TrustManager class that enhances the default SunJSSE X509TrustManager behavior by providing alternative authentication logic when the default X509TrustManager fails.

Once you have created such a trust manager, assign it to an SSLContext via the init() method, as in the following example. Future SocketFactories created from this SSLContext will use your new TrustManager when making trust decisions.

```
TrustManager[] myTMs = new TrustManager[] { new MyX509TrustManager() };
SSLContext ctx = SSLContext.getInstance("TLS");
ctx.init(null, myTMs, null);
```



Example 8-12 Sample Code for Creating a X509TrustManager

The following code sample illustrates MyX509TrustManager class that enhances the default SunJSSE X509TrustManager behavior by providing alternative authentication logic when the default X509TrustManager fails:

```
class MyX509TrustManager implements X509TrustManager {
      * The default PKIX X509TrustManager9. Decisions are delegated
      * to it, and a fall back to the logic in this class is performed
      * if the default X509TrustManager does not trust it.
     X509TrustManager pkixTrustManager;
     MyX509TrustManager() throws Exception {
         // create a "default" JSSE X509TrustManager.
         KeyStore ks = KeyStore.getInstance("JKS");
         ks.load(new FileInputStream("trustedCerts"), "passphrase".toCharArray());
         TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
         tmf.init(ks);
         TrustManager tms [] = tmf.getTrustManagers();
          * Iterate over the returned trust managers, looking
          \star for an instance of X509TrustManager. If found,
          * use that as the default trust manager.
          * /
         for (int i = 0; i < tms.length; i++) {
             if (tms[i] instanceof X509TrustManager) {
                 pkixTrustManager = (X509TrustManager) tms[i];
                 return;
          * Find some other way to initialize, or else the
          * constructor fails.
         throw new Exception("Couldn't initialize");
     }
      * Delegate to the default trust manager.
     public void checkClientTrusted(X509Certificate[] chain, String authType)
                 throws CertificateException {
         try {
             pkixTrustManager.checkClientTrusted(chain, authType);
         } catch (CertificateException excep) {
             // do any special handling here, or rethrow exception.
     }
      * Delegate to the default trust manager.
      * /
```

Updating the Keystore Dynamically

You can enhance MyX509TrustManager to handle dynamic keystore updates. When a checkClientTrusted or checkServerTrusted test fails and does not establish a trusted certificate chain, you can add the required trusted certificate to the keystore. You must create a new pkixTrustManager from the TrustManagerFactory initialized with the updated keystore. When you establish a new connection (using the previously initialized SSLContext), the newly added certificate will be used when making trust decisions.

X509ExtendedTrustManager Class

The X509ExtendedTrustManager class is an abstract implementation of the X509TrustManager interface. It adds methods for connection-sensitive trust management. In addition, it enables endpoint verification at the TLS layer.

In TLS 1.2 and later, both client and server can specify which hash and signature algorithms they will accept. To authenticate the remote side, authentication decisions must be based on both X509 certificates and the local accepted hash and signature algorithms. The local accepted hash and signature algorithms can be obtained using the ExtendedSSLSession.getLocalSupportedSignatureAlgorithms() method.

The ExtendedSSLSession object can be retrieved by calling the SSLSocket.getHandshakeSession() method or the SSLEngine.getHandshakeSession() method.

The X509TrustManager interface is not connection-sensitive. It provides no way to access SSLSocket Or SSLEngine session properties.

Besides TLS 1.2 support, the X509ExtendedTrustManager class also supports algorithm constraints and SSL layer host name verification. For JSSE providers and trust manager implementations, the X509ExtendedTrustManager class is highly recommended over the legacy X509TrustManager interface.

Creating an X509ExtendedTrustManager

You can either create an X509ExtendedTrustManager subclass yourself (which is outlined in the following section) or obtain one from a provider-based TrustManagerFactory (such

as that supplied by the SunJSSE provider). In Java SE 7, the PKIX or SunX509 TrustManagerFactory returns an X509ExtendedTrustManager instance.

Creating Your Own X509ExtendedTrustManager

This section outlines how to create a subclass of X509ExtendedTrustManager in nearly the same way as described for X509TrustManager.

Example 8-13 illustrates how to create a class that uses the PKIX TrustManagerFactory to locate a default x509ExtendedTrustManager that will be used to make decisions about trust.

Example 8-13 Sample Code for Creating a PKIX TrustManagerFactory

The following code sample illustrates how to create a class that uses the PKIX <code>TrustManagerFactory</code> to locate a default <code>X509ExtendedTrustManager</code> that will be used to make decisions about trust. If the default trust manager fails for any reason, then the subclass can add other behavior. In the sample, these locations are indicated by comments in the <code>catch</code> clauses.

```
import java.io.*;
import java.net.*;
import java.security.*;
import java.security.cert.*;
import javax.net.ssl.*;
public class MyX509ExtendedTrustManager extends X509ExtendedTrustManager {
  * The default PKIX X509ExtendedTrustManager. Decisions are
   * delegated to it, and a fall back to the logic in this class is
   * performed if the default X509ExtendedTrustManager does not
   * trust it.
  X509ExtendedTrustManager pkixTrustManager;
  MyX509ExtendedTrustManager() throws Exception {
    // create a "default" JSSE X509ExtendedTrustManager.
    KeyStore ks = KeyStore.getInstance("JKS");
    ks.load(new FileInputStream("trustedCerts"), "passphrase".toCharArray());
    TrustManagerFactory tmf = TrustManagerFactory.getInstance("PKIX");
    tmf.init(ks);
    TrustManager tms [] = tmf.getTrustManagers();
    * Iterate over the returned trust managers, looking
     * for an instance of X509ExtendedTrustManager. If found,
     * use that as the default trust manager.
    for (int i = 0; i < tms.length; i++) {</pre>
     if (tms[i] instanceof X509ExtendedTrustManager) {
       pkixTrustManager = (X509ExtendedTrustManager) tms[i];
        return;
    }
```



```
/*
     * Find some other way to initialize, or else we have to fail the
     * constructor.
    throw new Exception("Couldn't initialize");
   \ ^{\star} Delegate to the default trust manager.
 public void checkClientTrusted(X509Certificate[] chain, String authType)
    throws CertificateException {
    try {
      pkixTrustManager.checkClientTrusted(chain, authType);
    } catch (CertificateException excep) {
      \ensuremath{//} do any special handling here, or rethrow exception.
   * Delegate to the default trust manager.
 public void checkServerTrusted(X509Certificate[] chain, String authType)
    throws CertificateException {
    try {
      pkixTrustManager.checkServerTrusted(chain, authType);
    } catch (CertificateException excep) {
       \mbox{\scriptsize *} Possibly pop up a dialog box asking whether to trust the
       * cert chain.
  }
   * Connection-sensitive verification.
   * /
 public void checkClientTrusted(X509Certificate[] chain, String authType, Socket
socket)
    throws CertificateException {
      pkixTrustManager.checkClientTrusted(chain, authType, socket);
    } catch (CertificateException excep) {
      // do any special handling here, or rethrow exception.
  }
 public void checkClientTrusted(X509Certificate[] chain, String authType, SSLEngine
engine)
    throws CertificateException {
    try {
      pkixTrustManager.checkClientTrusted(chain, authType, engine);
    } catch (CertificateException excep) {
      // do any special handling here, or rethrow exception.
 public void checkServerTrusted(X509Certificate[] chain, String authType, Socket
socket)
    throws CertificateException {
    try {
```

```
pkixTrustManager.checkServerTrusted(chain, authType, socket);
} catch (CertificateException excep) {
    // do any special handling here, or rethrow exception.
}

public void checkServerTrusted(X509Certificate[] chain, String authType, SSLEngine engine)
    throws CertificateException {
    try {
        pkixTrustManager.checkServerTrusted(chain, authType, engine);
    } catch (CertificateException excep) {
        // do any special handling here, or rethrow exception.
    }
}

/*
    * Merely pass this through.
    */
    public X509Certificate[] getAcceptedIssuers() {
        return pkixTrustManager.getAcceptedIssuers();
    }
}
```

The KeyManager Interface

The primary responsibility of the <code>KeyManager</code> is to select the authentication credentials that will eventually be sent to the remote host. To authenticate yourself (a local secure socket peer) to a remote peer, you must initialize an <code>SSLContext</code> object with one or more <code>KeyManager</code> objects. You must pass one <code>KeyManager</code> for each different authentication mechanism that will be supported. If null is passed into the <code>SSLContext</code> initialization, then an empty <code>KeyManager</code> will be created. If the internal default context is used (for example, an <code>SSLContext</code> created by <code>SSLSocketFactory.getDefault()</code> or <code>SSLServerSocketFactory.getDefault()</code>), then a default <code>KeyManager</code> is created. See Customizing the Default Keystores and Truststores, Store Types, and Store Passwords. Typically, a single key manager supports authentication based on X.509 public key certificates. Some secure socket implementations may also support authentication based on shared secret keys, Kerberos, or other mechanisms.

KeyManager objects are created either by a KeyManagerFactory, or by providing a concrete implementation of the interface.

The KeyManagerFactory Class

The <code>javax.net.ssl.KeyManagerFactory</code> class is an engine class for a provider-based service that acts as a factory for one or more types of <code>KeyManager</code> objects. The SunJSSE provider implements a factory that can return a basic X.509 key manager. Because it is provider-based, additional factories can be implemented and configured to provide additional or alternative key managers.



Creating a KeyManagerFactory

You create an instance of this class in a similar manner to SSLContext, except for passing an algorithm name string instead of a protocol name to the getInstance() method:

```
KeyManagerFactory kmf = getInstance(String algorithm);
KeyManagerFactory kmf = getInstance(String algorithm, String provider);
KeyManagerFactory kmf = getInstance(String algorithm, Provider provider);
```

A sample call as follows:

```
KeyManagerFactory kmf = KeyManagerFactory.getInstance("SunX509", "SunJSSE");
```

The preceding call creates an instance of the SunJSSE provider's default key manager factory, which provides basic X.509-based authentication keys.

A newly created factory should be initialized by calling one of the init methods:

```
public void init(KeyStore ks, char[] password);
public void init(ManagerFactoryParameters spec);
```

Call whichever init method is appropriate for the KeyManagerFactory you are using. If you are not sure, then ask the provider vendor.

For many factories, such as the default SunX509 <code>KeyManagerFactory</code> from the SunJSSE provider, the <code>KeyStore</code> and password are the only information required to initialize the <code>KeyManagerFactory</code> and thus the first <code>init</code> method is the appropriate one to call. The <code>KeyManagerFactory</code> will query the <code>KeyStore</code> for information about which private key and matching public key certificates should be used for authenticating to a remote socket peer. The password parameter specifies the password that will be used with the methods for accessing keys from the <code>KeyStore</code>. All keys in the <code>KeyStore</code> must be protected by the same password.

Sometimes initialization parameters other than a KeyStore and password are needed by a provider. Users of that provider are expected to pass an implementation of the appropriate ManagerFactoryParameters as defined by the provider. The provider can then call the specified methods in the ManagerFactoryParameters implementation to obtain the needed information.

Some factories can provide access to authentication material without being initialized with a <code>KeyStore</code> object or any other parameters. For example, they may access key material as part of a login mechanism such as one based on JAAS, the Java Authentication and Authorization Service.

As previously indicated, the SunJSSE provider supports a SunX509 factory that must be initialized with a KeyStore parameter.

The X509KeyManager Interface

The <code>javax.net.ss1.X509KeyManager</code> interface extends the general <code>KeyManager</code> interface. It must be implemented by a key manager for X.509-based authentication. To support X. 509 authentication to remote socket peers through JSSE, an instance of this interface must be passed to the <code>init()</code> method of an <code>SSLContext</code> object.



Creating an X509KeyManager

You can either implement this interface directly yourself or obtain one from a provider-based <code>KeyManagerFactory</code> (such as that supplied by the SunJSSE provider). You could also implement your own interface that delegates to a factory-generated key manager. For example, you might do this to filter the resulting keys and query an end-user through a graphical user interface.

Creating Your Own X509KeyManager

If the default x509KeyManager behavior is not suitable for your situation, then you can create your own x509KeyManager in a way similar to that shown in Creating Your Own X509TrustManager.

The X509ExtendedKeyManager Class

The x509ExtendedKeyManager abstract class is an implementation of the x509KeyManager interface that allows for connection-specific key selection. It adds two methods that select a key alias for client or server based on the key type, allowed issuers, and current SSLEngine:

- public String chooseEngineClientAlias(String[] keyType, Principal[] issuers, SSLEngine engine)
- public String chooseEngineServerAlias(String keyType, Principal[] issuers, SSLEngine engine)

If a key manager is not an instance of the X509ExtendedKeyManager class, then it will not work with the SSLEngine class.

For JSSE providers and key manager implementations, the X509ExtendedKeyManager class is highly recommended over the legacy X509KeyManager interface.

In TLS 1.2 and later, both client and server can specify which hash and signature algorithms they will accept. To pass the authentication required by the remote side, local key selection decisions must be based on both X509 certificates and the remote accepted hash and signature algorithms. The remote accepted hash and signature algorithms can be retrieved using the

 ${\tt ExtendedSSLSession.getPeerSupportedSignatureAlgorithms()} \ \ {\tt method}.$

You can create your own X509ExtendedKeyManager subclass in a way similar to that shown in Creating Your Own X509TrustManager.

Support for the Server Name Indication (SNI) Extension on the server side enables the key manager to check the server name and select the appropriate key accordingly. For example, suppose there are three key entries with certificates in the keystore:

- cn=www.example.com
- cn=www.example.org
- cn=www.example.net

If the ClientHello message requests to connect to www.example.net in the SNI extension, then the server should be able to select the certificate with subject cn=www.example.net.



Relationship Between a TrustManager and a KeyManager

Historically, there has been confusion regarding the functionality of a ${\tt TrustManager}$ and a ${\tt KeyManager}$.

A TrustManager determines whether the remote authentication credentials (and thus the connection) should be trusted.

A KeyManager determines which authentication credentials to send to the remote host.

Secondary Support Classes and Interfaces

These classes are provided as part of the JSSE API to support the creation, use, and management of secure sockets. They are less likely to be used by secure socket applications than are the core and support classes. The secondary support classes and interfaces are part of the <code>javax.net.ssl</code> and <code>javax.security.cert</code> packages.

The SSLParameters Class

The SSLParameters class encapsulates the following parameters that affect a SSL/TLS/DTLS connection:

- The list of cipher suites to be accepted in an SSL/TLS/DTLS handshake
- The list of protocols to be allowed
- The endpoint identification algorithm during SSL/TLS/DTLS handshaking
- The server names and server name matchers (see Server Name Indication (SNI) Extension)
- The cipher suite preference to be used in an SSL/TLS/DTLS handshake
- Algorithm during SSL/TLS/DTLS handshaking
- The Server Name Indication (SNI)
- The maximum network packet size
- The algorithm constraints and whether SSL/TLS/DTLS servers should request or require client authentication

You can retrieve the current SSLParameters for an SSLSocket or SSLEngine by using the following methods:

- getSSLParameters() in an SSLSocket, SSLServerSocket, and SSLEngine
- getDefaultSSLParameters() and getSupportedSSLParamters() in an SSLContext

You can assign SSLParameters with the setSSLParameters() method in an SSLSocket, SSLServerSocket and SSLEngine.

You can explicitly set the server name indication with the

SSLParameters.setServerNames() method. The server name indication in client mode also affects endpoint identification. In the implementation of X509ExtendedTrustManager, it uses the server name indication retrieved by the

ExtendedSSLSession.getRequestedServerNames() method. See Example 8-14.



Example 8-14 Sample Code to Set Server Name Indication

This example uses the host name in the server name indication (www.example.com) to make endpoint identification against the peer's identity presented in the end-entity's X. 509 certificate.

```
SSLSocketFactory factory = ...
SSLSocket sslSocket = factory.createSocket("172.16.10.6", 443);
// SSLEngine sslEngine = sslContext.createSSLEngine("172.16.10.6", 443);
SNIHostName serverName = new SNIHostName("www.example.com");
List<SNIServerName> serverNames = new ArrayList<>(1);
serverNames.add(serverName);
SSLParameters params = sslSocket.getSSLParameters();
params.setServerNames(serverNames);
sslSocket.setSSLParameters(params);
// sslEngine.setSSLParameters(params);
```

Cipher Suite Preference

During TLS handshaking, the client requests to negotiate a cipher suite from a list of cryptographic options that it supports, starting with its first preference. Then, the server selects a single cipher suite from the list of cipher suites requested by the client. Normally, the selection honors the client's preference. However, to mitigate the risks of using weak cipher suites, the server may select cipher suites based on its own preference rather than the client's preference, by invoking the method SSLParameters.setUseCipherSuitesOrder(true).

The SSLSessionContext Interface

The <code>javax.net.ssl.SSLSessionContext</code> interface is a grouping of <code>SSLSession</code> objects associated with a single entity. For example, it could be associated with a server or client that participates in many sessions concurrently. The methods in this interface enable the enumeration of all sessions in a context and allow lookup of specific sessions via their session IDs.

An SSLSessionContext may optionally be obtained from an SSLSession by calling the SSLSession getSessionContext() method. The context may be unavailable in some environments, in which case the getSessionContext() method returns null.

The SSLSessionBindingListener Interface

The <code>javax.net.ssl.SSLSessionBindingListener</code> interface is implemented by objects that are notified when they are being bound or unbound from an <code>SSLSession</code>.

The SSLSessionBindingEvent Class

The <code>javax.net.ssl.SSLSessionBindingEvent</code> class defines the event communicated to an <code>SSLSessionBindingListener</code> (see The <code>SSLSessionBindingListener</code> Interface) when it is bound or unbound from an <code>SSLSession</code> (see <code>SSLSession</code> and <code>ExtendedSSLSession</code>).



The HandShakeCompletedListener Interface

The <code>javax.net.ssl.HandShakeCompletedListener</code> interface is an interface implemented by any class that is notified of the completion of an SSL protocol handshake on a given <code>SSLSocket</code> connection.

The HandShakeCompletedEvent Class

The <code>javax.net.ssl.HandShakeCompletedEvent</code> class defines the event communicated to a <code>HandShakeCompletedListener</code> (see The <code>HandShakeCompletedListener</code> Interface) upon completion of an SSL protocol handshake on a given <code>sslsocket</code> connection.

The HostnameVerifier Interface

If the SSL/TLS implementation's standard host name verification logic fails, then the implementation calls the $\mathtt{verify}()$ method of the class that implements this interface and is assigned to this $\mathtt{Httpsurlconnection}$ instance. If the callback class can determine that the host name is acceptable given the parameters, it reports that the connection should be allowed. An unacceptable response causes the connection to be terminated. See Example 8-15.

See httpsurlConnection for more information about how to assign the HostnameVerifier to the httpsurlConnection.

Example 8-15 Sample Code for Implementing the HostnameVerifier Interface

The following example illustrates a class that implements HostnameVerifier interface:

```
public class MyHostnameVerifier implements HostnameVerifier {
    public boolean verify(String hostname, SSLSession session) {
        // pop up an interactive dialog box
        // or insert additional matching logic
        if (good_address) {
            return true;
        } else {
               return false;
        }
    }
}

//...deleted...

HttpsURLConnection urlc = (HttpsURLConnection)
    (new URL("https://www.example.com/")).openConnection();
urlc.setHostnameVerifier(new MyHostnameVerifier());
```

The X509Certificate Class

Many secure socket protocols perform authentication using public key certificates, also called X.509 certificates. This is the default authentication mechanism for the SSL/TLS protocols.

The java.security.cert.X509Certificate abstract class provides a standard way to access the attributes of X.509 certificates.

✓ Note:

The <code>javax.security.cert.X509Certificate</code> class is supported only for backward compatibility with previous (1.0.x and 1.1.x) versions of JSSE. New applications should use the <code>java.security.cert.X509Certificate</code> class instead.

The AlgorithmConstraints Interface

The java.security.AlgorithmConstraints interface is used for controlling allowed cryptographic algorithms. AlgorithmConstraints defines three permits() methods. These methods tell whether an algorithm name or a key is permitted for certain cryptographic functions. Cryptographic functions are represented by a set of CryptoPrimitive, which is an enumeration containing fields like STREAM_CIPHER, MESSAGE_DIGEST, and SIGNATURE.

Thus, an AlgorithmConstraints implementation can answer questions like: Can I use this key with this algorithm for the purpose of a cryptographic operation?

An AlgorithmConstraints Object can be associated with an SSLParameters object by using the new setAlgorithmConstraints() method. The current AlgorithmConstraints Object for an SSLParameters Object is retrieved using the getAlgorithmConstraints() method.

The StandardConstants Class

The StandardConstants class is used to represent standard constants definitions in JSSE.

StandardConstants.SNI_HOST_NAME represents a domain name server (DNS) host name in a Server Name Indication (SNI) extension, which can be used when instantiating an SNIServerName Of SNIMatcher Object.

The SNIServerName Class

An instance of the abstract SNIServerName class represents a server name in the Server Name Indication (SNI) extension. It is instantiated using the type and encoded value of the specified server name.

You can use the $\mathtt{getType}()$ and $\mathtt{getEncoded}()$ methods to return the server name type and a copy of the encoded server name value, respectively. The $\mathtt{equals}()$ method can be used to check if some other object is "equal" to this server name. The $\mathtt{hashCode}()$ method returns a hash code value for this server name. To get a string representation of the server name (including the server name type and encoded server name value), use the $\mathtt{toString}()$ method.

The SNIMatcher Class

An instance of the abstract SNIMatcher class performs match operations on an SNIServerName object. Servers can use information from the Server Name Indication (SNI) extension to decide if a specific SSLSocket or SSLEngine should accept a connection. For example, when multiple "virtual" or "name-based" servers are hosted on a single underlying network address, the server application can use SNI information



to determine whether this server is the exact server that the client wants to access. Instances of this class can be used by a server to verify the acceptable server names of a particular type, such as host names.

The SNIMatcher class is instantiated using the specified server name type on which match operations will be performed. To match a given SNIServerName, use the matches() method. To return the server name type of the given SNIMatcher object, use the getType() method.

The SNIHostName Class

An instance of the SNIHOSTNAME class (which extends the SNIServerName class) represents a server name of type "host_name" (see The StandardConstants Class) in the Server Name Indication (SNI) Extension. To instantiate an SNIHOSTNAME, specify the fully qualified DNS host name of the server (as understood by the client) as a String argument. The argument is illegal in the following cases:

- The argument is empty.
- The argument ends with a trailing period.
- The argument is not a valid Internationalized Domain Name (IDN) compliant with the RFC 3490 specification.

You can also instantiate an SNIHOSTNAME by specifying the encoded host name value as a byte array. This method is typically used to parse the encoded name value in a requested SNI extension. Otherwise, use the SNIHOSTNAME(String hostname) constructor. The encoded argument is illegal in the following cases:

- The argument is empty.
- The argument ends with a trailing period.
- The argument is not a valid Internationalized Domain Name (IDN) compliant with the RFC 3490 specification.
- The argument is not encoded in UTF-8 or US-ASCII.



The encoded byte array passed in as an argument is cloned to protect against subsequent modification.

To return the host name of an SNIHOSTName object in US-ASCII encoding, use the getAsciiName() method. To compare a server name to another object, use the equals() method (comparison is *not* case-sensitive). To return a hash code value of an SNIHOSTNAME, use the hashCode() method. To return a string representation of an SNIHOSTNAME, including the DNS host name, use the toString() method.

You can create an SNIMatcher object for an SNIHostName object by passing a regular expression representing one or more host names to match to the createSNIMatcher() method.



Customizing JSSE

JSSE includes a standard implementation that can be customized by plugging in different implementations or specifying the default keystore, and so on.

Table 8-7 and Table 8-8 summarize which aspects can be customized, what the defaults are, and which mechanisms are used to provide customization.

Some of the customizations are done by setting system property or Security Property values. Sections following the table explain how to set such property values.



Many of the properties shown in this table are currently used by the JSSE implementation, but there is no guarantee that they will continue to have the same names and types (system or security) or even that they will exist at all in future releases. All such properties are flagged with an asterisk (*). They are documented here for your convenience for use with the JSSE implementation.

Table 8-7 shows items that are customized by setting the <code>java.security.Security</code> property. See How to Specify a <code>java.security.Security Property</code>

Table 8-7 Security Properties and Customized Items

Security Property	Customized Item	Default Value	Notes
cert.provider.x509v1	Customizing the X509Certificate Implementation	X509Certificate implementation from Oracle	None
security.provider.n	Cryptographic service provider; see Customizing	The first five providers in order of priority are:	Specify the provider in the security.provider.n= line
	the Provider Implementation and	1. SUN	in security properties file, where n is an integer
	Customizing the Encryption	2. SunRsaSign	whose value is equal or
	Algorithm Providers	3. SunEC	greater than 1.
		4. SunJSSE	
		5. SunJCE	
*ssl.SocketFactory.provider	Default SSLSocketFactory implementation	SSLSocketFactory implementation from Oracle	None
*ssl.ServerSocketFactory .provider	Default SSLServerSocketFactory implementation	SSLServerSocketFactory implementation from Oracle	None
ssl.KeyManagerFactory.a lgorithm	Default key manager factory algorithm name (see Customizing the Default Key Managers and Trust Managers)	SunX509	None



Table 8-7 (Cont.) Security Properties and Customized Items

Security Property	Customized Item	Default Value	Notes
*jdk.certpath.disabledAl gorithms	Disabled certificate verification cryptographic algorithm (see Disabled and Restricted Cryptographic Algorithms)	MD2, MD5, SHA1 jdkCA & usage TLSServer, RSA keySize < 1024, DSA keySize < 1024, EC keySize < 224 ⁴	None
ssl.TrustManagerFactory .algorithm	Default trust manager factory algorithm name (see Customizing the Default Key Managers and Trust Managers)	PKIX	None
JCE encryption algorithms used by the SunJSSE provider	Give alternative JCE algorithm providers a higher preference order than the SunJCE provider	SunJCE implementations	None
*jdk.tls.disabledAlgorit	Disabled and Restricted Cryptographic Algorithms	SSLv3, RC4, MD5withRSA, DH keySize < 1024, EC keySize < 224 ⁴	Disables specific algorithms (protocols versions, cipher suites, key exchange mechanisms, etc.) that will not be negotiated for SSL/TLS/DTLS connections, even if they are enabled explicitly in an application
jdk.tls.server.defaultD HEParameters	Diffie-Hellman groups	Safe prime Diffie-Hellman groups in OpenJDK SSL/TLS/DTLS implementation	Defines default finite field Diffie-Hellman ephemeral (DHE) parameters for Transport Layer Security (SSL/TLS/DTLS) processing

^{*} This property is currently used by the JSSE implementation, but it is not guaranteed to be examined and used by other implementations. If it *is* examined by another implementation, then that implementation should handle it in the same manner as the JSSE implementation does. There is no guarantee the property will continue to exist or be of the same type (system or security) in future releases.

1

Table 8-8 shows items that are customized by setting <code>java.lang.System</code> property. See How to Specify a <code>java.lang.System</code> Property.

Table 8-8 System Properties and Customized Items

System Property	Customized Item	Default	Notes
java.protocol.handl er.pkgs	Specifying an Alternative HTTPS Protocol Implementation	Implementation from Oracle	None

¹ The list of restricted algorithms specified in these Security Properties may change; see the java.security file in your JDK installation for the latest values.



Table 8-8 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
*javax.net.ssl.keyS tore	Default keystore (see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords)	None	The value NONE may be specified. This setting is appropriate if the keystore is not file-based (for example, it resides in a hardware token)
*javax.net.ssl.keyS torePassword	Default keystore password (see Customizing the Default Keystores and Truststores, Store Types, and Store	None	It is inadvisable to specify the password in a way that exposes it to discovery by other users.
	Passwords)		For example, specifying the password on the command line. To keep the password secure, have the application prompt for the password, or specify the password in a properly protected option file
*javax.net.ssl.keyS toreProvider	Default keystore provider (see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords)	None	None
*javax.net.ssl.keyS toreType	Default keystore type (see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords)	<pre>KeyStore.getDefault Type()</pre>	None
*javax.net.ssl.trus tStore	Default truststore (see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords)	jssecacerts, if it exists. Otherwise, cacerts	None



Table 8-8 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
*javax.net.ssl.trus tStorePassword	Default truststore password (see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords)	None	It is inadvisable to specify the password in a way that exposes it to discovery by other users. For example, specifying the password on the command line. To keep the password secure, have the application prompt for the password, or specify the password in a properly protected option file
*javax.net.ssl.trus tStoreProvider	Default truststore provider (see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords)	None	None
*javax.net.ssl.trus tStoreType	Default truststore type (see Customizing the Default Keystores and Truststores, Store Types, and Store Passwords)	<pre>KeyStore.getDefault Type()</pre>	The value NONE may be specified. This setting is appropriate if the truststore is not file-based (for example, it resides in a hardware token)
*https.proxyHost	Default proxy host	None	None
*https.proxyPort	Default proxy port	80	None
*jsse.enableSNIExte nsion	Server Name Indication option	true	Server Name Indication (SNI) is a TLS extension, defined in RFC 6066. It enables TLS connections to virtual servers, in which multiple servers for different network names are hosted at a single underlying network address. Some very old SSL/TLS vendors may not be able handle SSL/TLS extensions. In this case, set this property to false to disable the SNI extension



Table 8-8 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
*https.cipherSuites	Default cipher suites	Determined by the socket factory.	This contains a comma-separated list of cipher suite names specifying which cipher suites to enable for use on this HttpsURLConnection. See the SSLSocket.setEnable dCipherSuites(String[])
*https.protocols	Default handshaking protocols	Determined by the socket factory.	This contains a comma-separated list of protocol suite names specifying which protocol suites to enable on this HttpsURLConnection. See SSLSocket.setEnable dProtocols(String[])
* Customize via port field in the HTTPS URL.	Default HTTPS port	443	None
*jsse.SSLEngine.acc eptLargeFragments	Default sizing buffers for large SSL/TLS packets	None	Setting this system property to true, SSLSession will size buffers to handle large data packets by default. This may cause applications to allocate unnecessarily large SSLEngine buffers. Instead, applications should dynamically check for buffer overflow conditions and resize buffers as appropriate
*sun.security.ssl.a llowUnsafeRenegotia tion	Allow unsafe SSL/TLS Renegotaions (see Description of the Phase 2 Fix)	false	Setting this system property to true permits full (unsafe) legacy renegotiation. This system property is deprecated and might be removed in a future JDK release.



Table 8-8 (Cont.) System Properties and Customized Items

System Property	Customized Item	Default	Notes
*sun.security.ssl.a llowLegacyHelloMess ages	Allow legacy Hello messages (see Description of the Phase 2 Fix)	true	Setting this system property to true allows the peer to handshake without requiring the proper RFC 5746 messages.
			This system property is deprecated and might be removed in a future JDK release.
jdk.tls.client.prot ocols	The SunJSSE Provider	None	To enable specific SunJSSE protocols on the client, specify them in a comma- separated list within quotation marks; all other supported protocols are not enabled on the client For example, If jdk.tls.client. protocols="TLSV 1,TLSV1.1", then the default protocol settings on the client for TLSV1 and TLSV1.1 are enabled, while SSLV3, TLSV1.2, and SSLv2Hello are not enabled If jdk.tls.client. protocols="DTLS v1.2", then the protocol setting on the client for DTLS1.2 is enabled, while DTLS1.0 is not enabled
jdk.tls.ephemeralDH KeySize	Customizing Size of Ephemeral Diffie- Hellman Keys	1024 bits	None
jsse.enableMFLNExte nsion	Customizing Maximum Fragment Length Negotiation (MFLN) Extension	false	None

 $^{^{\}star}$ This property is currently used by the JSSE implementation, but it is not guaranteed to be examined and used by other implementations. If it *is* examined by another



implementation, then that implementation should handle it in the same manner as the JSSE implementation does. There is no guarantee the property will continue to exist or be of the same type (system or security) in future releases.

How to Specify a java.lang.System Property

You can customize some aspects of JSSE by setting system properties. There are several ways to set these properties:

To set a system property statically, use the -D option of the java command. For example, to run an application named MyApp and set the javax.net.ssl.trustStore system property to specify a truststore named MyCacertsFile. See truststore. Enter the following:

```
java -Djavax.net.ssl.trustStore=MyCacertsFile MyApp
```

To set a system property dynamically, call the java.lang.System.setProperty()
method in your code:

```
System.setProperty("propertyName", "propertyValue");
```

For example, a <code>setProperty()</code> call corresponding to the previous example for setting the <code>javax.net.ssl.trustStore</code> system property to specify a truststore named <code>"MyCacertsFile"</code> would be:

```
System.setProperty("javax.net.ssl.trustStore", "MyCacertsFile");
```

- In the Java Deployment environment (Plug-In/Web Start), there are several ways to set the system properties.
 - Use the Java Control Panel to set the Runtime Environment Property on a local or per-VM basis. This creates a local deployment.properties file.
 Deployers can also distribute an enterprise wide deployment.properties file by using the deployment.config mechanism.
 - To set a property for a specific applet, use the HTML subtag <param>
 "java_arguments" within the <applet> tag.
 - To set the property in a specific Java Web Start application or applet using Plugin2, use the JNLP property sub element of the resources element. See resources Element in the Java Platform, Standard Edition Deployment Guide.

How to Specify a java.security.Security Property

You can customize some aspects of JSSE by setting Security Properties. You can set a Security Property either statically or dynamically:

 To set a Security Property statically, add a line to the security properties file. The security properties file is located at java-home/conf/security/ java.security



java-home

See Terms and Definitions

To specify a Security Property value in the security properties file, you add a line of the following form:

propertyName=propertyValue

For example, suppose that you want to specify a different key manager factory algorithm name than the default SunX509. You do this by specifying the algorithm name as the value of a Security Property named ssl.KeyManagerFactory.algorithm. For example, to set the value to MyX509, add the following line to the security properties file:

ssl.KeyManagerFactory.algorithm=MyX509

• To set a Security Property dynamically, call the java.security.Security.setProperty method in your code:

```
Security.setProperty("propertyName," "propertyValue");
```

For example, a call to the <code>setProperty()</code> method corresponding to the previous example for specifying the key manager factory algorithm name would be:

Security.setProperty("ssl.KeyManagerFactory.algorithm", "MyX509");

Customizing the X509Certificate Implementation

The X509Certificate implementation returned by the $\tt X509Certificate.getInstance()$ method is by default the implementation from the JSSE implementation.

To cause a different implementation to be returned:

Specify the name (and package) of the other implementation's class as the value of a How to Specify a java.security.Security Property named cert.provider.x509v1.

MyX509CertificateImplcom.cryptox

cert.provider.x509v1=com.cryptox.MyX509CertificateImpl

Specifying an Alternative HTTPS Protocol Implementation

You can communicate securely with an SSL-enabled web server by using the HTTPS URL scheme for the java.net.URL class. The JDK provides a default HTTPS URL implementation.

If you want an alternative HTTPS protocol implementation to be used, set the <code>java.protocol.handler.pkgs</code> How to Specify a java.lang.System Property to include the new class name. This action causes the specified classes to be found and loaded before the JDK default classes. See the <code>URL</code> class for details.



In past JSSE releases, you had to set the <code>java.protocol.handler.pkgs</code> system property during JSSE installation. This step is no longer required unless you want to obtain an instance of <code>com.sun.net.ssl.HttpsURLConnection</code>.



Customizing the Provider Implementation

The JDK comes with a JSSE Cryptographic Service Provider, or *provider* for short, named SunJSSE. Providers are essentially packages that implement one or more engine classes for specific cryptographic algorithms.

The JSSE engine classes are SSLContext, KeyManagerFactory, and TrustManagerFactory. See Java Cryptography Architecture (JCA) Reference Guide to know more about providers and engine classes.

Before it can be used, a provider must be registered, either statically or dynamically. You do not need to register the SunJSSE provider because it is preregistered. If you want to use other providers, read the following sections to see how to register them.

Registering the Cryptographic Provider Statically

Register a provider statically by adding a line of the following form to the security properties file, <java-home>/conf/security/java.security:

security.provider.n=provName | className

This declares a provider, and specifies its preference order n. The preference order is the order in which providers are searched for requested algorithms when no specific provider is requested. The order is 1-based; 1 is the most preferred, followed by 2, and so on.

provName is the provider's name and className is the fully qualified class name of the provider.

Standard security providers are automatically registered for you in the <code>java.security</code> security properties file.

To use another JSSE provider, add a line registering the other provider, giving it whatever preference order you prefer.

You can have more than one JSSE provider registered at the same time. The registered providers may include different implementations for different algorithms for different engine classes, or they may have support for some or all of the same types of algorithms and engine classes. When a particular engine class implementation for a particular algorithm is searched for, if no specific provider is specified for the search, then the providers are searched in preference order and the implementation from the first provider that supplies an implementation for the specified algorithm is used.

See Step 8.1: Configure the Provider in Steps to Implement and Integrate a Provider.

Registering the Cryptographic Service Provider Dynamically

Instead of registering a provider statically, you can add the provider dynamically at runtime by calling either the addProvider or insertProviderAt method in the Security class. Note that this type of registration is not persistent and can only be done by code which is granted the insertProvider.cyprovider name> permission.

See Step 8.1: Configure the Provider in Steps to Implement and Integrate a Provider.



Provider Configuration

Some providers may require configuration. This is done using the <code>configure</code> method of the <code>Provider</code> class, prior to calling the <code>addProvider</code> method of the <code>Security</code> class. See <code>SunPKCS11</code> Configuration for an example. The <code>Provider.configure()</code> method is new to Java SE 9.

Configuring the Preferred Provider for Specific Algorithms

Specify the preferred provider for a specific algorithm in the <code>jdk.security.provider.preferred</code> Security Property. By specifying a preferred provider you can configure providers that offer performance gains for specific algorithms but are not the best performing provider for other algorithms. The ordered provider list specified using the <code>security.provider.n</code> property is not sufficient to order providers that offer performance gains for specific algorithms but are not the best performing provider for other algorithms. More flexibility is required for configuring the ordering of provider list to achieve performance gains.

The jdk.security.provider.preferred Security Property allows specific algorithms, or service types to be selected from a preferred set of providers before accessing the list of registered providers. See How to Specify a java.security.Security Property.

The jdk.security.provider.preferred Security Property does not register the providers. The ordered provider list must be Registering the Cryptographic Provider Statically using the security.provider.n property. Any provider that is not registered is ignored.

Specifying the Preferred Provider for an Algorithm

The syntax for specifying the preferred providers string in the jdk.security.provider.preferred Security Property is a comma-separated list of ServiceType.Algorithm:Provider

In this syntax:

ServiceType

The name of the service type. (for example: "MessageDigest")ServiceType is optional. If it isn't specified, the algorithm applies to all service types.

Algorithm

The standard algorithm name. See Java Security Standard Algorithm Names Specification. Algorithms can be specified as full standard name, (AES/CBC/PKCS5Padding) or as partial (AES, AES/CBC, AES//PKCS5Padding).

Provider

The name of the provider. Any provider that isn't listed in the registered list will be ignored. See JDK Providers.

Entries containing errors such as parsing errors are ignored. Use the command <code>java</code> - <code>Djava.security.debug=jca</code> to debug errors.



Preferred Providers and FIPS

If you add a FIPS provider to the <code>security.provider.n</code> property, and specify the preferred provider ordering in the <code>jdk.security.provider.preferred</code> property then the preferred provider specified in <code>jdk.security.provider.preferred</code> are selected first.

Hence, it is recommended that you don't configure jdk.security.provider.preferred property for FIPS provider configurations.

jdk.security.provider.preferred Default Values

The jdk.security.provider.preferred property is not set by default and is used only for application performance tuning.

Example 8-16 Sample jdk.security.provider.preferred Property

The syntax for specifying the jdk.security.provider.preferred property is as follows:

jdk.security.provider.preferred=AES/GCM/NoPadding:SunJCE,
MessageDigest.SHA-256:SUN

In this syntax:

ServiceType

MessageDigest

Algorithm

AES/GCM/NoPadding, SHA-256

Provider

SunJCE, SUN

Customizing the Default Keystores and Truststores, Store Types, and Store Passwords

Whenever a default SSLSocketFactory or SSLServerSocketFactory is created (via a call to SSLSocketFactory.getDefault or SSLServerSocketFactory.getDefault), and this default SSLSocketFactory (or SSLServerSocketFactory) comes from the JSSE reference implementation, a default SSLContext is associated with the socket factory. (The default socket factory will come from the JSSE implementation.)

This default SSLContext is initialized with a default KeyManager and a default TrustManager. If a keystore is specified by the javax.net.ssl.keyStore system property and an appropriate javax.net.ssl.keyStorePassword system property (see How to Specify a java.lang.System Property), then the KeyManager created by the default SSLContext will be a KeyManager implementation for managing the specified keystore. (The actual implementation will be as specified in Customizing the Default Key Managers and Trust Managers.) If no such system property is specified, then the keystore managed by the KeyManager will be a new empty keystore.

Generally, the peer acting as the server in the handshake will need a keystore for its KeyManager in order to obtain credentials for authentication to the client. However, if one of the anonymous cipher suites is selected, then the server's <code>KeyManager</code> keystore is not necessary. And, unless the server requires client authentication, the peer acting



as the client does not need a KeyManager keystore. Thus, in these situations it may be OK if no <code>javax.net.ssl.keyStore</code> system property value is defined.

Similarly, if a truststore is specified by the <code>javax.net.ssl.truststore</code> system property, then the <code>TrustManager</code> created by the default <code>SSLContext</code> will be a <code>TrustManager</code> implementation for managing the specified truststore. In this case, if such a property exists but the file it specifies does not, then no truststore is used. If no <code>javax.net.ssl.trustStore</code> property exists, then a default truststore is searched for. If a truststore named <code>java-home/lib/security/jssecacerts</code> is found, it is used. If not, then a truststore named <code>java-home/lib/security/cacerts</code> is searched for and used (if it exists). Finally, if a truststore is still not found, then the truststore managed by the <code>TrustManager</code> will be a new empty truststore.

Note:

The JDK ships with a limited number of trusted root certificates in the <code>java-home/lib/security/cacerts</code> file. As documented in keytool in <code>Java Platform</code>, <code>Standard Edition Tools Reference</code>, it is your responsibility to maintain (that is, add and remove) the certificates contained in this file if you use this file as a truststore.

Depending on the certificate configuration of the servers that you contact, you may need to add additional root certificates. Obtain the needed specific root certificates from the appropriate vendor.

If the <code>javax.net.ssl.keyStoreType</code> and/or <code>javax.net.ssl.keyStorePassword</code> system properties are also specified, then they are treated as the default <code>KeyManager</code> keystore type and password, respectively. If no type is specified, then the default type is that returned by the <code>KeyStore.getDefaultType()</code> method, which is the value of the <code>keystore.type</code> Security Property, or "jks" if no such Security Property is specified. If no keystore password is specified, then it is assumed to be a blank string "".

Similarly, if the <code>javax.net.ssl.trustStoreType</code> and/or <code>javax.net.ssl.trustStorePassword</code> system properties are also specified, then they are treated as the default truststore type and password, respectively. If no type is specified, then the default type is that returned by the <code>KeyStore.getDefaultType()</code> method. If no truststore password is specified, then it is assumed to be a blank string "".

Note:

This section describes the current JSSE reference implementation behavior. The system properties described in this section are not guaranteed to continue to have the same names and types (system or security) or even to exist at all in future releases. They are also not guaranteed to be examined and used by any other JSSE implementations. If they *are* examined by an implementation, then that implementation should handle them in the same manner as the JSSE reference implementation does, as described herein.



Customizing the Default Key Managers and Trust Managers

As noted in Customizing the Default Keystores and Truststores, Store Types, and Store Passwords, whenever a default SSLSocketFactory Or SSLServerSocketFactory is created, and this default SSLSocketFactory (Or SSLServerSocketFactory) comes from the JSSE reference implementation, a default SSLContext is associated with the socket factory.

This default SSLContext is initialized with a KeyManager and a TrustManager. The KeyManager and/or TrustManager supplied to the default SSLContext will be an implementation for managing the specified keystore or truststore, as described in the aforementioned section.

The <code>KeyManager</code> implementation chosen is determined by first examining the <code>ssl.KeyManagerFactory.algorithm</code> Security Property. If such a property value is specified, then a <code>KeyManagerFactory</code> implementation for the specified algorithm is searched for. The implementation from the first provider that supplies an implementation is used. Its <code>getKeyManagers()</code> method is called to determine the <code>KeyManager</code> to supply to the default <code>sslContext</code>. Technically, <code>getKeyManagers()</code> returns an array of <code>KeyManager</code> objects, one <code>KeyManager</code> for each type of key material. If no such Security Property value is specified, then the default value of <code>SunX509</code> is used to perform the search.



A KeyManagerFactory implementation for the SunX509 algorithm is supplied by the SunJSSE provider. The KeyManager that it specifies is a javax.net.ssl.X509KeyManager implementation.

Similarly, the TrustManager implementation chosen is determined by first examining the ssl.TrustManagerFactory.algorithm Security Property. If such a property value is specified, then a TrustManagerFactory implementation for the specified algorithm is searched for. The implementation from the first provider that supplies an implementation is used. Its getTrustManagers() method is called to determine the TrustManager to supply to the default SSLContext. Technically, getTrustManagers() returns an array of TrustManager objects, one TrustManager for each type of trust material. If no such Security Property value is specified, then the default value of PKIX is used to perform the search.

Note:

A TrustManagerFactory implementation for the PKIX algorithm is supplied by the SunJSSE provider. The TrustManager that it specifies is a javax.net.ssl.X509TrustManager implementation.



Note:

This section describes the current JSSE reference implementation behavior. The system properties described in this section are not guaranteed to continue to have the same names and types (system or security) or even to exist at all in future releases. They are also not guaranteed to be examined and used by any other JSSE implementations. If they *are* examined by an implementation, then that implementation should handle them in the same manner as the JSSE reference implementation does, as described herein.

Disabled and Restricted Cryptographic Algorithms

In some environments, certain algorithms or key lengths may be undesirable when using SSL/TLS/DTLS. The Oracle JDK uses the <code>jdk.certpath.disabledAlgorithms</code> and <code>jdk.tls.disabledAlgorithm</code> Security Properties to disable algorithms during SSL/TLS/DTLS protocol negotiation, including version negotiation, cipher suites selection, peer authentication, and key exchange mechanisms. Note that these Security Properties are not guaranteed to be used by other JDK implementations. See the <code><java-home>/conf/security/java.security</code> file for information about the syntax of these Security Properties and their current active values.

• jdk.certpath.disabledAlgorithms Property: CertPath code uses the jdk.certpath.disabledAlgorithms Security Property to determine which algorithms should not be allowed during CertPath checking. For example, when a TLS Server sends an identifying certificate chain, a client TrustManager that uses a CertPath implementation to verify the received chain will not allow the stated conditions. For example, the following line blocks any MD2-based certificate, as well as SHA1 TLSServer certificates that chain to trust anchors that are preinstalled in the cacaerts keystore. Likewise, this line blocks any RSA key less than 1024 bits.

jdk.certpath.disabledAlgorithms=MD2, SHA1 jdkCA & usage TLSServer, RSA keySize <

• jdk.tls.disabledAlgorithms Property: SunJSSE code uses the jdk.tls.disabledAlgorithms Security Property to disable SSL/TLS/DTLS protocols, cipher suites, keys, and so on. The syntax is similar to the jdk.certpath.disabledAlgorithms Security Property. For example, the following line disables the SSLv3 algorithm and all of the TLS * RC4 * cipher suites:

```
jdk.tls.disabledAlgorithms=SSLv3, RC4
```

If you require a particular condition, you can reactivate it by either removing the associated value in the Security Property in the <code>java.security</code> file or dynamically setting the proper Security Property before JSSE is initialized.

Note that these Security Properties effectively create a third set of cipher suites, Disabled. The following list describes these three sets:

- **Disabled**: If a cipher suite contains any components (for example, RC4) on the disabled list (for example, RC4 is specified in the jdk.tls.disabledAlgorithms Security Property), then that cipher suite is disabled and will **not** be considered for a connection handshake.
- Enabled: A list of specific cipher suites that will be considered for a connection.



• **Not Enabled**: A list of non-disabled cipher suites that will **not** be considered for a connection. To re-enable these cipher suites, call the appropriate setEnabledCipherSuites() or setSSLParameters() methods.

Customizing the Encryption Algorithm Providers

The SunJSSE provider uses the SunJCE implementation for all its cryptographic needs. Although it is recommended that you leave the provider at its regular position, you can use implementations from other JCA or JCE providers by registering them *before* the SunJCE provider.

The standard JCA mechanism (see How Provider Implementations Are Requested and Supplied) can be used to configure providers, either statically via the security properties file /conf/security/java.security, or dynamically via the addProvider() or insertProviderAt() method in the java.security.Security class.

Customizing Size of Ephemeral Diffie-Hellman Keys

In SSL/TLS/DTLS connections, ephemeral Diffie-Hellman (DH) keys may be used internally during the handshaking. The SunJSSE provider provides a flexible approach to customize the strength of the ephemeral DH key size during SSL/TLS/DTLS handshaking.

Diffie-Hellman (DH) keys of sizes less than 1024 bits have been deprecated because of their insufficient strength. You can customize the ephemeral DH key size with the system property jdk.tls.ephemeralDHKeySize. This system property does not impact DH key sizes in ServerKeyExchange messages for exportable cipher suites. It impacts only the DHE_RSA, DHE_DSS, and DH_anon-based cipher suites in the JSSE Oracle provider.

You can specify one of the following values for this property:

- Undefined: A DH key of size 1024 bits will be used always for non-exportable cipher suites. This is the default value for this property.
- legacy: The JSSE Oracle provider preserves the legacy behavior (for example, using ephemeral DH keys of sizes 512 bits and 768 bits) of JDK 7 and earlier releases.
- matched: For non-exportable anonymous cipher suites, the DH key size in ServerKeyExchange messages is 1024 bits. For X.509 certificate based authentication (of non-exportable cipher suites), the DH key size matching the corresponding authentication key is used, except that the size must be between 1024 bits and 2048 bits. For example, if the public key size of an authentication certificate is 2048 bits, then the ephemeral DH key size should be 2048 bits unless the cipher suite is exportable. This key sizing scheme keeps the cryptographic strength consistent between authentication keys and key-exchange keys.
- A valid integer between 1024 and 2048, inclusively: A fixed ephemeral DH key size of the specified value, in bits, will be used for non-exportable cipher suites.

The following table summaries the minimum and maximum acceptable DH key sizes for each of the possible values for the system property jdk.tls.ephemeralDHKeySize:



Table 8-9 DH Key Sizes for the System Property jdk.tls.ephemeralDHKeySize

Value of jdk.tls.ephemer alDHKeySize	Undefined	legacy	matched	Integer value (fixed)
Exportable DH key size	512	512	512	512
Non-exportable anonymous cipher suites	1024	768	1024	The fixed key size is specified by a valid integer property value, which must be between 1024 and 2048, inclusively.
Authentication certificate	1024	768	The key size is the same as the authentication certificate, but must be between 1024 bits and 2048 bits, inclusively. However, the only DH key size that the SunJCE provider supports that is larger than 1024 bits is 2048 bits. Consequently, you may use the values 1024 or 2048 only.	The fixed key size is specified by a valid integer property value, which must be between 1024 and 2048, inclusively.

Customizing Maximum Fragment Length Negotiation (MFLN) Extension

In order to negotiate smaller maximum fragment lengths, clients have an option to include an extension of type $max_fragment_length$ in the ClientHello message. A system property jsse.enableMFLNExtension, can be used to enable or disable the MFLN extension for SSL/TLS/DTLS.

Maximum Fragment Length Negotiation

It may be desirable for constrained SSL/TLS/DTLS clients to negotiate a smaller maximum fragment length due to memory limitations or bandwidth limitations. In order to negotiate smaller maximum fragment lengths, clients have an option to include an extension of type max_fragment_length in the (extended) ClientHello message. See RFC 6066.

Once a maximum fragment length has been successfully negotiated, the SSL/TLS/DTLS client and server can immediately begin fragmenting messages (including



handshake messages) to ensure that no fragment larger than the negotiated length is sent.

System Property jsse.enableMFLNExtension

A system property <code>jsse.enableMFLNExtension</code> is defined to enable or disable the MFLN extension. The <code>jsse.enableMFLNExtension</code> is disabled by default.

The value of the system property can be set as follows:

Table 8-10 jsse.enableMFLNExtension system property

System Property	Description
jsse.enableMFLNExtension=true	Enable the MFLN extension. If the returned value of
	SSLParameters.getMaximumPacketSize() is less than (2^12 + header-size) the maximum fragment length negotiation extension would be enabled.
jsse.enableMFLNExtension=false	Disable the MFLN extension.

Configuring the Maximum and Minimum Packet Size

Set the maximum expected network packet size in bytes for a SSL/TLS/DTLS record with the SSLParameters.setMaximumPacketSize method.

It is recommended that the packet size should not be less than 256 bytes so that small handshake messages, such as HelloVerifyRequests, are not fragmented.

Transport Layer Security (TLS) Renegotiation Issue

In the fall of 2009, a flaw was discovered in the SSL/TLS protocols. A fix to the protocol was developed by the IETF TLS Working Group, and current versions of the JDK contain this fix. This section describes the situation in much more detail, along with interoperability issues when communicating with older implementations that do not contain this protocol fix.

The vulnerability allowed for man-in-the-middle (MITM) attacks where chosen plain text could be injected as a prefix to a TLS connection. This vulnerability did not allow an attacker to decrypt or modify the intercepted network communication once the client and server have successfully negotiated a session between themselves.

Refer to the following links to know more about the SSL/TLS vulnerability:

- CVE-2009-3555 (posted on Mitre's Common Vulnerabilities and Exposures List, 2009)
- Understanding the TLS Renegotiation Attack (posted on Eric Rescorla's blog, Educated Guesswork, November 5, 2009).

Phased Approach to Fixing This Issue

The fix for this issue was handled in two phases:



- Phase 1: Until a protocol fix could be developed, an interim fix that disabled SSL/TLS renegotiations by default was made available in the March 30, 2010 Java SE and Java for Business Critical Patch Update.
- Phase 2: The IETF issued RFC 5746, which addresses the renegotiation protocol flaw. The following table lists the JDK and JRE releases that include the fix which implements RFC 5746 and supports secure renegotiation.

Table 8-11 JDK and JRE Releases With Fixes to the TLS Renegotiation Issue

JDK Family	Vulnerable Releases	Phase 1 Fix (Disable Renegotiations)	Phase 2 Fix (RFC 5746)
JDK and JRE 6	Update 18 and earlier	Updates 19 through 21	Update 22
JDK and JRE 5.0	Update 23 and earlier	Updates 24 through 25	Update 26
JDK and JRE 1.4.2	Update 25 and earlier	Updates 26 through 27	Update 28



Applications that do not require renegotiations are not affected by the Phase 2 default configuration. However applications that require renegotiations (for example, web servers that initially allow for anonymous client browsing, but later require SSL/TLS authenticated clients):

- Are not affected if the peer is also compliant with RFC 5746
- Are affected if the peer has not been upgraded to RFC 5746 (see next section for details)

Description of the Phase 2 Fix

The SunJSSE implementation reenables renegotiations by default for connections to peers compliant with RFC 5746. That is, both the client and server **must support RFC 5746** in order to securely renegotiate. SunJSSE provides some interoperability modes for connections with peers that have not been upgraded, but users are **strongly encouraged to update both their client and server implementations as soon as possible**.

With the Phase 2 fix, SunJSSE has three renegotiation interoperability modes. Each mode fully supports the RFC 5746 secure renegotiation, but has these added semantics when communicating with a peer that has not been upgraded:

- **Strict mode**: Requires both client and server be upgraded to RFC 5746 and to send the proper RFC 5746 messages. If not, the initial (or subsequent) handshaking will fail and the connection will be terminated.
- Interoperable mode (default): Use of the proper RFC 5746 messages is optional; however, legacy (original SSL/TLS specifications) renegotiations are disabled if the proper messages are not used. Initial legacy connections are still allowed, but legacy renegotiations are disabled. This is the best mix of security and interoperability, and is the default setting.
- **Insecure mode**: Permits full legacy renegotiation. Most interoperable with legacy peers but vulnerable to the original MITM attack.



The three mode distinctions only affect a connection with a peer that has not been upgraded. Ideally, strict (full RFC 5746) mode should be used for all clients and servers; however, it will take some time for all deployed SSL/TLS implementations to support RFC 5746, because the interoperable mode is the current default.

The following table contains interoperability information about the modes for various cases in which the client and/or server are either updated to support RFC 5746 or not.

Table 8-12 Interoperability Information

Client	Comron	Mode
	Server	
Updated	Updated	Secure renegotiation in all modes.
Legacy [1]	Updated	• Strict If clients do not send the proper RFC 5746 messages, then initial connections will immediately be terminated by the server (SSLHandshakeException or handshake_failure).
		 Interoperable Initial connections from legacy clients are allowed (missing RFC 5746 messages), but renegotiations will not be allowed by the server. [3] [2]
		 Insecure Connections and renegotiations with legacy clients are allowed, but are vulnerable to the original MITM attack.
Updated	Legacy [1]	• Strict If the server does not respond with the proper RFC 5746 messages, then the client will immediately terminate the connection (SSLHandshakeException Or handshake_failure).
		 Interoperable Initial connections from legacy servers are allowed (missing RFC 5746 messages), but renegotiations will not be allowed by the server. [2] [3]
		 Insecure Connections and renegotiations with legacy servers are allowed, but are vulnerable to the original MITM attack.
Legacy [1]	Legacy [1]	Existing SSL/TLS behavior, vulnerable to the MITM attack.

Footnote [1] "Legacy" means the original SSL/TLS specifications (that is, *not* RFC 5746).

Footnote [2] SunJSSE Phase 1 implementations reject renegotiations unless specifically reenabled. If renegotiations are reenabled, then they will be treated as "Legacy" by the peer that is compliant with RFC 5746, because they do not send the proper RFC 5746 messages.

Footnote [3] In SSL/TLS, renegotiations can be initiated by either side. Like the Phase 1 fix, applications communicating with a peer that has not been upgraded in Interoperable mode and that attempt to initiate renegotiation (via



SSLSocket.startHandshake() or SSLEngine.beginHandshake()) will receive an SSLHandshakeException (IOException) and the connection will be shut down (handshake_failure). Applications that receive a renegotiation request from a peer that has not been upgraded will respond according to the type of connection in place:

- TLSv1 A warning alert message of type no_renegotiation(100) will be sent to the peer and the connection will remain open. Older versions of SunJSSE will shut down the connection when a no renegotiation alert is received.
- SSLv3 The application will receive an SSLHandshakeException, and the connection will be closed (handshake_failure). The no_renegotiation alert is not defined in the SSLv3 specification.

Set the mode with the following system properties (see How to Specify a java.lang.System Property):

- sun.security.ssl.allowUnsafeRenegotiation (introduced in Phase 1) controls whether legacy (unsafe) renegotiations are permitted.
- sun.security.ssl.allowLegacyHelloMessages (introduced in Phase 2) allows the peer to perform the handshake process without requiring the proper RFC 5746 messages.



The system properties sun.security.ssl.allowUnsafeRenegotiation and sun.security.ssl.allowLegacyHelloMessages are deprecated and might be removed in a future JDK release.

Table 8-13 Values of the System Properties for Setting the Interoperability Mode

Mode	allowLegacyHelloMessages	allowUnsafeRenegotiation
Strict	false	false
Interoperable (default)	true	false
Insecure	true	true

A

Caution:

Do not reenable the insecure SSL/TLS renegotiation, as this would reestablish the vulnerability.

Workarounds and Alternatives to SSL/TLS Renegotiation

All peers should be updated to RFC 5746-compliant implementation as soon as possible. Even with this RFC 5746 fix, communications with peers that have not been upgraded will be affected if a renegotiation is necessary. Here are a few suggested options:

Restructure the peer to not require renegotiation.



Renegotiations are typically used by web servers that initially allow for anonymous client browsing but later require SSL/TLS authenticated clients, or that may initially allow weak cipher suites but later need stronger ones. The alternative is to require client authentication or strong cipher suites during the *initial* negotiation. There are a couple of options for doing so:

- If an application has a browse mode until a certain point is reached and a renegotiation is required, then you can restructure the server to eliminate the browse mode and require all initial connections be strong.
- Break the server into two entities, with the browse mode occurring on one entity, and using a second entity for the more secure mode. When the renegotiation point is reached, transfer any relevant information between the servers.

Both of these options require a fair amount of work, but will not reopen the original security flaw.

 Set renegotiation interoperability mode to "insecure" using the system properties.

See Description of the Phase 2 Fix.

TLS Implementation Details

RFC 5746 defines two new data structures, which are mentioned here for advanced users

- A pseudo-cipher suite called the Signaling Cipher Suite Value (SCSV),
 "TLS_EMPTY_RENEGOTIATION_INFO_SCSV"
- A TLS extension called the Renegotiation Info (RI).

Either of these can be used to signal that an implementation is RFC 5746-compliant and can perform secure renegotiations. See IETF email discussion from November 2009 to February 2010.

RFC 5746 enables clients to send either an SCSV or RI in the first clientHello. For maximum interoperability, SunJSSE uses the SCSV by default, as a few TLS/SSL servers do not handle unknown extensions correctly. The presence of the SCSV in the enabled cipher suites (SSLSocket.setEnabledCipherSuites()) or SSLEngine.setEnabledCipherSuites()) determines whether the SCSV is sent in the initial ClientHello, or if an RI should be sent instead.

SSLv2 does not support SSL/TLS extensions. If the SSLv2Hello protocol is enabled, then the SCSV is sent in the initial ClientHello.

Description of the Phase 1 Fix

As previously mentioned, the Phase 1 Fix was to disable renegotiations by default until a fix compliant with RFC 5746 could be developed. Renegotiations could be reenabled by setting the sun.security.ssl.allowUnsafeRenegotiation system property. The Phase 2 fix uses the same sun.security.ssl.allowUnsafeRenegotiation system property, but also requires it to use RFC 5746 messages.

All applications should upgrade to the Phase 2 RFC 5746 fix as soon as possible.



Note:

The system properties sun.security.ssl.allowUnsafeRenegotiation and sun.security.ssl.allowLegacyHelloMessages are deprecated and might be removed in a future JDK release.

Allow Unsafe Server Certificate Change in SSL/TLS Renegotiations

Server certificate change in an SSL/TLS renegotiation may be unsafe:

- 1. If endpoint identification is not enabled in an SSL/TLS handshaking; and
- If the previous handshake is a session-resumption abbreviated initial handshake; and
- **3.** If the identities represented by both certificates can be regarded as different.

Two certificates can be considered to represent the same identity:

- 1. If the subject alternative names of IP address are present in both certificates, they should be identical; otherwise,
- 2. If the subject alternative names of DNS name are present in both certificates, they should be identical; otherwise,
- 3. If the subject fields are present in both certificates, the certificate subjects and issuers should be identical.

Starting with JDK 8u25, unsafe server certificate change in SSL/TLS renegotiations is not allowed by default. The new system property

 ${\tt jdk.tls.allowUnsafeServerCertChange}, \ can \ be \ used \ to \ define \ whether \ unsafe \ server \ certificate \ change \ in \ an \ SSL/TLS \ renegotiation \ should \ be \ restricted \ or \ not.$

The default value of this system property is "false".



Caution:

DO NOT set the system property to "true" unless it is really necessary, as this would re-establish the unsafe server certificate change vulnerability.

Hardware Acceleration and Smartcard Support

The Java Cryptography Architecture (JCA) is a set of packages that provides a framework and implementations for encryption, key generation and key agreement, and message authentication code (MAC) algorithms. (See Java Cryptography Architecture (JCA) Reference Guide.) The SunJSSE provider uses JCA exclusively for all of its cryptographic operations and can automatically take advantage of JCE features and enhancements, including JCA's support for RSA PKCS#11. This support enables the SunJSSE provider to use hardware cryptographic accelerators for significant performance improvements and to use smartcards as keystores for greater flexibility in key and trust management.



Use of hardware cryptographic accelerators is automatic if JCA has been configured to use the Oracle PKCS#11 provider, which in turn has been configured to use the underlying accelerator hardware. The provider must be configured before any other JCA providers in the provider list. For details on how to configure the Oracle PKCS#11 provider, see PKCS#11 Reference Guide.

Configuring JSSE to Use Smartcards as Keystores and Truststores

Support for PKCS#11 in JCA also enables access to smartcards as a keystore. For details on how to configure the type and location of the keystores to be used by JSSE, see Customizing JSSE. To use a smartcard as a keystore or truststore, set the <code>javax.net.ssl.keyStoreType</code> and <code>javax.net.ssl.trustStoreType</code> system properties, respectively, to <code>pkcsl1</code>, and set the <code>javax.net.ssl.keyStore</code> and <code>javax.net.ssl.trustStore</code> system properties, respectively, to <code>NONE</code>. To specify the use of a specific provider, use the <code>javax.net.ssl.keyStoreProvider</code> and <code>javax.net.ssl.trustStoreProvider</code> system properties (for example, set them to <code>SunPKCSll-joe</code>). By using these properties, you can configure an application that previously depended on these properties to access a file-based keystore to use a smartcard keystore with no changes to the application.

Some applications request the use of keystores programmatically. These applications can continue to use the existing APIs to instantiate a Keystore and pass it to its key manager and trust manager. If the Keystore instance refers to a PKCS#11 keystore backed by a Smartcard, then the JSSE application will have access to the keys on the smartcard.

Multiple and Dynamic Keystores

Smartcards (and other removable tokens) have additional requirements for an x509KeyManager. Different smartcards can be present in a smartcard reader during the lifetime of a Java application, and they can be protected using different passwords.

The KeyStore.Builder class abstracts the construction and initialization of a KeyStore object. It supports the use of CallbackHandler for password prompting, and its subclasses can be used to support additional features as desired by an application. For example, it is possible to implement a Builder that allows individual KeyStore entries to be protected with different passwords. The KeyStoreBuilderParameters class then can be used to initialize a KeyManagerFactory using one or more of these Builder objects.

A x509KeyManager implementation in the SunJSSE provider called NewSunX509 supports these parameters. If multiple certificates are available, it attempts to pick a certificate with the appropriate key usage and prefers valid to expired certificates.

Example 8-17 illustrates how to tell JSSE to use both a PKCS#11 keystore (which might in turn use a smartcard) and a PKCS#12 file-based keystore.

Example 8-17 Sample Code to Use PKCS#11 and PKCS#12 File-based Keystore

```
import javax.net.ssl.*;
import java.security.KeyStore.*;
// ...
// Specify keystore builder parameters for PKCS#11 keystores
Builder scBuilder = Builder.newInstance("PKCS11", null,
```



Kerberos Cipher Suites

The SunJSSE provider has support for Kerberos cipher suites, as described in RFC 2712. The following cipher suites are supported but not enabled by default:



According to DTLS Version 1.0 and DTLS Version 1.2, RC4 cipher suites must not be used with DTLS.

- TLS_KRB5_WITH_RC4_128_SHA
- TLS_KRB5_WITH_RC4_128_MD5
- TLS KRB5 WITH 3DES EDE CBC SHA
- TLS_KRB5_WITH_3DES_EDE_CBC_MD5
- TLS_KRB5_WITH_DES_CBC_SHA
- TLS_KRB5_WITH_DES_CBC_MD5
- TLS_KRB5_EXPORT_WITH_RC4_40_SHA
- TLS_KRB5_EXPORT_WITH_RC4_40_MD5
- TLS_KRB5_EXPORT_WITH_DES_CBC_40_SHA
- TLS_KRB5_EXPORT_WITH_DES_CBC_40_MD5

To enable the use of these cipher suites, you must do so explicitly. See the API documentation for <code>SSLEngine.setEnabledCipherSuites(String[])</code> and <code>SSLSocket.setEnabledCipherSuites(String[])</code> methods. As with all other SSL/TLS/DTLS cipher suites, if a cipher suite is not supported by the peer, then it will not be selected during cipher negotiation. Furthermore, if the application and/or server cannot acquire the necessary Kerberos credentials, then the Kerberos cipher suites also will not be selected.



The following is an example of a TLS client that will only use the TLS_KRB5_WITH_DES_CBC_SHA cipher suite:

```
// Create socket
SSLSocketFactory sslsf = (SSLSocketFactory) SSLSocketFactory.getDefault();
SSLSocket sslSocket = (SSLSocket) sslsf.createSocket(tlsServer, serverPort);

// Enable only one cipher suite
String enabledSuites[] = { "TLS_KRB5_WITH_DES_CBC_SHA" };
sslSocket.setEnabledCipherSuites(enabledSuites);
```

Kerberos Requirements

You must have the Kerberos infrastructure set up in your deployment environment before you can use the Kerberos cipher suites with JSSE. In particular, both the TLS client and server must have accounts set up with the Kerberos Key Distribution Center (KDC). At runtime, if one or more of the Kerberos cipher suites have been enabled, then the TLS client and server will acquire their Kerberos credentials associated with their respective account from the KDC. For example, a TLS server running on the machine machi.imc.org in the Kerberos realm IMC.ORG must have an account with the name host/machl.imc.org@IMC.ORG and be configured to use the KDC for IMC.ORG. See Kerberos Requirements.

An application can acquire its Kerberos credentials by using the Java Authentication and Authorization Service (JAAS) Reference Guide and a Kerberos login module. The JDK comes with a Krb5LoginModule. You can use the Kerberos cipher suites with JSSE with or without JAAS programming, similar to how you can use the JAAS and Java GSS-API Tutorial with or without JAAS programming.

To use the Kerberos cipher suites with JSSE without JAAS programming, you must use the index names <code>com.sun.net.ssl.server</code> or other for the TLS server JAAS configuration entry, and <code>com.sun.net.ssl.client</code> or other for the TLS client, and set the <code>javax.security.auth.useSubjectCredsOnly</code> system property to false. For example, a TLS server that is not using JAAS programming might have the following JAAS configuration file:

```
com.sun.net.ssl.server {
  com.sun.security.auth.module.Krb5LoginModule required
  principal="host/mach1.imc.org@IMC.ORG"
  useKeyTab=true
  keyTab=mach1.keytab
  storeKey=true;
};
```

An example of how to use Java GSS and Kerberos without JAAS programming is described in the tutorial Use of Java GSS-API for Secure Message Exchanges Without JAAS Programming in the JDK 8 documentation. You can adapt it to use JSSE by replacing Java GSS calls with JSSE calls.

To use the Kerberos cipher suites with JAAS programming, you can use any index name because your application is responsible for creating the JAAS LoginContext using the index name, and then wrapping the JSSE calls inside of a Subject.doAs() or Subject.doAsPrivileged() call. An example of how to use JAAS with Java GSS and Kerberos is described in the tutorial Use of JAAS Login Utility and Java GSS-API for Secure Message Exchange in the JDK 8 documentation. You can adapt it to use JSSE by replacing Java GSS calls with JSSE calls.



If you have trouble using or configuring the JSSE application to use Kerberos, see Troubleshooting in the Java GSS Tutorial in the JDK 8 documentation.

Peer Identity Information

To determine the identity of the peer of an SSL/TLS/DTLS connection, use the getPeerPrincipal() method in the following classes:

- javax.net.ssl.SSLSession
- javax.net.ssl.HttpsURLConnection
- javax.net.HandshakeCompletedEvent

Similarly, to get the identity that was sent to the peer (to identify the local entity), use the <code>getLocalPrincipal()</code> method in these classes. For X509-based cipher suites, these methods will return an instance of <code>javax.security.auth.x500.X500Principal</code>; for Kerberos cipher suites, these methods will return an instance of <code>javax.security.auth.kerberos.KerberosPrincipal</code>.

JSSE applications use <code>getPeerCertificates()</code> and similar methods in <code>javax.net.ssl.SSLSession</code>, <code>javax.net.ssl.HttpsURLConnection</code>, and <code>javax.net.HandshakeCompletedEvent</code> classes to obtain information about the peer. When the peer does not have any certificates, <code>SSLPeerUnverifiedException</code> is thrown.

If the application must determine only the identity of the peer or identity sent to the peer, then it should use the <code>getPeerPrincipal()</code> and <code>getLocalPrincipal()</code> methods, respectively. It should use <code>getPeerCertificates()</code> and <code>getLocalCertificates()</code> methods only if it must examine the contents of those certificates. Furthermore, the application must be prepared to handle the case where an authenticated peer might not have any certificate.

Security Manager

When the security manager has been enabled, in addition to the <code>SocketPermission</code> needed to communicate with the peer, a TLS client application that uses the Kerberos cipher suites also needs the following permission:

```
javax.security.auth.kerberos.ServicePermission(serverPrincipal, "initiate");
```

Where,

serverPrincipal

Indicates the Kerberos principal name of the TLS server that the TLS client will be communicating with (such as host/mach1.imc.org@IMC.ORG).

A TLS server application needs the following permission:

```
javax.security.auth.kerberos.ServicePermission(serverPrincipal, "accept");
```

Where,

serverPrincipal

Indicates the Kerberos principal name of the TLS server (such as host/machl.imc.org@IMC.ORG).

If the server or client must contact the KDC (for example, if its credentials are not cached locally), then it also needs the following permission:



javax.security.auth.kerberos.ServicePermission(tgtPrincipal, "initiate");

Where,

tgtPrincipal

Indicates the principal name of the KDC (such as krbtgt/IMC.ORG@IMC.ORG).

Additional Keystore Formats (PKCS12)

The PKCS#12 (Personal Information Exchange Syntax Standard) specifies a portable format for storage and/or transport of a user's private keys, certificates, miscellaneous secrets, and other items. The SunJSSE provider supplies a complete implementation of the PKCS12 <code>java.security.KeyStore</code> format for reading and writing PKCS12 files. This format is also supported by other toolkits and applications for importing and exporting keys and certificates, such as Mozilla Firefox, Microsoft Internet Explorer, and OpenSSL. For example, these implementations can export client certificates and keys into a file using the .p12 file name extension.

With the SunJSSE provider, you can access PKCS12 keys through the KeyStore API with a keystore type of PKCS12. In addition, you can list the installed keys and associated certificates by using the keytool command with the -storetype option set to pkcs12. See keytool in *Java Platform, Standard Edition Tools Reference*.

Server Name Indication (SNI) Extension

The SNI extension is a feature that extends the SSL/TLS/DTLS protocols to indicate what server name the client is attempting to connect to during handshaking. Servers can use server name indication information to decide if specific <code>sslsocket</code> or <code>sslengine</code> instances should accept a connection. For example, when multiple virtual or name-based servers are hosted on a single underlying network address, the server application can use SNI information to determine whether this server is the exact server that the client wants to access. Instances of this class can be used by a server to verify the acceptable server names of a particular type, such as host names. See section 3 of TLS Extensions (RFC 6066).

Developers of client applications can explicitly set the server name indication using the SSLParameters.setServerNames(List<SNIServerName> serverNames) method. See Example 8-18.

Developers of server applications can use the SNIMatcher class to decide how to recognize server name indication. Example 8-19 and Example 8-20 illustrate this functionality:

Example 8-18 Sample Code to Set the Server Name Indication

The following code sample illustrates how to set the server name indication using the method SSLParameters.setServerNames(List<SNIServerName> serverNames):

```
SSLSocketFactory factory = ...
SSLSocket sslSocket = factory.createSocket("172.16.10.6", 443);
// SSLEngine sslEngine = sslContext.createSSLEngine("172.16.10.6", 443);
SNIHostName serverName = new SNIHostName("www.example.com");
List<SNIServerName> serverName> = new ArrayList<>(1);
```



```
serverNames.add(serverName);

SSLParameters params = sslSocket.getSSLParameters();
params.setServerNames(serverNames);
sslSocket.setSSLParameters(params);
// sslEngine.setSSLParameters(params);
```

Example 8-19 Sample Code Using SSLSocket Class to Recognize SNI

The following code sample illustrates how the server applications can use the SNIMatcher class to decide how to recognize server name indication:

```
SSLSocket sslSocket = sslServerSocket.accept();

SNIMatcher matcher = SNIHostName.createSNIMatcher("www\\.example\\.(com|org)");
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);

SSLParameters params = sslSocket.getSSLParameters();
params.setSNIMatchers(matchers);
sslSocket.setSSLParameters(params);
```

Example 8-20 Sample Code Using SSLServerSocket Class to Recognize SNI

The following code sample illustrates how the server applications can use the SNIMatcher class to decide how to recognize server name indication:

```
SSLServerSocket sslServerSocket = ...;
SNIMatcher matcher = SNIHostName.createSNIMatcher("www\\.example\\.(com|org)");
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);

SSLParameters params = sslServerSocket.getSSLParameters();
params.setSNIMatchers(matchers);
sslServerSocket.setSSLParameters(params);

SSLSocket sslSocket = sslServerSocket.accept();
```

The following list provides examples for the behavior of the SNIMatcher when receiving various server name indication requests in the ClientHello message:

- Matcher configured to www\\.example\\.com:
 - If the requested host name is www.example.com, then it will be accepted and a confirmation will be sent in the ServerHello message.
 - If the requested host name is www.example.org, then it will be rejected with an unrecognized_name fatal error.
 - If there is no requested host name or it is empty, then the request will be accepted but no confirmation will be sent in the ServerHello message.
- Matcher configured to www\\.invalid\\.com:
 - If the requested host name is www.example.com, then it will be rejected with an unrecognized_name fatal error.
 - If the requested host name is www.example.org, then it will be accepted and a confirmation will be sent in the ServerHello message.



- If there is no requested host name or it is empty, then the request will be accepted but no confirmation will be sent in the ServerHello message.
- Matcher is not configured:

Any requested host name will be accepted but no confirmation will be sent in the ServerHello message.

For descriptions of new classes that implement the SNI extension, see:

- The StandardConstants Class
- The SNIServerName Class
- The SNIMatcher Class
- The SNIHostName Class

For examples, see Using the Server Name Indication (SNI) Extension.

TLS Application Layer Protocol Negotiation

Negotiate an application protocol for a TLS connection with Application Layer Protocol Negotiation (ALPN).

What is ALPN?

Some applications might want or need to negotiate a shared application level value before a TLS handshake has completed. For example, HTTP/2 uses the Application Layer Protocol Negotiation mechanism to help establish which HTTP version ("h2", "spdy/3", "http/1.1") can or will be used on a particular TCP or UDP port. ALPN (RFC 7301) does this without adding network round-trips between the client and the server. In the case of HTTP/2 the protocol must be established before the connection is negotiated, as client and server need to know what version of HTTP to use before they start communicating. Without ALPN it would not be possible to have application protocols HTTP/1 and HTTP/2 on the same port.

The client uses the ALPN extension at the beginning of the TLS handshake to send a list of supported application protocols to the server as part of the <code>clientHello</code>. The server reads the list of supported application protocols in the <code>clientHello</code>, and determines which of the supported protocols it prefers. It then sends a <code>serverHello</code> message back to the client with the negotiation result. The message may contain either the name of the protocol that has been chosen or that no protocol has been chosen.

The application protocol negotiation can thus be accomplished within the TLS handshake, without adding network round-trips, and allows the server to associate a different certificate with each application protocol, if desired.

Unlike many other TLS extensions, this extension does not establish properties of the session, only of the connection. That's why you'll find the negotiated values in the SSLSocket/SSLEngine, not the SSLSession. When session resumption or session tickets are used (see TLS Session Resumption without Server-Side State), the previously negotiated values are irrelevant, and only the values in the new handshake messages are considered.



Setting up ALPN on the Client

Set the Application Layer Protocol Negotiation (ALPN) values supported by the client. During the handshake with the server, the server will read the client's list of application protocols and will determine which is most suitable.

For the client, use the <code>SSLParameters.setApplicationProtocols(String[])</code> method, followed by the <code>setSSLParameters</code> method of either <code>SSLSocket</code> or <code>SSLEngine</code> to set up the application protocols to send to the server.

Example 8-21 Sample Code for Setting and Getting ALPN Values in a Java Client

For example, here are the steps to set ALPN values of "three" and "two", on the client.

To run the code the property <code>javax.net.ssl.trustStore</code> must be set to a valid root certificate. (This can be done on the command line).

```
import java.io.*;
import java.util.*;
import javax.net.ssl.*;
public class SSLClient {
    public static void main(String[] args) throws Exception {
        // Code for creating a client side SSLSocket
        SSLSocketFactory sslsf = (SSLSocketFactory) SSLSocketFactory.getDefault();
        SSLSocket sslSocket = (SSLSocket) sslsf.createSocket("localhost", 9999);
        // Get an SSLParameters object from the SSLSocket
        SSLParameters sslp = sslSocket.getSSLParameters();
        // Populate SSLParameters with the ALPN values
        // On the client side the order doesn't matter as
        // when connecting to a JDK server, the server's list takes priority
        String[] clientAPs = {"three", "two"};
        sslp.setApplicationProtocols(clientAPs);
        // Populate the SSLSocket object with the SSLParameters object
        // containing the ALPN values
        sslSocket.setSSLParameters(sslp);
        sslSocket.startHandshake();
        // After the handshake, get the application protocol that has been negotiated
        String ap = sslSocket.getApplicationProtocol();
        System.out.println("Application Protocol client side: \"" + ap + "\"");
        // Do simple write/read
        InputStream sslIS = sslSocket.getInputStream();
        OutputStream sslOS = sslSocket.getOutputStream();
        sslOS.write(280);
       sslOS.flush();
       sslIS.read();
        sslSocket.close();
```

When this code is run and sends a ClientHello to a Java server that has set the ALPN values one, two, and three, the output will be:

}

```
Application Protocol client side: two
```

See The SSL Handshake for further details on handshaking. It is also possible to check the results of the negotiation during handshaking. See Determining Negotiated ALPN Value during Handshaking.

Setting up Default ALPN on the Server

Use the default ALPN mechanism to determine a suitable application protocol by setting ALPN values on the server.

To use the default mechanism for ALPN on the server, populate an SSLParameters object with the ALPN values you wish to set, and then use this SSLParameters object to populate either the SSLSocket object or the SSLEngine object with these parameters as you have done when you set up ALPN on the client (see the section Setting up ALPN on the Client). The first value of the ALPN values set on the server that matches any of the ALPN values contained in the ClientHello will be chosen and returned to the client as part of the ServerHello.

Example 8-22 Sample Code for Default ALPN Value Negotiation on the Server

Here is the code for a Java server that uses the default approach for protocol negotiation. To run the code the property <code>javax.net.ssl.keyStore</code> must be set to a valid keystore. (This can be done on the command line, see Creating a Keystore to Use with <code>JSSE</code>).

```
import java.util.*;
import javax.net.ssl.*;
public class SSLServer {
    public static void main(String[] args) throws Exception {
        // Code for creating a server side SSLSocket
        SSLServerSocketFactory sslssf =
            (SSLServerSocketFactory) SSLServerSocketFactory.getDefault();
        SSLServerSocket sslServerSocket =
            (SSLServerSocket) sslssf.createServerSocket(9999);
        SSLSocket sslSocket = (SSLSocket) sslServerSocket.accept();
        // Get an SSLParameters object from the SSLSocket
        SSLParameters sslp = sslSocket.getSSLParameters();
        // Populate SSLParameters with the ALPN values
        // As this is server side, put them in order of preference
        String[] serverAPs ={ "one", "two", "three" };
        sslp.setApplicationProtocols(serverAPs);
        // If necessary at any time, get the ALPN values set on the
        // SSLParameters object with:
        // String serverAPs = sslp.setApplicationProtocols();
        // Populate the SSLSocket object with the ALPN values
        sslSocket.setSSLParameters(sslp);
        sslSocket.startHandshake();
        // After the handshake, get the application protocol that
        // has been negotiated
        String ap = sslSocket.getApplicationProtocol();
        System.out.println("Application Protocol server side: \"" + ap + "\"");
```



```
// Continue with the work of the server
InputStream sslIS = sslSocket.getInputStream();
OutputStream sslOS = sslSocket.getOutputStream();
sslIS.read();
sslOS.write(85);
sslOS.flush();
sslSocket.close();
}
```

When this code is run and a Java client sends a ClientHello with ALPN values three and two, the output is:

```
Application Protocol server side: two
```

See The SSL Handshake for further details on handshaking. It is also possible to check the results of the negotiation during handshaking. See Determining Negotiated ALPN Value during Handshaking.

Setting up Custom ALPN on the Server

Use the custom ALPN mechanism to determine a suitable application protocol by setting up a callback method.

If you do not want to use the server's default negotiation protocol, you can use the <code>setHandshakeApplicationProtocolSelector</code> method of <code>SSLEngine</code> or <code>SSLSocket</code> to register a <code>BiFunction</code> (lambda) callback that can examine the handshake state so far, and then make your selection based on the client's list of application protocols and any other relevant information. For example, you may consider using the cipher suite suggested, or the Server Name Indication (SNI) or any other data you can obtain in making the choice. If custom negotiation is used, the values set by the <code>setApplicationProtocols</code> method (default negotiation) will be ignored.

Example 8-23 Sample Code for Custom ALPN Value Negotiation on the Server

Here is the code for a Java server that uses the custom mechanism for protocol negotiation. To run the code the property <code>javax.net.ssl.keyStore</code> must be set to a valid certificate. (This can be done on the command line, see Creating a Keystore to Use with <code>JSSE</code>).

```
import java.util.*;
import javax.net.ssl.*;
public class SSLServer {
    public static void main(String[] args) throws Exception {
        // Code for creating a server side SSLSocket
        SSLServerSocketFactory sslssf =
            (SSLServerSocketFactory) SSLServerSocketFactory.getDefault();
        SSLServerSocket sslServerSocket =
            (SSLServerSocket) sslssf.createServerSocket(9999);
        SSLSocket sslSocket = (SSLSocket) sslServerSocket.accept();
        // Code to set up a callback function
        // Pass in the current SSLSocket to be inspected and client AP values
        sslSocket.setHandshakeApplicationProtocolSelector(
            (serverSocket, clientProtocols) -> {
                SSLSession handshakeSession = serverSocket.getHandshakeSession();
                // callback function called with current SSLSocket and client AP
```



```
values
                // plus any other useful information to help determine appropriate
                // application protocol. Here the protocol and ciphersuite are also
                // passed to the callback function.
                return chooseApplicationProtocol(
                    serverSocket.
                    clientProtocols.
                    handshakeSession.getProtocol(),
                    handshakeSession.getCipherSuite());
        });
        sslSocket.startHandshake();
        // After the handshake, get the application protocol that has been
        // returned from the callback method.
        String ap = sslSocket.getApplicationProtocol();
        System.out.println("Application Protocol server side: \"" + ap + "\"");
        // Continue with the work of the server
        InputStream sslIS = sslSocket.getInputStream();
        OutputStream ssloS = sslSocket.getOutputStream();
        sslIS.read();
        sslOS.write(85);
       sslOS.flush();
        sslSocket.close();
    // The callback method. Note how the parameters match the call within
    // the setHandshakeApplicationProtocolSelector method above.
    public static String chooseApplicationProtocol(SSLSocket serverSocket,
            List<String> clientProtocols, String protocol, String cipherSuite ) {
        // For example, check the cipher suite and return an application protocol
        // value based on that.
        if (cipherSuite.equals("<--a_particular_ciphersuite-->")) {
            return "three";
        } else {
            return "";
```

If the cipher suite matches the one you specify in the condition statement when this code is run , then the value ${\tt three}$ will be returned. Otherwise an empty string will be returned.

Note that the BiFunction object's return value is a string, which will be the application protocol name, or null to indicate that none of the advertised names are acceptable. If the return value is an empty string then application protocol indications will not be used. If the return value is null (no value chosen) or is a value that was not advertised by the peer, the underlying protocol will determine what action to take. (For example, the server code will send a "no_application_protocol" alert and terminate the connection.)

After handshaking completes on both client and server, you can check the result of the negotiation by calling the <code>getApplicationProtocol</code> method on either the <code>SSLSocket</code> object or the <code>SSLEngine</code> object. See The SSL Handshake for further details on handshaking.

Determining Negotiated ALPN Value during Handshaking

To determine the ALPN value that has been negotiated during the handshaking, create a custom KeyManager Or TrustManager class, and include in this custom class a call to the getHandshakeApplicationProtocol method.

There are some use cases where the selected ALPN and SNI values will affect the choices made by a <code>KeyManager</code> or <code>TrustManager</code>. For example, an application might want to select different certificate/private key sets depending on the attributes of the server and the chosen ALPN/SNI/ciphersuite values.

The sample code given illustrates how to call the <code>getHandshakeApplicationProtocol</code> method from within a custom <code>x509ExtendedKeyManager</code> that you create and register as the <code>KeyManager</code> object.

Example 8-24 Sample Code for a Custom KeyManager

This example shows the entire code for a custom KeyManager that extends X509ExtendedKeyManager. Most methods simply return the value returned from the KeyManager class that is being wrapped by this MyX509ExtendedKeyManager class. However the chooseServerAlias method calls the getHandshakeApplicationProtocol on the SSLSocket object and therefore can determine the current negotiated ALPN value.

```
import java.net.Socket;
import java.security.*;
import javax.net.ssl.*;
public class MyX509ExtendedKeyManager extends X509ExtendedKeyManager {
    // X509ExtendedKeyManager is an abstract class so your new class
    // needs to implement all the abstract methods in this class.
    // The easiest way to do this is to wrap an existing KeyManager
    // and call its methods for each of the methods you need to implement.
    X509ExtendedKeyManager akm;
    public MyX509ExtendedKeyManager(X509ExtendedKeyManager akm) {
        this.akm = akm;
    @Override
    public String[] getClientAliases(String keyType, Principal[] issuers) {
       return akm.getClientAliases(keyType, issuers);
    public String chooseClientAlias(String[] keyType, Principal[] issuers,
       Socket socket) {
       return akm.chooseClientAlias(keyType, issuers, socket);
    @Override
    public String chooseServerAlias(String keyType, Principal[] issuers,
       Socket socket) {
        // This method has access to a Socket, so it is possible to call the
        // getHandshakeApplicationProtocol method here. Note the cast from
```



```
// a Socket to an SSLSocket
   String ap = ((SSLSocket) socket).getHandshakeApplicationProtocol();
   System.out.println("In chooseServerAlias, ap is: " + ap);
   return akm.chooseServerAlias(keyType, issuers, socket);
}

@Override
public String[] getServerAliases(String keyType, Principal[] issuers) {
   return akm.getServerAliases(keyType, issuers);
}

@Override
public X509Certificate[] getCertificateChain(String alias) {
   return akm.getCertificateChain(alias);
}

@Override
public PrivateKey getPrivateKey(String alias) {
   return akm.getPrivateKey(alias);
}
```

When this code is registered as the KeyManager for a Java server and a Java client sends a ClientHello with ALPN values, the output will be:

In chooseServerAlias, ap is: <negotiated value>

Example 8-25 Sample Code for Using a Custom KeyManager in a Java Server

This example shows a simple Java server that uses the default ALPN negotiation strategy and the custom <code>KeyManager</code>, <code>MyX509ExtendedKeyManager</code>, shown in the prior code sample.

```
import java.io.*;
import java.util.*;
import javax.net.ssl.*;
import java.security.KeyStore;
public class SSLServerHandshake {
    public static void main(String[] args) throws Exception {
        SSLContext ctx = SSLContext.getInstance("TLS");
        // You need to explicitly create a create a custom KeyManager
        // Keystores
        KeyStore keyKS = KeyStore.getInstance("PKCS12");
        keyKS.load(new FileInputStream("serverCert.p12"),
            "password".toCharArray());
        // Generate KeyManager
        KeyManagerFactory kmf = KeyManagerFactory.getInstance("PKIX");
        kmf.init(keyKS, "password".toCharArray());
        KeyManager[] kms = kmf.getKeyManagers();
        // Code to substitute MyX509ExtendedKeyManager
        if (!(kms[0] instanceof X509ExtendedKeyManager)) {
            throw new Exception("kms[0] not X509ExtendedKeyManager");
        // Create a new KeyManager array and set the first index
```



```
// of the array to an instance of MyX509ExtendedKeyManager.
        // Notice how creating this object is done by passing in the
        // existing default X509ExtendedKeyManager
        kms = new KeyManager[] {
            new MyX509ExtendedKeyManager((X509ExtendedKeyManager) kms[0])};
        // Initialize SSLContext using the new KeyManager
        ctx.init(kms, null, null);
        // Instead of using SSLServerSocketFactory.getDefault(),
        // get a SSLServerSocketFactory based on the SSLContext
        SSLServerSocketFactory sslssf = ctx.getServerSocketFactory();
        SSLServerSocket sslServerSocket =
            (SSLServerSocket) sslssf.createServerSocket(9999);
        SSLSocket sslSocket = (SSLSocket) sslServerSocket.accept();
        SSLParameters sslp = sslSocket.getSSLParameters();
        String[] serverAPs ={"one","two","three"};
        sslp.setApplicationProtocols(serverAPs);
        sslSocket.setSSLParameters(sslp);
        sslSocket.startHandshake();
        String ap = sslSocket.getApplicationProtocol();
        System.out.println("Application Protocol server side: \"" + ap + "\"");
        InputStream sslIS = sslSocket.getInputStream();
       OutputStream ssloS = sslSocket.getOutputStream();
        sslIS.read();
        sslOS.write(85);
        sslOS.flush();
        sslSocket.close();
        sslServerSocket.close();
}
```

With the custom X509ExtendedKeyManager in place, when chooseServerAlias is called during handshaking the KeyManager has the opportunity to examine the negotiated application protocol value. In the case of the example shown, this value is output to the console.

For example, when this code is run and a Java client sends a ClientHello with ALPN values three and two, the output will be:

Application Protocol server side: two

ALPN Related Classes and Methods

These classes and methods are used when working with Application Layer Protocol Negotiation (ALPN).

Classes and Methods to Use

SSLEngine and SSLSocket contain the same ALPN related methods and they have the same functionality.



Class	Method	Purpose					
SSLParameters	<pre>public String[] getApplicationProtocols();</pre>	Client-side and server-side: use the method to return a String array containing each protocol set.					
SSLParameters	<pre>public void setApplicationProtocols([] prot ocols);</pre>	Client-side : use the method to set the protocols that can be chosen by the server.					
		Server-side : use the method to set the protocols that the server can use. The String array should contain the protocols in order of preference.					
SSLEngine	public String	Client-side and server-side: use					
SSLSocket	<pre>getApplicationProtocol();</pre>	the method <i>after</i> TLS protocol negotiation has completed to return a String containing the protocol that has been chosen for the connection.					
SSLEngine	public String	Client-side and server-side: use					
SSLSocket	<pre>getHandshakeApplicationProtoco 1();</pre>	the method <i>during</i> handshaking to return a String containing the protocol that has been chosen for the connection. If this method is called before or after handshaking, it will return null. See Determining Negotiated ALPN Value during Handshaking for instructions on how to call this method.					
SSLEngine	public void	Server-side: use the method to					
SSLSocket	<pre>setHandshakeApplicationProtocol Selector(BiFunction,String> selector)</pre>	register a callback function. The application protocol value can then be set in the callback based on any information available, for example the protocol or cipher suite. See Setting up Custom ALPN on the Server for instructions on how to use this method.					

Troubleshooting JSSE

This section contains information for troubleshooting JSSE. It provides solutions to common configuration problem.

First, it provides some common Configuration Problems and ways to solve them, and then it describes helpful Debugging Utilities.

Configuration Problems

Solutions to some common configuration problems.

CertificateException While Handshaking

Problem: When negotiating an SSL/TLS/DTLS connection, the client or server throws a CertificateException.

Cause 1: This is generally caused by the remote side sending a certificate that is unknown to the local side.

Solution 1: The best way to debug this type of problem is to turn on debugging (see Debugging Utilities) and watch as certificates are loaded and when certificates are received via the network connection. Most likely, the received certificate is unknown to the trust mechanism because the wrong trust file was loaded.

Refer to the following sections:

- JSSE Classes and Interfaces
- The TrustManager Interface
- The KeyManager Interface

Cause 2: The system clock is not set correctly. In this case, the perceived time may be outside the validity period on one of the certificates, and unless the certificate can be replaced with a valid one from a truststore, the system must assume that the certificate is invalid, and therefore throw the exception.

Solution 2: Correct the system clock time.

Runtime Exception: SSL Service Not Available

Problem: When running a program that uses JSSE, an exception occurs indicating that an SSL service is not available. For example, an exception similar to one of the following is thrown:

```
Exception in thread "main" java.net.SocketException:
    no SSL Server Sockets

Exception in thread "main":
    SSL implementation not available
```

Cause: There was a problem with SSLContext initialization, for example, due to an incorrect password on a keystore or a corrupted keystore (a JDK vendor once shipped a keystore in an unknown format, and that caused this type of error).

Solution: Check initialization parameters. Ensure that any keystores specified are valid and that the passwords specified are correct. One way that you can check this is by trying to use keytool to examine the keystores and the relevant contents. See keytool in *Java Platform, Standard Edition Tools Reference*.

Runtime Exception: "No available certificate corresponding to the SSL cipher suites which are enabled"

Problem: When trying to run a simple SSL server program, the following exception is thrown:

```
Exception in thread "main" javax.net.ssl.SSLException:

No available certificate corresponding to the SSL cipher suites which are enabled...
```

Cause: Various cipher suites require certain types of key material. For example, if an RSA cipher suite is enabled, then an RSA keyEntry must be available in the keystore.



If no such key is available, then this cipher suite cannot be used. This exception is thrown if there are no available key entries for all of the cipher suites enabled.

Solution: Create key entries for the various cipher suite types, or use an anonymous suite. Anonymous cipher suites are inherently dangerous because they are vulnerable to MITM (man-in-the-middle) attacks. See RFC 2246.

Refer to the following sections to learn how to pass the correct keystore and certificates:

- JSSE Classes and Interfaces
- Customizing the Default Keystores and Truststores, Store Types, and Store Passwords
- Additional Keystore Formats (PKCS12)

Runtime Exception: No Cipher Suites in Common

Problem 1: When handshaking, the client and/or server throw this exception.

Cause 1: Both sides of an SSL connection must agree on a common cipher suite. If the intersection of the client's cipher suite set with the server's cipher suite set is empty, then you will see this exception.

Solution 1: Configure the enabled cipher suites to include common cipher suites, and be sure to provide an appropriate keyEntry for asymmetric cipher suites. Also see Runtime Exception: "No available certificate corresponding to the SSL cipher suites which are enabled" in this section.)

Problem 2: When using Mozilla Firefox or Microsoft Internet Explorer to access files on a server that only has DSA-based certificates, a runtime exception occurs indicating that there are no cipher suites in common.

Cause 2: By default, keyEntries created with keytool use DSA public keys. If only DSA keyEntries exist in the keystore, then only DSA-based cipher suites can be used. By default, Navigator and Internet Explorer send only RSA-based cipher suites. Because the intersection of client and server cipher suite sets is empty, this exception is thrown.

Solution 2: To interact with Navigator or Internet Explorer, you should create certificates that use RSA-based keys. To do this, specify the -keyalg RSA option when using keytool. For example:

keytool -genkeypair -alias duke -keystore testkeys -keyalg rsa

Socket Disconnected After Sending ClientHello Message

Problem: A socket attempts to connect, sends a ClientHello message, and is immediately disconnected.

Cause: Some SSL/TLS servers will disconnect if a ClientHello message is received in a format they do not understand or with a protocol version number that they do not support.

Solution: Try adjusting the enabled protocols on the client side. This involves modifying or invoking some of the following system properties and methods:

System property https.protocols for the HttpsURLConnection class



- System property jdk.tls.client.protocols
- SSLContext.getInstance method
- SSLEngine.setEnabledProtocols method
- SSLSocket.setEnabledProtocols method
- SSLParameters.setProtocols and SSLEngine.setSSLParameters methods
- SSLParameters.setProtocols and SSLSocket.setSSLParameters methods

For backwards compatibility, some SSL/TLS implementations (such as SunJSSE) can send SSL/TLS ClientHello messages encapsulated in the SSLv2 ClientHello format. The SunJSSE provider supports this feature. If you want to use this feature, add the "SSLv2Hello" protocol to the enabled protocol list, if necessary. (See Protocols in the JDK Providers, which lists the protocols that are enabled by default for the SunJSSE provider.)

The SSL/TLS RFC standards require that implementations negotiate to the latest version both sides speak, but some non-conforming implementation simply hang up if presented with a version they don't understand. For example, some older server implementations that speak only SSLv3 will shutdown if TLSv1.2 is requested. In this situation, consider using a SSL/TLS version fallback scheme:

- 1. Fall back from TLSv1.2 to TLSv1.1 if the server does not understand TLSv1.2.
- 2. Fall back from TLSv1.1 to TLSv1.0 if the previous step does not work.

For example, if the enabled protocol list on the client is TLSv1, TLSv1.1, and TLSv1.2, a typical SSL/TLS version fallback scheme may look like:

- 1. Try to connect to server. If server rejects the SSL/TLS connection request immediately, go to step 2.
- 2. Try the version fallback scheme by removing the highest protocol version (for example, TLSv1.2 for the first failure) in the enabled protocol list.
- 3. Try to connect to the server again. If server rejects the connection, go to step 2 unless there is no version to which the server can fall back.
- 4. If the connection fails and SSLv2Hello is not on the enabled protocol list, restore the enable protocol list and enable SSLv2Hello. (For example, the enable protocol list should be SSLv2Hello, TLSv1, TLSv1.1, and TLSv1.2.) Start again from step 1.

Note:

A fallback to a previous version normally means security strength downgrading to a weaker protocol. It is not suggested to use a fallback scheme unless it is really necessary, and you clearly know that the server does not support a higher protocol version.

Note:

As part of disabling SSLv3, some servers have also disabled SSLv2Hello, which means communications with SSLv2Hello-active clients (JDK 6u95) will fail. Starting with JDK 7, SSLv2Hello default to disabled on clients, enabled on servers.



SunJSSE Cannot Find a JCA Provider That Supports a Required Algorithm and Causes a NoSuchAlgorithmException

Problem: A handshake is attempted and fails when it cannot find a required algorithm. Examples might include:

```
Exception in thread ...deleted...
...deleted...
Caused by java.security.NoSuchAlgorithmException: Cannot find any provider supporting RSA/ECB/PKCS1Padding

Or
Caused by java.security.NoSuchAlgorithmException: Cannot find any provider supporting AES/CBC/NoPadding
```

Cause: SunJSSE uses JCE for all its cryptographic algorithms. If the SunJCE provider has been deregistered from the Provider mechanism and an alternative implementation from JCE is not available, then this exception will be thrown.

Solution: Ensure that the SunJCE is available by checking that the provider is registered with the Provider interface. Try to run the following code in the context of your SSL connection:

```
import javax.crypto.*;
System.out.println("=====Where did you get AES=====");
Cipher c = Cipher.getInstance("AES/CBC/NoPadding");
System.out.println(c.getProvider());
```

FailedDownloadException Thrown When Trying to Obtain Application Resources from Web Server over SSL

Problem: If you receive a <code>com.sun.deploy.net.FailedDownloadException</code> when trying to obtain application resources from your web server over SSL, and your web server uses the virtual host with Server Name Indication (SNI) extension (such as Apache HTTP Server), then you may have not configured your web server correctly.

Cause: Because Java SE 7 supports the SNI extension in the JSSE client, the requested host name of the virtual server is included in the first message sent from the client to the server during the SSL handshake. The server may deny the client's request for a connection if the requested host name (the server name indication) does not match the expected server name, which should be specified in the virtual host's configuration. This triggers an SSL handshake unrecognized name alert, which results in a FailedDownloadException being thrown.

Solution: To better diagnose the problem, enable tracing through the Java Console. See Debugging and Java Console in *Java Platform, Standard Edition Deployment Guide*. If the cause of the problem is <code>javax.net.ssl.SSLProtocolException: handshake alert: unrecognized_name</code>, it is likely that the virtual host configuration for SNI is incorrect. If you are using Apache HTTP Server, see Name-based Virtual Host Support about configuring virtual hosts. In particular, ensure that the <code>ServerName</code> directive is configured properly in a <code>VirtualHost</code> block.

See the following:

- SSL with Virtual Hosts Using SNI from Apache HTTP Server Wiki
- SSL/TLS Strong Encryption: FAQ from Apache HTTP Server Documentation
- RFC 3546, Transport Layer Security (TLS) Extensions
- Bug 7194590: SSL handshaking error caused by virtual server misconfiguration

IllegalArgumentException When RC4 Cipher Suites are Configured for DTLS

Problem: An IllegalArgumentException exception is thrown when RC4 cipher suite algorithm is specified in SSLEngine.setEnabledCipherSuites(String[] suites) method and the SSLEngine is a DTLS engine.

```
sslContext = SSLContext.getInstance("DTLS");

// Create the engine
SSLEngine engine = sslContext.createSSLengine(hostname, port);

String enabledSuites[] = { "SSL_RSA_WITH_RC4_128_SHA" };
engine.setEnabledCipherSuites(enabledSuites);
```

Cause: According to DTLS Version 1.0 and DTLS Version 1.2, RC4 cipher suites must not be used with DTLS.

Solution: Do not use RC4 based cipher suites for DTLS connections. See "JSSE Cipher Suite Names" in Java Security Standard Algorithm Names.

Debugging Utilities

JSSE provides dynamic debug tracing support. This is similar to the support used for debugging access control failures in the Java SE platform. The generic Java dynamic debug tracing support is accessed with the <code>java.security.debug</code> system property, whereas the JSSE-specific dynamic debug tracing support is accessed with the <code>java.net.debug</code> system property.



The debug utility is not an officially supported feature of JSSE.

To view the options of the JSSE dynamic debug utility, use the following command-line option on the <code>java</code> command:

-Djavax.net.debug=help



If you specify the value help with either dynamic debug utility when running a program that does not use any classes that the utility was designed to debug, you will not get the debugging options.



The following complete example shows how to get a list of the debug options for an application named MyApp that uses some of the JSSE classes:

```
java -Djavax.net.debug=help MyApp
```

The MyApp application will not run after the debug help information is printed, as the help code causes the application to exit.

Current options are:

- all: Turn on all debugging
- ssl: Turn on SSL debugging

The following can be used with the ssl option:

- record: Enable per-record tracing
- handshake: Print each handshake message
- keygen: Print key generation data
- session: Print session activity
- defaultctx: Print default SSL initialization
- sslctx: Print SSLContext tracing
- sessioncache: Print session cache tracing
- keymanager: Print key manager tracing
- trustmanager: Print trust manager tracing

Messages generated from the handshake option can be widened with these options:

- data: Hex dump of each handshake message
- verbose: Verbose handshake message printing

Messages generated from the record option can be widened with these options:

- plaintext: Hex dump of record plaintext
- packet: Print raw SSL/TLS packets

The <code>javax.net.debug</code> property value must be either <code>all</code> or <code>ssl</code>, optionally followed by debug specifiers. You can use one or more options. You do *not* have to have a separator between options, although a separator such as a colon (:) or a comma (,) helps readability. It does not matter what separators you use, and the ordering of the option keywords is also not important.

For an introduction to reading this debug information, see the guide, Debugging SSL/TLS Connections.

The following are examples of using the <code>javax.net.debug</code> property:

To view all debugging messages:

```
java -Djavax.net.debug=all MyApp
```

• To view the hexadecimal dumps of each handshake message, enter the following (the colons are optional):

```
java -Djavax.net.debug=ssl:handshake:data MyApp
```



• To view the hexadecimal dumps of each handshake message, and to print trust manager tracing, enter the following (the commas are optional):

```
java -Djavax.net.debug=SSL,handshake,data,trustmanager MyApp
```

Debugging SSL/TLS Connections

Understanding SSL/TLS connection problems can sometimes be difficult, especially when it is not clear what messages are actually being sent and received. The SunJSSE has a built-in debug facility and is activated by the System property javax.net.debug.

To know more about javax.net.debug System property, see Debugging Utilities.

What follows is a brief example how to read the debug output. Please be aware that the output is non-standard, and may change from release to release. We are using the default SunJSSE X509KeyManager and X509TrustManager which prints debug information.

This example assumes a basic understanding of the SSL/TLS protocol. To know more about protocols (handshake messages, etc.), see Secure Sockets Layer (SSL) Protocol Overview.

In this example, we connect using the SSLSocketClientWithClientAuth sample application to a simple HTTPS server that requires client authentication, then send a HTTPS request and receive the reply.

```
java -Djavax.net.debug=all \
    -Djavax.net.ssl.trustStore=trustStore
    SSLSocketClientWithClientAuth bongos 2001 /index.html
```

First, the X509KeyManager is initialized and discovers there is one keyEntry in the supplied KeyStore for a subject called "duke". If a server requests a client to authenticate itself, the X509KeyManager will search its list of keyEntries for an appropriate credential.

```
found key for : duke
chain [0] = [
 Version: V1
 Subject: CN=Duke, OU=Java Software, O="Sun Microsystems, Inc.",
 L=Cupertino, ST=CA, C=US
 Signature Algorithm: MD5withRSA, OID = 1.2.840.113549.1.1.4
 Key: Sun RSA public key, 1024 bits
 modulus: 134968166047563266914058280571444028986498087544923991226919517
 704208851688212064483570055963805034839797994154526862998272017486468599
 962268346037652120279791547218281230795146025359480589335682217749874703
 510467348902637769973696151441
 public exponent: 65537
 Validity: [From: Tue May 22 16:46:46 PDT 2001,
             To: Sun May 22 16:46:46 PDT 2011]
 Issuer: CN=Duke, OU=Java Software, O="Sun Microsystems, Inc.",
 L=Cupertino, ST=CA, C=US
 SerialNumber: [
                  3b0afa66]
```



The X509TrustManager is next initialized, and finds a certificate for a Certificate Authority (CA) named "JSSE Test CA". Any server presenting

valid

credentials signed by this CA will be trusted.

```
trustStore is: trustkeys
trustStore type is : jks
trustStore provider is :
init truststore
adding as trusted cert:
   Subject: CN=JSSE Test CA, OU=JWS, O=Sun,
        L=Santa Clara, ST=CA, C=US
   Issuer: CN=JSSE Test CA, OU=JWS, O=Sun,
        L=Santa Clara, ST=CA, C=US
   Algorithm: RSA; Serial number: 0x0
   Valid from Mon Jul 19 13:30:15 PDT 2004 until Fri Dec 05 12:30:15 PST 2031
```

We finish some additional initialization code, and after this, we are now finally ready to make the connection to the server.

```
trigger seeding of SecureRandom
done seeding SecureRandom
export control - checking the cipher suites
export control - no cached value available...
export control - storing legal entry into cache...
%% No cached client session
```

The connection to the server is made, and we see the initial ClientHello message, which contains:

- random information to initialize the cryptographic routines,
- the SessionID, which if non-null, would be used in reestablishing a previous session,
- · the list of ciphersuites that the client requests,
- and no compression algorithms.

This is followed by the output of various filters, such as encapsulating the TLSv1 header into the SSLv2Hello header format (See setEnabledProtocols()).

```
*** ClientHello, TLSv1
```



```
RandomCookie: GMT: 1073239164 bytes = { 10, 80, 71, 86, 124, 135, 104,
151, 72, 153, 70, 28, 97, 232, 160, 217, 146, 178, 87, 255, 122, 147, 83,
197, 60, 187, 227, 76 }
Session ID: {}
Cipher Suites: [SSL_RSA_WITH_RC4_128_MD5, SSL_RSA_WITH_RC4_128_SHA,
TLS_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA,
TLS_DHE_DSS_WITH_AES_128_CBC_SHA, SSL_RSA_WITH_3DES_EDE_CBC_SHA,
SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA, SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA,
SSL_RSA_WITH_DES_CBC_SHA, SSL_DHE_RSA_WITH_DES_CBC_SHA,
SSL_DHE_DSS_WITH_DES_CBC_SHA, SSL_RSA_EXPORT_WITH_RC4_40_MD5,
SSL_RSA_EXPORT_WITH_DES40_CBC_SHA, SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA,
SSL DHE DSS EXPORT WITH DES40 CBC SHA]
Compression Methods: { 0 }
[write] MD5 and SHA1 hashes: len = 73
0000: 01 00 00 45 03 01 40 F8 54 7C 0A 50 47 56 7C 87 ...E..@.T..PGV..
0010: 68 97 48 99 46 1C 61 E8 A0 D9 92 B2 57 FF 7A 93 h.H.F.a....W.z.
0040: 03 00 08 00 14 00 11 01
                          0.0
main, WRITE: TLSv1 Handshake, length = 73
[write] MD5 and SHA1 hashes: len = 98
0000: 01 03 01 00 39 00 00 00 20 00 04 01 00 80 00 ....9...
0010: 00 05 00 00 2F 00 00 33 00 00 32 00 00 0A 07 00 ..../..3..2.....
0030: 00 00 12 00 00 03 02 00 80 00 00 08 00 00 14 00 .....
0040: 00 11 40 F8 54 7C 0A 50 47 56 7C 87 68 97 48 99 ..@.T..PGV..h.H.
0050: 46 1C 61 E8 A0 D9 92 B2 57 FF 7A 93 53 C5 3C BB F.a....W.z.S.<.
0060: E3 4C
                                                .L
main, WRITE: SSLv2 client hello message, length = 98
```

Section labeled "[Raw write]" represent the actual data sent to the raw output object (in this case, an OutputStream).

```
[Raw write]: length = 100

0000: 80 62 01 03 01 00 39 00 00 00 00 00 04 01 00 .b...9.....

0010: 80 00 00 05 00 00 2F 00 00 33 00 00 32 00 00 0A ..../.3..2...

0020: 07 00 00 00 16 00 00 13 00 00 99 06 00 40 00 .....@.

0030: 00 15 00 00 12 00 00 03 02 00 80 00 00 80 00 ..........

0040: 14 00 00 11 40 F8 54 7C 0A 50 47 56 7C 87 68 97 ....@.T.PGV.h.

0050: 48 99 46 1C 61 E8 A0 D9 92 B2 57 FF 7A 93 53 C5 H.F.a....W.z.S.

0060: 3C BB E3 4C <..L
```

After sending the initial ClientHello, we wait for the server's response, a ServerHello. "[Raw read]" displays the raw data read from the input device (InputStream), before any processing has been performed.

```
[Raw read]: length = 5
0000: 16 03 01 06 F0
[Raw read]: length = 1776
0000: 02 00 00 46 03 01 40 FC 31 10 79 AB 17 66 FA 8B ...F..@.l.y..f..
0010: 3F AA FD 5E 48 23 FA 90
                             31 D8 3C B9 A3 2C 8C F5 ?..^H#..1.<....
0020: E9 81 9B A2 63 6C 20 40
                              FC 31 10 BD 8D A5 91 06 ....cl @.1.....
0030: 8B E1 E6 80 C6 5A 5C D9
                              8D 0A AE CA 58 4A BA 36 .....Z\.....XJ.6
0040: B1 3D 04 8D 82 21 B4 00
                              04 00 0B 00 06 1B 00 06 .=...!.....
0050: 18 00 03 11 30 82 03 0D
                              30 82 02 76 A0 03 02 01 ....0...v....
0060: 02 02 01 01 30 0D 06 09
                              2A 86 48 86 F7 0D 01 01 ....0...*.H.....
0070: 04 05 00 30 63 31 0B 30
                             09 06 03 55 04 06 13 02 ...0c1.0...U....
0080: 55 53 31 0B 30 09 06 03 55 04 08 13 02 43 41 31 US1.0...U....CA1
```



0090:	14	30	12	06	03	55	04	07	13	0В	53	61	6E	74	61	20	.0USanta
00A0:	43	6C	61	72	61	31	0C	30	0A	06	0.3	55	04	0A	13	0.3	Clara1.0U
00B0:														03			
																	Sun1.0UJW
00C0:	53	31	15	30	13	06	03	55	04	03	13	0C	4A	53	53	45	S1.0UJSSE
00D0:	20	54	65	73	74	20	43	41	3.0	1 E	17	OΠ	30	34	30	37	Test CA00407
00E0:	3 I	39	32	30	33	30	35	31	5A	17	UD	33	3 L	31	32	30	19203051Z31120
00F0:	35	32	30	33	30	35	31	5A	30	48	31	0B	30	09	06	03	5203051Z0H1.0
0100:	55	٥4	06	12	0.2	55	53	31	ΛR	3 0	nα	06	ΛZ	55	Λ4	ΛR	UUS1.0U
0110:	13	02	43	41	31	0C	30	0A	06	03	55	04	0A	13	03	53	CA1.0US
0120:	75	6E	31	0D	30	0B	06	0.3	55	04	0B	13	04	4A	61	76	un1.0UJav
0130:														6F			a1.0Ubong
0140:	бF	73	30	81	9F	30	0D	06	09	2A	86	48	86	F7	0D	01	os00*.H
0150:	01	01	05	0.0	03	81	ЯD	0.0	3.0	81	89	02	81	81	0.0	CC	0
0160:	09	/4	CB	43	AB	6D	ΕD	F.0	35	AA	UΕ	49	29	D9	ΕU	F.O	.t.C.m5I)
0170:	Α1	D5	E2	3E	8F	5E	C5	CE	F4	DE	C1	Α4	F3	CB	8C	45	>.^E
0180:	Π۵	٥٣	6 F	21	r 1	٥٥	65	CD	30	ח1	50	55	67	FB	50	96	n!e.<.\.j.].
0190:	93	F4	71	41	41	45	FF	37	86	4C	AΒ	F9	EΑ	9A	3F	Α5	qAAE.7.L?.
01A0:	82	60	BF	0A	81	84	C9	3E	AC	0F	3D	20	3D	AC	Α0	69	.`>= =i
01B0:														25			J7fR.%C.
01C0:	10	44	07	1E	93	74	D9	68	01	D7	06	48	C9	0D	52	2D	.Dt.hHR-
01D0:	D5	64	2₽	Δ6	48	4C	59	E2	50	C6	C1	50	СВ	4C	1 R	02	.jHLY.\\.L
01E0:	03	01	00	01	Α3	81	EB	30	81	E8	30	09	06	03	55	1D	
01F0:	13	04	02	30	00	30	2C	06	09	60	86	48	01	86	F8	42	0.0,`.HB
0200:	Λ1	Δ٦	0.4	1 🗁	16	1 D	1 E	70	65	6 E	E 2	E 2	10	20	17	65	OpenSSL Ge
0210:	6E	65	72	61	74	65	64	20	43	65	72	74	69	66	69	63	nerated Certific
0220:	61	74	65	30	1 D	06	03	55	ח1	0E	04	16	0.4	14	58	D7	ate0UX.
0230:														8E			:.7.>.'E*.
0240:	77	28	30	81	8D	06	03	55	1D	23	04	81	85	30	81	82	w(0U.#0
0250:	٩n	14	ΛR	Δ3	7 E	35	96	15	FΔ	RΛ	F5	1 R	58	CD	4 F	54	5OT
0260:	EF.	3 L	33	70	Ŀ4	Α'/	AΙ	6.7	A4	65	30	63	3 L	0B	30	09	.13pg.e0c1.0.
0270:	06	03	55	04	06	13	02	55	53	31	0B	30	09	06	03	55	UUS1.0U
0280:														04			
																	CA1.0U
0290:	0В	53	61	6E	74	61	20	43	6C	61	72	61	31	0C	30	0A	.Santa Clara1.0.
02A0:	06	03	55	0.4	0Α	13	03	53	75	6E	31	0C	30	0A	06	03	USun1.0
02B0:	55	04	ÛВ	13	03	4A	5 /	53						03			UJWS1.0U.
02C0:	03	13	0C	4A	53	53	45	20	54	65	73	74	20	43	41	82	JSSE Test CA.
02D0:	Λ1	٥٥	3 0	UΠ	06	nα	27	86	4 Q	86	F7	UΠ	0.1	01	Λ4	05	0*.H
02E0:	00	03	81	81	00	05	3E	17	DA	F2	05	СВ	4E	9E	BF	12	>N
02F0:	CE	13	76	FF	В2	FB	7F	9C	3D	45	28	43	6C	98	28	E3	v=E(Cl.(.
0300:	0.2	17	an	06	₽1	62	CΛ	60	an	DΛ	ъc	E6	70	4C	25	aa	b.`L/.
0310:	40	FE	06	CB	34	60	В1	F4	26	1C	E8	46	39	24	Ε1	8A	@4`&F9\$
0320:	71	F2	13	90	Α4	0.A	7в	0B	13	AB	51	68	53	D9	7A	31	qQhS.z1
0330:																	
														5D			Z <d,ipw%].b< td=""></d,ipw%].b<>
0340:	7F	3E	61	1F	29	Α3	31	46	02	C6	D2	8C	27	79	40	76	.>a.).1F'y@v
0350:	97	B6	25	19	BE	60	6 A	92	חכ	EF	11	BE	F.7	4A	FF	2 A	%ljJ.*
0360:														82			910f
0370:	Α0	03	02	01	02	02	01	00	30	0D	06	09	2A	86	48	86	0*.H.
0380:	F7	UΠ	01	0.1	0.4	05	٥٥	3.0	63	31	ΛR	3.0	n q	06	03	55	Oc1.0U
0390:	04	06	13	02	55	53	31	UВ	30	09	06	03	55	04	08	13	US1.0U
03A0:	02	43	41	31	14	30	12	06	0.3	55	04	07	13	0B	53	61	.CA1.0USa
03B0:														06			nta Clara1.0U
03C0:	04	0A	13	03	53	75	6E	31	0C	30	0A	06	03	55	04	0B	Sun1.0U
03D0:	13	03	4 A	57	53	31	15	30	13	06	03	55	04	03	13	0C	JWS1.0U
03E0:	4A													1E			JSSE Test CA0
0250.			2 0	37	31	39	32	30	33	30	31	35	5A	17	0D	33	040719203015z3
03F0:		34	30										30				
	30				35	30	3 (1	3.3									1 1 2015 2013 0 15 20 21
0400:	30 31	31	32	30													11205203015Z0c1.
	30 31	31	32	30										30			0UUS1.0
0400: 0410:	30 31 30	31 09	32 06	30 03	55	04	06	13	02	55	53	31	0B		09	06	0UUS1.0
0400: 0410: 0420:	30 31 30 03	31 09 55	32 06 04	30 03 08	55 13	04 02	06 43	13 41	02 31	55 14	53 30	31 12	0B 06	30 03	09 55	06 04	0UUS1.0 .UCA1.0U.
0400: 0410: 0420: 0430:	30 31 30 03 07	31 09 55 13	32 06 04 0B	30 03 08 53	55 13 61	04 02 6E	06 43 74	13 41 61	02 31 20	55 14 43	53 30 6C	31 12 61	0B 06 72	30 03 61	09 55 31	06 04 0C	0UUS1.0 .UCA1.0U. Santa Clara1.
0400: 0410: 0420:	30 31 30 03 07	31 09 55 13	32 06 04 0B	30 03 08 53	55 13 61	04 02 6E	06 43 74	13 41 61	02 31 20	55 14 43	53 30 6C	31 12 61	0B 06 72	30 03	09 55 31	06 04 0C	0UUS1.0 .UCA1.0U.
0400: 0410: 0420: 0430:	30 31 30 03 07 30	31 09 55 13 0A	32 06 04 0B 06	30 03 08 53 03	55 13 61 55	04 02 6E 04	06 43 74 0A	13 41 61 13	02 31 20 03	55 14 43 53	53 30 6C 75	31 12 61 6E	0B 06 72 31	30 03 61	09 55 31 30	06 04 0C 0A	0UUS1.0 .UCA1.0U. Santa Clara1.

```
0460: 55 04 03 13 0C 4A 53 53 45 20 54 65 73 74 20 43 U....JSSE Test C
0480: 01 05 00 03 81 8D 00 30 81 89 02 81 81 00 9A 0A .....0.....
0490: B6 45 66 D5 DE 4A D9 3C 8C AC A6 B5 A5 88 B4 CF .Ef..J.<.....
04A0: 14 E1 A6 1B 25 25 4F 44 C9 1F 22 38 32 29 CF A1 ....%*OD.."82)..
04B0: 7C 18 30 93 DC 2B EC 2B 67 EE 2E 08 66 2D 0F 47 ..0..+.+g...f-.G
04D0: 55 7D 32 3E F9 66 A2 DD 55 EB 4D 0A 67 C7 5D 21 U.2>.f..U.M.q.]!
04E0: 9B 29 EA 2E 51 C5 83 A3 55 FF 35 CA A6 99 8F 46 .)..Q...U.5....F
04F0: F8 8E 56 BB A2 B1 39 83 D8 61 42 79 E0 95 78 FA ..V...9..aBy..x.
0500: C6 E3 65 B0 FD 74 2D 64 51 71 04 F2 A1 91 02 03 ..e..t-dQq......
                         BD 30 1D 06 03 55 1D 0E .....0...U...
0510: 01 00 01 A3 81 C0 30 81
                         96 15 FA BO F5 1B 5F CD .....5....._.
0520: 04 16 04 14 08 A3 7E 35
                         30 81 8D 06 03 55 1D 23 OT.13p..0...U.#
0530: 4F 54 EF 31 33 70 E4 A7
0540: 04 81 85 30 81 82 80 14 08 A3 7E 35 96 15 FA B0 ...0.....5....
0550: F5 1B 5F CD 4F 54 EF 31 33 70 E4 A7 A1 67 A4 65 .._.OT.13p...g.e
0560: 30 63 31 0B 30 09 06 03 55 04 06 13 02 55 53 31 0c1.0...U....US1
0580: 06 03 55 04 07 13 0B 53 61 6E 74 61 20 43 6C 61 ..U....Santa Cla
0590: 72 61 31 0C 30 0A 06 03 55 04 0A 13 03 53 75 6E ral.0...U....Sun
05A0: 31 OC 30 OA 06 03 55 O4 OB 13 O3 4A 57 53 31 15 1.0...U....JWS1.
05B0: 30 13 06 03 55 04 03 13 0C 4A 53 53 45 20 54 65 0...U....JSSE TE
05CO: 73 74 20 43 41 82 01 00 30 0C 06 03 55 1D 13 04 st CA...O...U...
05D0: 05 30 03 01 01 FF 30 0D 06 09 2A 86 48 86 F7 0D .0...0...*.H...
05E0: 01 01 04 05 00 03 81 81 00 73 6A 46 A2 05 E3 D8 .....sjF....
0600: C7 FA E4 ED 5C 4F 71 22 FB 26 E3 01 3D 0C 10 AA ....\Oq".&..=...
0620: D6 FB 42 6C B4 A9 A9 57 A5 84 00 42 6D 37 37 6D ..Bl...W...Bm77m
0630: C7 6C 23 BC DC 60 D1 9D 6F B3 75 47 3A 15 33 1A .1#..`..o.uG:.3.
0640: EC 90 09 9D F9 EB BD 88
                         96 E7 1D 41 BC 01 8D CA .....A....
0650: 88 D9 5B 04 09 8F 3E EA
                         C8 15 A0 AA 4E 85 95 AE ..[...>....N...
0660: 2F 0E 31 92 AC 3C FB 2F
                         C4 OD OO OO 7F O2 O1 O2 /.1..<./....
0670: 00 7A 00 78 30 76 31 0B
                         30 09 06 03 55 04 06 13 .z.x0v1.0...U...
0680: 02 55 53 31 0B 30 09 06
                         03 55 04 08 13 02 43 41 .US1.0...U....CA
0690: 31 12 30 10 06 03 55 04
                         07 13 09 43 75 70 65 72 1.0...U....Cuper
                         06 03 55 04 0A 13 16 53 tino1.0...U....S
06A0: 74 69 6E 6F 31 1F 30 1D
06B0: 75 6E 20 4D 69 63 72 6F
                         73 79 73 74 65 6D 73 2C un Microsystems.
06D0: 4A 61 76 61 20 53 6F 66 74 77 61 72 65 31 0D 30 Java Softwarel.0
06E0: 0B 06 03 55 04 03 13 04 44 75 6B 65 0E 00 00 00 ...U....Duke....
main, READ: TLSv1 Handshake, length = 1776
```

The data is unpackaged, and if the message is in the SSL/TLS format, it is parsed into a ServerHello. If you connected to a non-SSL/TLS socket (plaintext?), the received data will not be in SSL/TLS format, and you'll have problems connecting.

The ServerHello specifies several things:

- The server's random data, also used to initialize the cryptographic algorithms,
- the identifier of this session (if the client wants to try to rejoin this session using a
 different connection, it can send this ID in its ClientHello. If the client session ID
 equals the server session ID, an abbreviated handshake takes place, and the
 previously established parameters are used),
- the selected cipher suite,
- and the compression method (none in this case).

Lastly note that the ServerHello has specified that the connection should use "TLSv1", rather than "SSLv3."

```
*** ServerHello, TLSv1
RandomCookie: GMT: 1073492240 bytes = { 121, 171, 23, 102, 250, 139, 63,
170, 253, 94, 72, 35, 250, 144, 49, 216, 60, 185, 163, 44, 140, 245, 233,
129, 155, 162, 99, 108 }
Session ID: {64, 252, 49, 16, 189, 141, 165, 145, 6, 139, 225, 230, 128,
198, 90, 92, 217, 141, 10, 174, 202, 88, 74, 186, 54, 177, 61, 4, 141, 130,
33, 180}
Cipher Suite: SSL_RSA_WITH_RC4_128_MD5
Compression Method: 0
%% Created: [Session-1, SSL_RSA_WITH_RC4_128_MD5]
** SSL RSA WITH RC4 128 MD5
[read] MD5 and SHA1 hashes: len = 74
0000: 02 00 00 46 03 01 40 FC 31 10 79 AB 17 66 FA 8B ...F..@.1.y..f..
0010: 3F AA FD 5E 48 23 FA 90 31 D8 3C B9 A3 2C 8C F5 ?..^H#..1.<....
0030: 8B E1 E6 80 C6 5A 5C D9 8D 0A AE CA 58 4A BA 36 .....Z\.....XJ.6
0040: B1 3D 04 8D 82 21 B4 00 04 00
                                                   .=...!....
```

The server next identifies itself to the client by passing a Certificate chain. In this example, we have a certificate for the subject "bongos", signed by the issuer "JSSE Test CA". We know that "JSSE Test CA" is a trusted CA, so if the certificate chain verifies correctly by our X509TrustManager, we can accept this connection.

There are many different ways of establishing trust, so if the default X509TrustManager is not doing the types of trust management you need, you can supply your own X509TrustManager to the SSLContext.

```
*** Certificate chain
chain [0] = [
[
  Version: V3
 Subject:
  CN=bongos, OU=Java, O=Sun, ST=CA, C=US
  Signature Algorithm: MD5withRSA, OID = 1.2.840.113549.1.1.4
 Key: Sun RSA public key, 1024 bits
  modulus: 143279610700427050704216702734995283650706638118826657356308087
  682552751165540126665070195006746918193702313900836063802045448392771274
  463088345157808670190122017153821642985630288017629294930800445939721128
  735250668515619736933648548512047941708018130926985936894512063397816602
  623867976474763783110866258971
  public exponent: 65537
 Validity: [From: Mon Jul 19 13:30:51 PDT 2004,
             To: Fri Dec 05 12:30:51 PST 2031]
  Issuer: CN=JSSE Test CA, OU=JWS, O=Sun,
     L=Santa Clara, ST=CA, C=US
  SerialNumber: [
                   011
Certificate Extensions: 3
[1]: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0010: 2A 8B 77 28
                                                     *.w(
]
]
```



```
[2]: ObjectId: 2.5.29.35 Criticality=false
AuthorityKeyIdentifier [
KeyIdentifier [
0000: 08 A3 7E 35 96 15 FA B0 F5 1B 5F CD 4F 54 EF 31 ...5......OT.1
0010: 33 70 E4 A7
                                              3p..
[CN=JSSE Test CA, OU=JWS, O=Sun, L=Santa Clara, ST=CA, C=US]
SerialNumber: [
              0.01
[3]: ObjectId: 2.5.29.19 Criticality=false
BasicConstraints:[
CA:false
PathLen: undefined
 Algorithm: [MD5withRSA]
 Signature:
0000: 05 3E 17 DA F2 05 CB 4E 9E BF 12 CE 13 76 FF B2 .>....N....v..
0010: FB 7F 9C 3D 45 28 43 6C 98 28 E3 92 17 C2 C6 F1 ...=E(Cl.(.....
0040: 0A 7B 0B 13 AB 51 68 53 D9 7A 31 5A C1 7E 3C 44 .....Ohs.z1z..<D
0050: 2C 49 70 57 25 F9 18 FE 5D A5 42 7F 3E 61 1F 29 ,IpW%...].B.>a.)
chain [1] = [
 Version: V3
 Subject: CN=JSSE Test CA, OU=JWS, O=Sun,
    L=Santa Clara, ST=CA, C=US
 Signature Algorithm: MD5withRSA, OID = 1.2.840.113549.1.1.4
 Key: Sun RSA public key, 1024 bits
 modulus: 108171861314934294923646852258201093253619460299818135230481040
 615923025149195168140458238629251726950220398889722590740552079782864577
 976838691751841449679901644183317203824143803940037883199193775839934767
 304560313841716869284745769157293013188246601563271959824290073095150730
 505329011956986145636520993169
 public exponent: 65537
 Validity: [From: Mon Jul 19 13:30:15 PDT 2004,
           To: Fri Dec 05 12:30:15 PST 2031]
 Issuer: CN=JSSE Test CA, OU=JWS, O=Sun,
    L=Santa Clara, ST=CA, C=US
 SerialNumber: [
                001
Certificate Extensions: 3
[1]: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0010: 33 70 E4 A7
                                              3p..
1
[2]: ObjectId: 2.5.29.35 Criticality=false
```

```
AuthorityKeyIdentifier [
KeyIdentifier [
0000: 08 A3 7E 35 96 15 FA B0 F5 1B 5F CD 4F 54 EF 31 ...5......OT.1
0010: 33 70 E4 A7
                                                   3p..
[CN=JSSE Test CA, OU=JWS, O=Sun, L=Santa Clara, ST=CA, C=US]
SerialNumber: [
               00]
[3]: ObjectId: 2.5.29.19 Criticality=false
BasicConstraints:[
CA:true
PathLen: 2147483647
]
 Algorithm: [MD5withRSA]
 Signature:
0000: 73 6A 46 A2 05 E3 D8 6E 5C F4 18 A2 74 BC CF EB sjF....n\...t...
0010: 0C 5B FF 81 1C 28 85 C7 FA E4 ED 5C 4F 71 22 FB .[...(....\Oq".
0020: 26 E3 01 3D 0C 10 AA BB 3E 90 ED 0E 1F 0C 9B B1 &..=...>....
0030: 8C 49 6A 51 E4 C3 52 D6 FB 42 6C B4 A9 A9 57 A5 .ijq..R..Bl...W.
0040: 84 00 42 6D 37 37 6D C7 6C 23 BC DC 60 D1 9D 6F ..Bm77m.1#..`..o
0050: B3 75 47 3A 15 33 1A EC 90 09 9D F9 EB BD 88 96 .uG:.3......
0060: E7 1D 41 BC 01 8D CA 88 D9 5B 04 09 8F 3E EA C8 ..A....[...>..
]
```

We recognize this cert! We can trust it, and continue on with the handshake.

```
Found trusted certificate:
ſ
```

[

]]

```
Version: V3
  Subject: CN=JSSE Test CA, OU=JWS, O=Sun, L=Santa Clara, ST=CA, C=US
  Signature Algorithm: MD5withRSA, OID = 1.2.840.113549.1.1.4
  Key: Sun RSA public key, 1024 bits
  modulus: 108171861314934294923646852258201093253619460299818135230481040
  615923025149195168140458238629251726950220398889722590740552079782864577
  976838691751841449679901644183317203824143803940037883199193775839934767
  304560313841716869284745769157293013188246601563271959824290073095150730
  505329011956986145636520993169
  public exponent: 65537
  Validity: [From: Mon Jul 19 13:30:15 PDT 2004,
             To: Fri Dec 05 12:30:15 PST 2031]
  Issuer: CN=JSSE Test CA, OU=JWS, O=Sun, L=Santa Clara, ST=CA, C=US
  SerialNumber: [
Certificate Extensions: 3
[1]: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0010: 33 70 E4 A7
                                                    3p..
```



```
[2]: ObjectId: 2.5.29.35 Criticality=false
AuthorityKeyIdentifier [
KeyIdentifier [
0010: 33 70 E4 A7
                                               3p..
[CN=JSSE Test CA, OU=JWS, O=Sun, L=Santa Clara, ST=CA, C=US]
SerialNumber: [
              0.01
[3]: ObjectId: 2.5.29.19 Criticality=false
BasicConstraints:[
CA:true
PathLen: 2147483647
1
 Algorithm: [MD5withRSA]
 Signature:
0000: 73 6A 46 A2 05 E3 D8 6E 5C F4 18 A2 74 BC CF EB sjF....n\...t...
0010: 0C 5B FF 81 1C 28 85 C7 FA E4 ED 5C 4F 71 22 FB .[...(.....\Oq".
0020: 26 E3 01 3D 0C 10 AA BB 3E 90 ED 0E 1F 0C 9B B1 &..=...>....
0030: 8C 49 6A 51 E4 C3 52 D6 FB 42 6C B4 A9 A9 57 A5 .ijo..R..Bl...W.
0040: 84 00 42 6D 37 37 6D C7 6C 23 BC DC 60 D1 9D 6F ..Bm77m.l#..`..o
0050: B3 75 47 3A 15 33 1A EC 90 09 9D F9 EB BD 88 96 .uG:.3......
0060: E7 1D 41 BC 01 8D CA 88 D9 5B 04 09 8F 3E EA C8 ..A.....[...>..
```

We read the next few bytes of the handshake...

```
[read] MD5 and SHA1 hashes: len = 1567
0020: 48 86 F7 0D 01 01 04 05
                        00 30 63 31 0B 30 09 06 H......0c1.0..
0030: 03 55 04 06 13 02 55 53
                        31 OB 30 O9 O6 O3 55 O4
                                           .U....US1.0...U.
0040: 08 13 02 43 41 31 14 30
                        12 06 03 55 04 07 13 0B
                                            ...CA1.0...U....
0050: 53 61 6E 74 61 20 43 6C
                        61 72 61 31 0C 30 0A 06 Santa Claral.0..
0060: 03 55 04 0A 13 03 53 75
                        6E 31 0C 30 0A 06 03 55 .U....Sun1.0...U
0070: 04 OB 13 03 4A 57 53 31
                        15 30 13 06 03 55 04 03 ....JWS1.0...U..
                        65 73 74 20 43 41 30 1E ...JSSE Test CAO.
0080: 13 OC 4A 53 53 45 20 54
0090: 17 OD 30 34 30 37 31 39
                        32 30 33 30 35 31 5A 17 ..040719203051z.
00A0: 0D 33 31 31 32 30 35 32 30 33 30 35 31 5A 30 48 .311205203051Z0H
00B0: 31 0B 30 09 06 03 55 04 06 13 02 55 53 31 0B 30 1.0...U....US1.0
00C0: 09 06 03 55 04 08 13 02 43 41 31 0C 30 0A 06 03 ...U....CA1.0...
00D0: 55 04 0A 13 03 53 75 6E 31 0D 30 0B 06 03 55 04 U....Sun1.0...U.
00F0: 13 06 62 6F 6E 67 6F 73 30 81 9F 30 0D 06 09 2A ..bongos0..0...*
0100: 86 48 86 F7 0D 01 01 01 05 00 03 81 8D 00 30 81 .H............
0110: 89 02 81 81 00 CC 09 74 CB 43 AB 6D ED F6 35 AA .....t.C.m..5.
0130: C1 A4 F3 CB 8C 45 0B 0F 6E 21 E1 00 65 CB 3C D1
                                            .....E..n!..e.<.
0140: 5C EF 6A FB 5D 96 93 F4
                        71 41 41 45 FF 37 86 4C \.j.]...qAAE.7.L
0150: AB F9 EA 9A 3F A5 82 60 BF 0A 81 84 C9 3E AC 0F ....?..`....>..
0160: 3D 20 3D AC AO 69 EF CA 4A A7 94 AD C8 A5 CE 37 = =..i..J......7
0170: 66 52 D1 25 43 CB 10 44 07 1E 93 74 D9 68 01 D7 fR.%C..D...t.h..
```



0190:	Cl	5C	C8	4C	1B	02	03	01	00	01	A3	81	EB	30	81	E8	.\.L0
01A0:	30	09	06	03	55	1D	13	04	02	30	00	30	2C	06	09	60	0U0.0,`
01B0:	86	48	01	86	F8	42	01	0D	04	1F	16	1D	4F	70	65	6E	.HBOpen
01C0:	53	53	4C	20	47	65	6E	65	72	61	74	65	64	20	43	65	SSL Generated Ce
01D0:									65		1D						rtificate0U
01E0:	04	16	04	14	58	D7	3A	Α9	37	AA	3E	14	27	FC	EC	CC	X.:.7.>.'
01F0:	45	08	04	8E	2A	8B	77	28	30	81	8D	06	03	55	1D	23	E*.w(0U.#
0200:	04	81	85	30	81	82	80	14	0.8	Α3	7E	35	96	15	FΑ	в0	05
0210:											E4						OT.13pg.e
0220:									55	04	06	13	02	55	53	3 L	0c1.0UUS1
0230:	0B	30	09	06	03	55	04	80	13	02	43	41	31	14	30	12	.0UCA1.0.
0240:	06	03	55	04	07	13	0B	53	61	бE	74	61	20	43	6C	61	USanta Cla
0250:	72	61	31	വറ	30	ΛΩ	06	03	55	0.4	0A	13	03	53	75	6E	ra1.0USun
0260:											03						1.0UJWS1.
0270:									0C	4A	53	53	45	20	54	65	0UJSSE Te
0280:	73	74	20	43	41	82	01	00	30	0D	06	09	2A	86	48	86	st CA0*.H.
0290:	F7	0D	01	01	04	05	00	03	81	81	00	05	3E	17	DA	F2	
02A0:	05	СВ	4 E	9E	BF	12	CE	13	76	FF	В2	FB	7F	9C	3D	45	N=E
02B0:											F1						(Cl.(b.`
02C0:											34						L/.@4`&.
02D0:	E8	46	39	24	Ε1	A8	71	F2	13	90	Α4	0A	7В	0В	13	AB	.F9\$q
02E0:	51	68	53	D9	7A	31	5A	C1	7E	3C	44	2C	49	70	57	25	QhS.z1Z <d,ipw%< td=""></d,ipw%<>
02F0:	F9	18	FE	5D	Α5	42	7F	3E	61	1 F	29	Α3	31	46	02	C6].B.>a.).1F
0300:											BE						'y@v%lj
																	=
0310:									AC	39	31	00	03	0.1	30	82	J.*910.
0320:	02	FD	30	82	02	66	Α0	03	02	01	02	02	01	00	30	0D	0f0.
0330:	06	09	2A	86	48	86	F7	0D	01	01	04	05	00	30	63	31	*.H0c1
0340:	0в	30	09	06	0.3	55	04	06	13	02	55	53	31	0B	30	09	.0UUS1.0.
0350:											14						UCA1.0U
0360:								74			43						Santa Claral
0370:	0C	30	0A	06	03	55	04	0A	13	03	53	75	6E	31	0C	30	.0USun1.0
0380:	0A	06	03	55	04	0B	13	03	4A	57	53	31	15	30	13	06	UJWS1.0
0390:	0.3	55	04	0.3	13	0C	4 A	53	53	45	20	54	65	73	74	20	.UJSSE Test
03A0:											31						CA00407192030
03B0:											35						15Z31120520301
03C0:	35	5A	30	63	31	0B	30	09	06	03	55	04	06	13	02	55	5Z0c1.0UU
03D0:	53	31	0B	30	09	06	03	55	04	80	13	02	43	41	31	14	S1.0UCA1.
03E0:	30	12	06	03	55	04	07	13	0B	53	61	6E	74	61	20	43	0USanta C
03F0:											55						lara1.0US
0400:																	
											0B						un1.0UJWS
0410:									03	13	0C	4A	53	53	45	20	1.0UJSSE
0420:	54	65	73	74	20	43	41	30	81	9F	30	0D	06	09	2A	86	Test CA00*.
0430:	48	86	F7	0D	01	01	01	05	00	03	81	8D	00	30	81	89	н
0440:	02	81	81	0.0	9A	0.A	В6	45	66	D5	DE	4 A	D9	3C	8C	AC	EfJ.<
0450:											25						%%OD
0460:											DC						"82)0+.+g.
0470:									3A	DC	ΕO	03	Ε9	65	16	F6	fG∶e
0480:	18	C6	16	14	56	24	55	7D	32	3E	F9	66	A2	DD	55	EΒ	V\$U.2>.fU.
0490:	4D	0A	67	C7	5D	21	9В	29	EA	2E	51	C5	83	A3	55	FF	M.g.]!.)QU.
04A0:											A2						5FV9a
04B0:											FD						Byxet-dQq
04C0:											81						00
04D0:	1D	06	03	55	1D	ΟE	04	16	04	14	08	A3	7E	35	96	15	U5
04E0:	FA	В0	F5	1в	5F	CD	4F	54	EF	31	33	70	E4	Α7	30	81	OT.13p0.
04F0:											81						U.#0
0500:											4F						.5OT.13p
0510:											30						g.e0c1.0U.
0520:	06	13	02	55	53	31	0B	30	09	06	03	55	04	08	13	02	US1.0U
0530:	43	41	31	14	30	12	06	03	55	04	07	13	0B	53	61	6E	CA1.0USan
0540:											30						ta Clara1.0U.
0550:											06						Sun1.0U
00000	υA	τ2	U S	JS	13	ОŢ	эΤ	UC	30	υA	00	U.S	22	υ 1	שט	т Э	built.uu

```
0560: 03 4A 57 53 31 15 30 13 06 03 55 04 03 13 0C 4A .JWS1.0...U...J
0570: 53 53 45 20 54 65 73 74 20 43 41 82 01 00 30 0C SSE Test CA...O.
0580: 06 03 55 1D 13 04 05 30 03 01 01 FF 30 0D 06 09 ..U...O...O...
0590: 2A 86 48 86 F7 0D 01 01 01 04 05 00 03 81 81 00 73 *.H.......s
05A0: 6A 46 A2 05 E3 D8 6E 5C F4 18 A2 74 BC CF EB 0C jF...n\...t...
05B0: 5B FF 81 1C 28 85 C7 FA E4 ED 5C 4F 71 22 FB 26 [...(....\Oq".&
05C0: E3 01 3D 0C 10 AA BB 3E 90 ED 0E 1F 0C 9B B1 8C ..=...>......
05D0: 49 6A 51 E4 C3 52 D6 FB 42 6C B4 A9 A9 57 A5 84 IjQ..R.Bl..W..
05E0: 00 42 6D 37 37 6D C7 6C 23 BC DC 60 D1 9D 6F B3 .Bm77m.l#..\...o.
05F0: 75 47 3A 15 33 1A EC 90 09 9D F9 EB BD 88 96 E7 uG:.3.........
0600: 1D 41 BC 01 8D CA 88 D9 5B 04 09 8F 3E EA C8 15 .A....[...>...
```

...and parse it to find that it's a CertificateRequest message.

The server is asking the client to identify itself with a X509 certificate subject having the common name (CN=) "Duke". The server's X509TrustManager has the option of rejecting any credentials provided by the client (or lack thereof). In a real-world situation, you'd probably use a certificate signed by a CA, and the list of trusted CA's would be included in this message instead.

```
*** CertificateRequest
Cert Types: RSA, DSS,
Cert Authorities:
CN=Duke, OU=Java Software, O="Sun Microsystems, Inc.",
L=Cupertino, ST=CA, C=US>
[read] MD5 and SHA1 hashes: len = 131
0010: 09 06 03 55 04 06 13 02 55 53 31 0B 30 09 06 03 ...U....US1.0...
0030: 13 09 43 75 70 65 72 74 69 6E 6F 31 1F 30 1D 06 ...Cupertino1.0..
0040: 03 55 04 0A 13 16 53 75 6E 20 4D 69 63 72 6F 73 .U....Sun Micros
0050: 79 73 74 65 6D 73 2C 20 49 6E 63 2E 31 16 30 14 ystems, Inc.1.0.
0060: 06 03 55 04 0B 13 0D 4A 61 76 61 20 53 6F 66 74 ...U....Java Soft
0070: 77 61 72 65 31 0D 30 0B 06 03 55 04 03 13 04 44 ware1.0...U....D
0080: 75 6B 65
                                              uke
*** ServerHelloDone
[read] MD5 and SHA1 hashes: len = 4
0000: 0E 00 00 00
```

We need to send client credentials back to the server, so the client's X509KeyManager is now consulted. We look for a match between the list of accepted issuers (above), and the certificates we have in our KeyStore. In this case (luckily?), there is a match: we have credentials for "duke". It's now up to the server's X509TrustManager to decide whether to accept these credentials.

```
matching alias: duke
*** Certificate chain
chain [0] = [
[
    Version: V1
    Subject: CN=Duke, OU=Java Software, O="Sun Microsystems, Inc.",
    L=Cupertino, ST=CA, C=US
    Signature Algorithm: MD5withRSA, OID = 1.2.840.113549.1.1.4

Key: Sun RSA public key, 1024 bits
    modulus: 134968166047563266914058280571444028986498087544923991226919517
667593269213420979048109900052353578998293280426361122296881234393722020
```



```
704208851688212064483570055963805034839797994154526862998272017486468599
 962268346037652120279791547218281230795146025359480589335682217749874703
 510467348902637769973696151441
 public exponent: 65537
 Validity: [From: Tue May 22 16:46:46 PDT 2001,
           To: Sun May 22 16:46:46 PDT 2011]
 Issuer: CN=Duke, OU=Java Software, O="Sun Microsystems, Inc.",
 L=Cupertino, ST=CA, C=US
 SerialNumber: [
               3b0afa661
 Algorithm: [MD5withRSA]
 Signature:
0020: 62 17 08 48 14 68 80 CF DD 89 11 EA 92 7F CE DD b..H.h......
0030: B4 FD 12 A8 71 C7 9E D7 C3 D0 E3 BD BB DE 20 92 ....q.........
0040: C2 3B C8 DE CB 25 23 C0 8B B6 92 B9 0B 64 80 63 .;...%#.....d.c
0050: D9 09 25 2D 7A CF 0A 31 B6 E9 CA C1 37 93 BC 0D ..%-z..1....7...
0060: 4E 74 95 4F 58 31 DA AC DF D8 BD 89 BD AF EC C8 Nt.OX1......
1
+++
```

In the case of this particular cipher suite, we must now pass a message called a ClientKeyExchange, which helps establish a shared secret between the two parties.

All of this data is eventually collected and written to the raw device.

```
*** ClientKeyExchange, RSA PreMasterSecret, TLSv1
Random Secret: { 3, 1, 132, 84, 245, 214, 235, 245, 168, 8, 186, 250,
122, 34, 97, 45, 117, 220, 64, 232, 152, 249, 14, 178, 135, 128, 184,
26, 143, 104, 37, 184, 81, 208, 84, 69, 97, 138, 80, 201, 187, 14, 57,
83, 69, 120, 190, 121 }
[write] MD5 and SHA1 hashes: len = 754
0000: 0B 00 02 68 00 02 65 00 02 62 30 82 02 5E 30 82 ...h..e..b0..^0.
0010: 01 C7 02 04 3B 0A FA 66 30 0D 06 09 2A 86 48 86 ....;..f0...*.H.
0020: F7 0D 01 01 04 05 00 30 76 31 0B 30 09 06 03 55 ......0v1.0...U
0030: 04 06 13 02 55 53 31 0B 30 09 06 03 55 04 08 13 ....US1.0...U...
0040: 02 43 41 31 12 30 10 06
                        03 55 04 07 13 09 43 75 .CA1.0...U....Cu
0050: 70 65 72 74 69 6E 6F 31
                         1F 30 1D 06 03 55 04 0A pertinol.0...U..
0060: 13 16 53 75 6E 20 4D 69
                         63 72 6F 73 79 73 74 65 ...Sun Microsyste
0070: 6D 73 2C 20 49 6E 63 2E 31 16 30 14 06 03 55 04 ms, Inc.1.0...U.
0080: 0B 13 0D 4A 61 76 61 20 53 6F 66 74 77 61 72 65 ...Java Software
0090: 31 0D 30 0B 06 03 55 04 03 13 04 44 75 6B 65 30 1.0...U....Duke0
OOAO: 1E 17 OD 30 31 30 35 32 32 33 34 36 34 36 5A ...010522234646Z
00B0: 17 0D 31 31 30 35 32 32 32 32 34 36 34 36 5A 30 ..110522234646Z0
00C0: 76 31 0B 30 09 06 03 55 04 06 13 02 55 53 31 0B v1.0...US1.
OODO: 30 09 06 03 55 04 08 13 02 43 41 31 12 30 10 06 0...U....CA1.0..
0100: 63 72 6F 73 79 73 74 65 6D 73 2C 20 49 6E 63 2E crosystems, Inc.
0120: 53 6F 66 74 77 61 72 65 31 0D 30 0B 06 03 55 04 Software1.0...U.
0130: 03 13 04 44 75 6B 65 30 81 9F 30 0D 06 09 2A 86 ...Duke0..0...*.
0140: 48 86 F7 0D 01 01 01 05 00 03 81 8D 00 30 81 89 H.................
0170: F3 30 E9 A0 C6 07 B7 C8 55 2D B9 5B 57 7A 4C AD .0.....U-.[WzL.
```

```
0180: 1A 30 63 5C 7D 6D 16 BF ED 54 13 49 8A 1B E6 29 .0c\.m...T.I...)
0190: 26 20 85 F9 5E 2B 2F A7
                             12 9C 98 2D 83 F6 EE B1 &..^+/....
01A0: 85 68 DA B5 8E 4C 1D 2D
                             8E 21 97 B0 30 C8 3A 57
                                                     .h...L.-.!..0.:W
01B0: F4 E1 18 9E F6 98 B2 D5
                              3D 8E D5 2B 09 E2 E1 A0
                                                     . . . . . . . . = . . + . . . .
01CO: 49 C1 A6 43 CE EA 57 7F
                              3B 5C 3A C9 BA DB B7 F0 I..C..W.;\:....
01D0: 89 69 BF 91 02 03 01 00
                              01 30 0D 06 09 2A 86 48
                                                     .i.......0...*.H
01E0: 86 F7 0D 01 01 04 05 00
                              03 81 81 00 5F B5 62 E9
                                                      .....b.
01F0: A0 26 1D 8E A2 7E 7C 02
                              08 36 3A 3E C9 C2 45 03
                                                      .&.....6:>..E.
0200: DD F9 BC 06 FC 25 CF 30
                              92 91 B1 4E 62 17 08 48
                                                      ....%.0...Nb..H
0210: 14 68 80 CF DD 89 11 EA
                              92 7F CE DD B4 FD 12 A8
                                                      .h.........
0220: 71 C7 9E D7 C3 D0 E3 BD
                              BB DE 20 92 C2 3B C8 DE
                                                     q....i..
0230: CB 25 23 CO 8B B6 92 B9
                              OB 64 80 63 D9 09 25 2D
                                                     .%#.....d.c..%-
0240: 7A CF 0A 31 B6 E9 CA C1
                              37 93 BC 0D 4E 74 95 4F
                                                     z..1....7...Nt.0
0250: 58 31 DA AC DF D8 BD 89
                              BD AF EC C8 2D 18 A2 BC
                                                     X1........
0260: B2 15 4F B7 28 6F D3 00
                             E1 72 9B 6C 10 00 00 82
                                                     ..O.(o...r.l....
0270: 00 80 4E DD E7 77 F1 91
                              6B 31 4E FA D6 61 D9 69
                                                     ..N..w..klN..a.i
0280: 82 BD 22 40 83 FD 76 E6
                             FF A7 18 95 A0 04 28 0D
                                                     .."@..v.....(.
0290: 0D F7 44 6F 0C 42 4F 17
                              77 A0 99 56 2A 13 77 28 ..Do.BO.w..V*.w(
02A0: 0B 09 48 C1 B9 8C 09 ED
                             9F C6 2E 32 18 DB BD ED ..H.....2....
02B0: AF C3 AB E7 AD 8F DF 9E
                             02C0: EA FC E8 4D C9 DA FC B0
                             B2 C7 D4 83 50 B5 84 B8 ...M......P....
02D0: 44 86 7B 5D 8A C2 F8 04 80 06 E6 84 42 33 B2 EE D..]......B3..
02E0: 05 E6 D3 48 0E 23 E5 1F 63 4C 53 98 B8 8C 45 BA ...H.#..cLS...E.
main, WRITE: TLSv1 Handshake, length = 754
[Raw write]: length = 759
0000: 16 03 01 02 F2 0B 00 02
                             68 00 02 65 00 02 62 30
                                                     ....h..e..b0
0020: 09 2A 86 48 86 F7 0D 01
                             01 04 05 00 30 76 31 0B .*.H......0v1.
0030: 30 09 06 03 55 04 06 13
                              02 55 53 31 0B 30 09 06 0...U....US1.0..
0040: 03 55 04 08 13 02 43 41
                              31 12 30 10 06 03 55 04
                                                     .U....CA1.0...U.
0050: 07 13 09 43 75 70 65 72
                              74 69 6E 6F 31 1F 30 1D
                                                     ...Cupertinol.0.
0060: 06 03 55 04 0A 13 16 53
                              75 6E 20 4D 69 63 72 6F
                                                     ..U....Sun Micro
0070: 73 79 73 74 65 6D 73 2C
                              20 49 6E 63 2E 31 16 30 systems, Inc.1.0
0080: 14 06 03 55 04 0B 13 0D
                              4A 61 76 61 20 53 6F 66
                                                     ...U....Java Sof
0090: 74 77 61 72 65 31 0D 30
                              OB 06 03 55 04 03 13 04 tware1.0...U....
00A0: 44 75 6B 65 30 1E 17 0D
                              30 31 30 35 32 32 32 33 Duke0...01052223
00B0: 34 36 34 36 5A 17 0D 31
                              31 30 35 32 32 32 33 34 4646Z..110522234
00C0: 36 34 36 5A 30 76 31 0B
                              30 09 06 03 55 04 06 13 646Z0v1.0...U...
00D0: 02 55 53 31 0B 30 09 06
                             03 55 04 08 13 02 43 41 .US1.0...U....CA
                             07 13 09 43 75 70 65 72 1.0...U....Cuper
00E0: 31 12 30 10 06 03 55 04
00F0: 74 69 6E 6F 31 1F 30 1D
                             06 03 55 04 0A 13 16 53 tino1.0...U....S
0100: 75 6E 20 4D 69 63 72 6F
                             73 79 73 74 65 6D 73 2C un Microsystems,
0110: 20 49 6E 63 2E 31 16 30
                             14 06 03 55 04 0B 13 0D Inc.1.0...U....
0120: 4A 61 76 61 20 53 6F 66
                             74 77 61 72 65 31 0D 30 Java Softwarel.0
0130: 0B 06 03 55 04 03 13 04
                             44 75 6B 65 30 81 9F 30
                                                     ...U....Duke0..0
0140: 0D 06 09 2A 86 48 86 F7
                              OD 01 01 01 05 00 03 81
                                                     ...*.H.......
0150: 8D 00 30 81 89 02 81 81
                             00 CO 33 77 E7 1F DO CE
                                                     ..0.....3w....
0160: CE BD 43 2F 8D EB C6 D3
                             07 A9 00 F5 75 4D C8 4B
                                                     ..C/....uM.K
0170: 04 52 42 EE 69 F3 30 E9
                             A0 C6 07 B7 C8 55 2D B9
                                                     .RB.i.0.....U-.
0180: 5B 57 7A 4C AD 1A 30 63
                              5C 7D 6D 16 BF ED 54 13 [WzL..0c\.m...T.
0190: 49 8A 1B E6 29 26 20 85
                              F9 5E 2B 2F A7 12 9C 98 I...)& ..^+/....
01A0: 2D 83 F6 EE B1 85 68 DA
                              B5 8E 4C 1D 2D 8E 21 97
                                                     -....h...L.-.!.
01B0: B0 30 C8 3A 57 F4 E1 18
                              9E F6 98 B2 D5 3D 8E D5
                                                     .O.:W.....
01CO: 2B 09 E2 E1 A0 49 C1 A6
                              43 CE EA 57 7F 3B 5C 3A
                                                     +....I..C..W.;\:
01D0: C9 BA DB B7 F0 89 69 BF
                              91 02 03 01 00 01 30 0D
                                                      .....i.....0.
01E0: 06 09 2A 86 48 86 F7 0D
                              01 01 04 05 00 03 81 81
                                                     ..*.H.......
                                                     ._.b..&.....6:
01F0: 00 5F B5 62 E9 A0 26 1D
                              8E A2 7E 7C 02 08 36 3A
0200: 3E C9 C2 45 03 DD F9 BC
                             06 FC 25 CF 30 92 91 B1 >..E.....%.0...
0210: 4E 62 17 08 48 14 68 80
                             CF DD 89 11 EA 92 7F CE Nb..H.h.....
0220: DD B4 FD 12 A8 71 C7 9E D7 C3 D0 E3 BD BB DE 20 ....q.....
```

```
0230: 92 C2 3B C8 DE CB 25 23 C0 8B B6 92 B9 0B 64 80 ..;...\%#.....d.
0240: 63 D9 09 25 2D 7A CF 0A 31 B6 E9 CA C1 37 93 BC c..%-z..1....7..
0250: 0D 4E 74 95 4F 58 31 DA AC DF D8 BD 89 BD AF EC .Nt.OX1......
0260: C8 2D 18 A2 BC B2 15 4F B7 28 6F D3 00 E1 72 9B .-....O.(o...r.
0270: 6C 10 00 00 82 00 80 4E DD E7 77 F1 91 6B 31 4E 1.....N..w..kln
0290: 95 A0 04 28 0D 0D F7 44
                      6F OC 42 4F 17 77 AO 99
                                         ...(...Do.BO.w..
02C0: B4 50 EF 74 98 EA FC E8
                      4D C9 DA FC B0 B2 C7 D4 .P.t...M......
02D0: 83 50 B5 84 B8 44 86 7B
                      5D 8A C2 F8 04 80 06 E6
                                         .P...D..]......
02E0: 84 42 33 B2 EE 05 E6 D3 48 0E 23 E5 1F 63 4C 53 .B3.....H.#..cLS
02F0: 98 B8 8C 45 BA C8 19
```

At this point, we have everything we need to generate the actual secrets.

```
SESSION KEYGEN:
PreMaster Secret:
0000: 03 01 84 54 F5 D6 EB F5 A8 08 BA FA 7A 22 61 2D ...T....z"a-
0010: 75 DC 40 E8 98 F9 0E B2
                         87 80 B8 1A 8F 68 25 B8 u.@.....h%.
0020: 51 D0 54 45 61 8A 50 C9 BB 0E 39 53 45 78 BE 79 Q.TEa.P...9SEx.y
CONNECTION KEYGEN:
Client Nonce:
0010: 05 F6 A7 7B 37 BB 72 E1 FC 1D 1B 6A F5 1C C8 9F ....7.r....j....
Server Nonce:
0000: 40 FC 31 10 79 AB 17 66 FA 8B 3F AA FD 5E 48 23 @.1.y..f..?..^H#
0010: FA 90 31 D8 3C B9 A3 2C 8C F5 E9 81 9B A2 63 6C ..1.<....cl
Master Secret:
Server MAC write Secret:
0000: 1E 4D D1 D3 0A 78 EE B7 4F EC 15 79 B2 59 18 40 .M...x..O..y.Y.@
Client write key:
0000: 10 D0 D6 C2 D9 B7 62 CB 2C 74 BF 5F 85 3C 6F E7 .....b.,t._.<o.
Server write key:
0000: 06 65 DF BD 16 4B A5 7D 8C 66 2A 80 C1 45 B4 F3 .e...K...f*..E..
... no IV for cipher
```

Send a quick confirmation to the server verifying that we know the private key corresponding to the client certificate we just sent.

*** CertificateVerify

```
[write] MD5 and SHA1 hashes: len = 134
0010: AE F9 B3 2C 1F B9 FE 7B 3E 91 50 C5 0F F1 57 4F ...,.....>.P...WO
0020: 55 F1 4B C3 79 16 A8 F1 72 6B 10 CA CC 83 02 FC U.K.y...rk.....
0030: 97 3D 04 29 44 4C 58 74 84 94 19 63 BB 8A 2C 78 .=.)DLXt...c..,x
0050: F2 DA AE 2D 8F FB 50 44 9E E2 1F 7D C9 C5 CB A0 ...-..PD......
0080: 1B 7E 6E D5 8C 62
                                   ..n..b
main, WRITE: TLSv1 Handshake, length = 134
[Raw write]: length = 139
0000: 16 03 01 00 86 0F 00 00 82 00 80 45 41 43 4B 47 ......EACKG
0020: C5 0F F1 57 4F 55 F1 4B C3 79 16 A8 F1 72 6B 10 ...WOU.K.y...rk.
0030: CA CC 83 02 FC 97 3D 04 29 44 4C 58 74 84 94 19 .....=.)DLXt...
```



```
0040: 63 BB 8A 2C 78 43 A0 DD 5E 54 52 AA 97 15 92 1C c..,xC..^TR.....
0050: 39 6B 10 2E BF F2 DA AE 2D 8F FB 50 44 9E E2 1F 9k....-..PD...
0060: 7D C9 C5 CB A0 31 A0 F9 AA 93 2D 1B 07 1B FA E0 ....1...-....
0070: EE 95 E7 88 D7 AD 4A 3A 40 DC FB DF 9E EB 75 04 .....J:@....u.
0080: 14 E2 F2 BB DC 1B 7E 6E D5 8C 62 .....n.b
```

Almost finished! Tell the server we're changing to the newly established cipher suite. All further messages will be encrypted using the parameters we just established. We send an encrypted Finished message to verify everything worked.

Note next that when the message above is actually written to the raw output device (following the 5 bytes of header information), the message is now encrypted.

```
[Raw write]: length = 37

0000: 16 03 01 00 20 15 8C 25 BA 4E 73 F5 27 79 49 B1 ....%.Ns.'yI.

0010: E9 F5 7E C8 48 A7 D3 A6 9B BD 6F 8E A5 8E 2B B7 ....H.................+.

0020: EE DC BD F4 D7 .....
```

We now wait for the server to send the same (Change Cipher Spec/Finshed), so we can know it completed negotiations successfully.

```
[Raw read]: length = 5
0000: 14 03 01 00 01
                                                   . . . . .
[Raw read]: length = 1
0000: 01
main, READ: TLSv1 Change Cipher Spec, length = 1
[Raw read]: length = 5
0000: 16 03 01 00 20
[Raw read]: length = 32
0000: 1F F5 FA C8 79 20 CE 91 AA 68 7F 6C F3 5A DB 7B ....y ...h.l.z..
main, READ: TLSv1 Handshake, length = 32
Padded plaintext after DECRYPTION: len = 32
0000: 14 00 00 0C 07 38 46 5F 62 AD 41 B3 DC 79 30 FD .....8F_b.A..y0.
0010: 34 F2 3B 54 1C D4 68 0E 92 0B 9C 7E ED 47 9F 3B 4.;T..h......G.;
*** Finished
verify_data: { 7, 56, 70, 95, 98, 173, 65, 179, 220, 121, 48, 253 }
[read] MD5 and SHA1 hashes: len = 16
0000: 14 00 00 0C 07 38 46 5F 62 AD 41 B3 DC 79 30 FD .....8F_b.A..y0.
```

Everything completed successfully! Let's cache the established session in case we want to reestablish this session after this connection is dropped.

At this point, a SSL/TLS client should examine the credentials of the peer to make sure that it is communicating with the expected server. A HttpsURLConnection would check the hostname and call HostnameVerifier if there was a problem, but the raw SSLSocket doesn't. This verification should be done by hand, but we're ignoring this for now

So, after all that, we're finally ready to exchange application data. We send a "GET / index.html HTTP1.1" command.

```
%% Cached client session: [Session-1, SSL_RSA_WITH_RC4_128_MD5]

Padded plaintext before ENCRYPTION: len = 42

0000: 47 45 54 20 2F 69 6E 64 65 78 2E 68 74 6D 6C 20 GET /index.html

0010: 48 54 54 50 2F 31 2E 31 0A 0A CA CB D1 10 9D 0E HTTP/1.1......

0020: 13 3C D9 66 6B 9E 36 87 ED 9B .<.fk.6...

main, WRITE: TLSv1 Application Data, length = 42
```

Note again the data over the wire is encrypted (skipping the 5 header bytes).

We get the application data back. First the HTTPS header, then the actual data.

```
[Raw read]: length = 5
0000: 17 03 01 00 50
                                                                                                                ....P
[Raw read]: length = 80
0010: E7 CD A5 D6 2D 1B 10 FD 48 56 9C 1B B5 EC 8F 1E ....-...HV......
0020: DB DA F9 83 62 52 15 38 70 4B C1 85 13 EF CA 17 ....br.8pK......
0030: 89 37 D3 45 C0 88 BB 92 63 F6 9C DE 69 E6 60 3E .7.E....c...i.`>
0040: 1F F7 4D C1 56 61 79 01 49 55 FB 38 6B 16 81 BC ..M.Vay.IU.8k...
main, READ: TLSv1 Application Data, length = 80
Padded plaintext after DECRYPTION: len = 80
0000: 48 54 54 50 2F 31 2E 30 20 32 30 30 20 4F 4B 0D HTTP/1.0 200 OK.
0010: 0A 43 6F 6E 74 65 6E 74 2D 4C 65 6E 67 74 68 3A .Content-Length:
0020: 20 35 38 0D 0A 43 6F 6E 74 65 6E 74 2D 54 79 70
                                                                                                              58..Content-Typ
0030: 65 3A 20 74 65 78 74 2F 68 74 6D 6C 0D 0A 0D 0A e: text/html....
0040: 40 18 A1 FF 1D 5A D4 55 98 DB E3 95 01 A0 91 AF @...Z.U......
HTTP/1.0 200 OK
Content-Length: 58
Content-Type: text/html
[Raw read]: length = 5
0000: 17 03 01 00 4A
                                                                                                                ...J
[Raw read]: length = 74
0010: EB E9 2B A2 D7 82 5C 6E C9 9C 58 84 D7 A8 C6 F8 ..+...\n..X.....
0030: 3E 67 4E 95 30 AA 91 0B E4 4F 5C C7 BF 50 AC 61 >gN.0....0\..P.a
0040: 87 B7 80 75 F0 81 F1 00 63 C9
                                                                                                               ...u...c.
main, READ: TLSv1 Application Data, length = 74
Padded plaintext after DECRYPTION: len = 74
0000: 3C 48 54 4D 4C 3E 0A 3C 48 31 3E 48 65 6C 6C 6F <hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>
<hr/>

0010: 20 57 6F 72 6C 64 3C 2F 48 31 3E 0A 54 68 65 20 World</hd>
.The
0020: 74 65 73 74 20 69 73 20
                                                            63 6F 6D 70 6C 65 74 65 test is complete
0030: 21 0A 3C 2F 48 54 4D 4C
                                                             3E 0A 38 2E 68 72 F1 47 !.</html>.8.hr.G
0040: E8 56 D1 EA A6 FC 3C 30
                                                            6F F3
                                                                                                                . V. . . . < 00.
```



```
<HTML>
<H1>Hello World<H1>
The test is complete!
<HTML>
```

Read from the socket again to see if there is any more data. We get a close_notify message, which means this connection is shutting down properly. We send our own in turn, then close the socket.

```
[Raw read]: length = 5
0000: 15 03 01 00 12
[Raw read]: length = 18
0010: 76 49
main, READ: TLSv1 Alert, length = 18
Padded plaintext after DECRYPTION: len = 18
0000: 01 00 FA 44 D5 57 71 5C CC C7 D9 D0 04 23 10 D8 ...D.Wq\....#..
0010: 21 7B
main, RECV TLSv1 ALERT: warning, close_notify
main, called closeInternal(false)
main, SEND TLSv1 ALERT: warning, description = close_notify
Padded plaintext before ENCRYPTION: len = 18
0000: 01 00 8A 2C A2 36 9C 88 22 50 6E BC 95 3B B2 C4 ...,.6.."Pn..;..
0010: FE F2
main, WRITE: TLSv1 Alert, length = 18
[Raw write]: length = 23
0000: 15 03 01 00 12 19 BE 10 8D FA F1 CA DD AB CC 91 ......
0010: 2E 49 08 71 2B C1 05
                                                    .I.q+..
main, called close()
main, called closeInternal(true)
main, called close()
main, called closeInternal(true)
main, called close()
main, called closeInternal(true)
```

Code Examples

The following code examples are included in this section:

Topics

- Converting an Unsecure Socket to a Secure Socket
- Running the JSSE Sample Code
- Creating a Keystore to Use with JSSE
- Using the Server Name Indication (SNI) Extension

Converting an Unsecure Socket to a Secure Socket

Code examples that illustrate how to use JSSE to convert an unsecure socket connection to a secure socket connection. The code samples are excerpted from the book Java SE 6 Network Security by Marco Pistoia, et. al.

Example 8-26 shows sample code that can be used to set up communication between a client and a server using unsecure sockets. This code is then modified in Example 8-27 to use JSSE to set up secure socket communication.



Example 8-26 Socket Example Without SSL

The following examples demonstrates server-side and client-side code for setting up an unsecure socket connection.

In a Java program that acts as a server and communicates with a client using sockets, the socket communication is set up with code similar to the following:

```
import java.io.*;
import java.net.*;

...
int port = availablePortNumber;

ServerSocket s;

try {
    s = new ServerSocket(port);
    Socket c = s.accept();

    OutputStream out = c.getOutputStream();
    InputStream in = c.getInputStream();

    // Send messages to the client through
    // the OutputStream
    // Receive messages from the client
    // through the InputStream
} catch (IOException e) { }
```

The client code to set up communication with a server using sockets is similar to the following:

```
import java.io.*;
import java.net.*;

. . . .

int port = availablePortNumber;
String host = "hostname";

try {
    s = new Socket(host, port);

    OutputStream out = s.getOutputStream();
    InputStream in = s.getInputStream();

    // Send messages to the server through
    // the OutputStream
    // Receive messages from the server
    // through the InputStream
} catch (IOException e) { }
```

Example 8-27 Socket Example with SSL

The following examples demonstrate server-side and client-side code for setting up a secure socket connection.

In a Java program that acts as a server and communicates with a client using secure sockets, the socket communication is set up with code similar to the following. Differences between this program and the one for communication using unsecure sockets are highlighted in bold.

```
import java.io.*;
import javax.net.ssl.*;
int port = availablePortNumber;
SSLServerSocket s;
try {
    SSLServerSocketFactory sslSrvFact =
        (SSLServerSocketFactory)SSLServerSocketFactory.getDefault();
    s = (SSLServerSocket)sslSrvFact.createServerSocket(port);
    SSLSocket c = (SSLSocket)s.accept();
    OutputStream out = c.getOutputStream();
    InputStream in = c.getInputStream();
    // Send messages to the client through
    // the OutputStream
    // Receive messages from the client
    // through the InputStream
catch (IOException e) {
```

The client code to set up communication with a server using secure sockets is similar to the following, where differences with the unsecure version are highlighted in bold:



```
catch (IOException e) {
}
```

Running the JSSE Sample Code

The JSSE sample programs illustrate how to use JSSE.

- Sample Code Illustrating a Secure Socket Connection Between a Client and a Server
- Sample Code Illustrating HTTPS Connections
- Sample Code Illustrating a Secure RMI Connection
- · Sample Code Illustrating the Use of an SSLEngine

When you use the sample code, be aware that the sample programs are designed to illustrate how to use JSSE. They are not designed to be robust applications.



Setting up secure communications involves complex algorithms. The sample programs provide no feedback during the setup process. When you run the programs, be patient: you may not see any output for a while. If you run the programs with the <code>javax.net.debug</code> system property set to <code>all</code>, you will see more feedback. For an introduction to reading this debug information, see the guide, <code>Debugging SSL/TLS Connections</code>.

Where to Find the Sample Code

JSSE Sample Code in the JDK 8 documentation lists all the sample code files and text files. That page also provides a link to a ZIP file that you can download to obtain all the sample code files.

The following sections describe the samples. See README. txt in JSSE Sample Code in the JDK 8 documentation.

Sample Code Illustrating a Secure Socket Connection Between a Client and a Server

The sample programs in the samples/sockets directory illustrate how to set up a secure socket connection between a client and a server.

When running the sample client programs, you can communicate with an existing server, such as a web server, or you can communicate with the sample server program, <code>classFileServer</code>. You can run the sample client and the sample server programs on different machines connected to the same network, or you can run them both on one machine but from different terminal windows.

All the sample <code>SSLSocketClient*</code> programs in the samples/sockets/client directory (and <code>URLReader*</code> programs described in Sample Code Illustrating HTTPS Connections) can be run with the <code>ClassFileServer</code> sample server program. An example of how to do this is shown in Running <code>SSLSocketClientWithClientAuth</code> with <code>ClassFileServer</code>. You can make similar changes to run <code>URLReader</code>, <code>SSLSocketClient</code>, or <code>SSLSocketClientWithTunneling</code> with <code>ClassFileServer</code>.



If an authentication error occurs during communication between the client and the server (whether using a web server or <code>classFileServer</code>), it is most likely because the necessary keys are not in the truststore (trust key database). See Terms and Definitions. For example, the <code>classFileServer</code> uses a keystore called <code>testkeys</code> containing the private key for <code>localhost</code> as needed during the SSL handshake. The <code>testkeys</code> keystore is included in the same samples/sockets/server directory as the <code>classFileServer</code> source. If the client cannot find a certificate for the corresponding public key of <code>localhost</code> in the truststore it consults, then an authentication error will occur. Be sure to use the <code>samplecacerts</code> truststore (which contains the public key and certificate of the <code>localhost</code>), as described in the next section.

Configuration Requirements

When running the sample programs that create a secure socket connection between a client and a server, you will need to make the appropriate certificates file (truststore) available. For both the client and the server programs, you should use the certificates file samplecacerts from the samples directory. Using this certificates file will allow the client to authenticate the server. The file contains all the common Certificate Authority (CA) certificates shipped with the JDK (in the cacerts file), plus a certificate for localhost needed by the client to authenticate localhost when communicating with the sample server ClassFileServer. The ClassFileServer uses a keystore containing the private key for localhost that corresponds to the public key in samplecacerts.

To make the <code>samplecacerts</code> file available to both the client and the server, you can either copy it to the file <code>java-home/lib/security/jssecacerts</code>, rename it to cacerts, and use it to replace the <code>java-home/lib/security/cacerts</code> file, or add the following option to the command line when running the <code>java</code> command for both the client and the server:

-Djavax.net.ssl.trustStore=path_to_samplecacerts_file

To know more about *java-home*, see Terms and Definitions.

The password for the samplecacerts truststore is changeit. You can substitute your own certificates in the samples by using the keytool utility.

If you use a browser, such as Mozilla Firefox or Microsoft Internet Explorer, to access the sample SSL server provided in the <code>classFileServer</code> example, then a dialog box may pop up with the message that it does not recognize the certificate. This is normal because the certificate used with the sample programs is self-signed and is for testing only. You can accept the certificate for the current session. After testing the SSL server, you should exit the browser, which deletes the test certificate from the browser's namespace.

For client authentication, a separate duke certificate is available in the appropriate directories. The public key and certificate is also stored in the samplecacerts file.

Running SSLSocketClient

The SSLSocketClient.java program in JSSE Sample Code in the JDK 8 documentation demonstrates how to create a client that uses an SSLSocket to send an HTTP request and to get a response from an HTTPS server. The output of this program is the HTML source for https://www.verisign.com/index.html.

You must not be behind a firewall to run this program as provided. If you run it from behind a firewall, you will get an UnknownHostException because JSSE cannot find a path through your firewall to www.verisign.com. To create an equivalent client that can



run from behind a firewall, set up proxy tunneling as illustrated in the sample program SSLSocketClientWithTunneling.

Running SSLSocketClientWithTunneling

The SSLSocketClientWithTunneling.java program in JSSE Sample Code in the JDK 8 documentation illustrates how to do proxy tunneling to access a secure web server from behind a firewall. To run this program, you must set the following Java system properties to the appropriate values:

java -Dhttps.proxyHost=webproxy
-Dhttps.proxyPort=ProxyPortNumber
SSLSocketClientWithTunneling



Proxy specifications with the $\neg D$ options are optional. Replace *webproxy* with the name of your proxy host and *ProxyPortNumber* with the appropriate port number.

The program will return the HTML source file from https://www.verisign.com/index.html.

Running SSLSocketClientWithClientAuth

The SSLSocketClientWithClientAuth.java program in JSSE Sample Code in the JDK 8 documentation shows how to set up a key manager to do client authentication if required by a server. This program also assumes that the client is not outside a firewall. You can modify the program to connect from inside a firewall by following the example in SSLSocketClientWithTunneling.

To run this program, you must specify three parameters: host, port, and requested file path. To mirror the previous examples, you can run this program without client authentication by setting the host to www.verisign.com, the port to 443, and the requested file path to https://www.verisign.com. The output when using these parameters is the HTML for the website https://www.verisign.com.

To run SSLSocketClientWithClientAuth to do client authentication, you must access a server that requests client authentication. You can use the sample program ClassFileServer as this server. This is described in the following sections.

Running ClassFileServer

The program referred to herein as <code>ClassFileServer</code> is made up of two files: <code>ClassFileServer.java</code> and <code>ClassServer.java</code> in <code>JSSE Sample Code</code> in the <code>JDK 8</code> documentation.

To execute them, run classFileServer.class, which requires the following parameters:

- port can be any available unused port number, for example, you can use the number 2001.
- docroot indicates the directory on the server that contains the file you want to retrieve. For example, on Solaris, you can use /home/userid/ (where userid refers to your particular UID), whereas on Microsoft Windows systems, you can use c:\.
- TLS is an optional parameter that indicates that the server is to use SSL or TLS.



• true is an optional parameter that indicates that client authentication is required. This parameter is only consulted if the TLS parameter is set.



The TLS and true parameters are optional. If you omit them, indicating that an ordinary (not TLS) file server should be used, without authentication, then nothing happens. This is because one side (the client) is trying to negotiate with TLS, while the other (the server) is not, so they cannot communicate.

Note:

The server expects GET requests in the form <code>GET /path_to_file</code>.

Running SSLSocketClientWithClientAuth with ClassFileServer

You can use the sample programs SSLSocketClientWithClientAuth.java and ClassFileServer in JSSE Sample Code in the JDK 8 documentation to set up authenticated communication, where the client and server are authenticated to each other. You can run both sample programs on different machines connected to the same network, or you can run them both on one machine but from different terminal windows or command prompt windows. To set up both the client and the server, do the following:

- 1. Run the program ClassFileServer from one machine or terminal window. See Running ClassFileServer.
- 2. Run the program SSLSocketClientWithClientAuth on another machine or terminal window. SSLSocketClientWithClientAuth requires the following parameters:
 - host is the host name of the machine that you are using to run ClassFileServer.
 - port is the same port that you specified for ClassFileServer.
 - requestedfilepath indicates the path to the file that you want to retrieve from
 the server. You must give this parameter as /filepath. Forward slashes are
 required in the file path because it is used as part of a GET statement, which
 requires forward slashes regardless of what type of operating system you are
 running. The statement is formed as follows:

```
"GET " + requestedfilepath + " HTTP/1.0"
```

Note:

You can modify the other SSLClient* applications' GET commands to connect to a local machine running ClassFileServer.

Sample Code Illustrating HTTPS Connections

There are two primary APIs for accessing secure communications through JSSE. One way is through a socket-level API that can be used for arbitrary secure communications, as illustrated by the SSLSocketClient, SSLSocketClientWithTunneling,



and SSLSocketClientWithClientAuth (with and without ClassFileServer) sample programs.

A second, and often simpler, way is through the standard Java URL API. You can communicate securely with an SSL-enabled web server by using the HTTPS URL protocol or scheme using the <code>java.net.URL</code> class.

Support for HTTPS URL schemes is implemented in many of the common browsers, which allows access to secured communications without requiring the socket-level API provided with JSSE.

An example URL is https://www.verisign.com.

The trust and key management for the HTTPS URL implementation is environment-specific. The JSSE implementation provides an HTTPS URL implementation. To use a different HTTPS protocol implementation, set the <code>java.protocol.handler.pkgs</code>. See How to Specify a java.lang.System Property to the package name. See the <code>java.net.URL</code> class documentation for details.

The samples that you can download with JSSE include two sample programs that illustrate how to create an HTTPS connection. Both of these sample programs (URLReader.java and URLReaderWithOptions.java) are in the samples/urls directory.

Running URLReader

The URLReader.java program in JSSE Sample Code in the JDK 8 documentation illustrates using the URL class to access a secure site. The output of this program is the HTML source for https://www.verisign.com/. By default, the HTTPS protocol implementation included with JSSE is used. To use a different implementation, set the system property java.protocol.handler.pkgs value to be the name of the package containing the implementation.

If you are running the sample code behind a firewall, then you must set the https.proxyHost and https.proxyPort system properties. For example, to use the proxy host "webproxy" on port 8080, you can use the following options for the java command:

```
-Dhttps.proxyHost=webproxy
-Dhttps.proxyPort=8080
```

Alternatively, you can set the system properties within the source code with the <code>java.lang.System</code> method <code>setProperty()</code>. For example, instead of using the command-line options, you can include the following lines in your program:

```
System.setProperty("java.protocol.handler.pkgs", "com.ABC.myhttpsprotocol");
System.setProperty("https.proxyHost", "webproxy");
System.setProperty("https.proxyPort", "8080");
```

Running URLReaderWithOptions

The URLReaderWithOptions.java program in JSSE Sample Code in the JDK 8 documentation is essentially the same as the URLReader.java program, except that it allows you to optionally input any or all of the following system properties as arguments to the program when you run it:

- java.protocol.handler.pkgs
- https.proxyHost



- https.proxyPort
- https.cipherSuites

To run URLReaderWithOptions, enter the following command:

java URLReaderWithOptions [-h proxyhost -p proxyport] [-k protocolhandlerpkgs] [-c
ciphersarray]

Note:

Multiple protocol handlers can be included in the protocolhandlerpkgs argument as a list with items separated by vertical bars. Multiple SSL cipher suite names can be included in the ciphersarray argument as a list with items separated by commas. The possible cipher suite names are the same as those returned by the SSLSocket.getSupportedCipherSuites() method. The suite names are taken from the SSL and TLS protocol specifications.

You need a protocolhandlerpkgs argument only if you want to use an HTTPS protocol handler implementation other than the default one provided by Oracle.

If you are running the sample code behind a firewall, then you must include arguments for the proxy host and the proxy port. Additionally, you can include a list of cipher suites to enable.

Here is an example of running URLReaderWithOptions and specifying the proxy host "webproxy" on port 8080:

java URLReaderWithOptions -h webproxy -p 8080

Sample Code Illustrating a Secure RMI Connection

The sample code in the samples/rmi directory illustrates how to create a secure Java Remote Method Invocation (RMI) connection. The sample code is basically a "Hello World" example modified to install and use a custom RMI socket factory.

Sample Code Illustrating the Use of an SSLEngine

SSLEngine gives application developers flexibility when choosing I/O and compute strategies. Rather than tie the SSL/TLS implementation to a specific I/O abstraction (such as single-threaded SSLSockets), SSLEngine removes the I/O and compute constraints from the SSL/TLS implementation.

As mentioned earlier, <code>sslengine</code> is an advanced API, and is not appropriate for casual use. Some introductory sample code is provided here that helps illustrate its use. The first demo removes most of the I/O and threading issues, and focuses on many of the <code>SSlengine</code> methods. The second demo is a more realistic example showing how <code>sslengine</code> might be combined with Java NIO to create a rudimentary <code>HTTP/HTTPS</code> server.

Running SSLEngineSimpleDemo

The SSLEngineSimpleDemo.java program in JSSE Sample Code in the JDK 8 documentation is a very simple application that focuses on the operation of the SSLEngine while simplifying the I/O and threading issues. This application creates two SSLEngine objects that exchange SSL/TLS messages via common ByteBuffer objects. A single loop serially performs all of the engine operations and demonstrates how a



secure connection is established (handshaking), how application data is transferred, and how the engine is closed.

The SSLEngineResult provides a great deal of information about the current state of the SSLEngine. This example does not examine all of the states. It simplifies the I/O and threading issues to the point that this is not a good example for a production environment; nonetheless, it is useful to demonstrate the overall function of the SSLEngine.

Running the NIO-Based Server

To fully exploit the flexibility provided by <code>sslengine</code>, you must first understand complementary APIs, such as I/O and threading models.

An I/O model that large-scale application developers find of use is the NIO <code>SocketChannel</code>. NIO was introduced in part to solve some of the scaling problem inherent in the <code>java.net.Socket API</code>. <code>SocketChannel</code> has many different modes of operation including:

- Blocking
- Nonblocking
- Nonblocking with selectors

Sample code for a basic HTTP server is provided that not only demonstrates many of the new NIO APIs, but also shows how SSLEngine can be employed to create a secure HTTPS server. The server is not production quality, but does show many of these new APIs in action.

Inside the samples directory is a README.txt file that introduces the server, explains how to build and configure the server, and provides a brief overview of the code layout. The files of most interest for SSLEngine users are Channello.java and ChannelloSecure.java.



The server example discussed in this section is included in the JDK. You can find the code bundled in the jdk-home/samples/nio/server directory.

Creating a Keystore to Use with JSSE

The procedure as to how you can use the keytool utility to create a simple PKCS12 keystore suitable for use with JSSE.

First you make a keyEntry (with public and private keys) in the keystore, and then you make a corresponding trustedCertEntry (public keys only) in a truststore. For client authentication, you follow a similar process for the client's certificates.



Storing trust anchors and secret keys in PKCS12 is supported since JDK 8.



Note:

It is beyond the scope of this example to explain each step in detail. See keytool.

User input is shown in bold.

 Create a new keystore and self-signed certificate with corresponding public and private keys.

% keytool -genkeypair -alias duke -keyalg RSA -validity 7 -keystore keystore

```
Enter keystore password: <password>
What is your first and last name?
[Unknown]: Duke
What is the name of your organizational unit?
[Unknown]: Java Software
What is the name of your organization?
[Unknown]: Oracle, Inc.
What is the name of your City or Locality?
[Unknown]: Palo Alto
What is the name of your State or Province?
[Unknown]: CA
What is the two-letter country code for this unit?
[Unknown]: US
Is CN=Duke, OU=Java Software, O="Oracle, Inc.",
L=Palo Alto, ST=CA, C=US correct?
[no]: yes
```

2. Examine the keystore. Notice that the entry type is PrivatekeyEntry, which means that this entry has a private key associated with it).

% keytool -list -v -keystore keystore

```
Enter keystore password: <password>
Keystore type: PKCS12
Keystore provider: SUN
Your keystore contains 1 entry
Alias name: duke
Creation date: Jul 25, 2016
Entry type: PrivateKeyEntry
Certificate chain length: 1
Certificate[1]:
Owner: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto, ST=CA, C=US
Issuer: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto, ST=CA, C=US
Serial number: 210cccfc
Valid from: Mon Jul 25 10:33:27 IST 2016 until: Mon Aug 01 10:33:27 IST 2016
Certificate fingerprints:
     SHA1: 80:E5:8A:47:7E:4F:5A:70:83:97:DD:F4:DA:29:3D:15:6B:2A:45:1F
     SHA256: ED:3C:70:68:4E:86:35:9C:63:CC:B9:59:35:58:94:1F:7E:B8:B0:EE:D2:
4B:9D:80:31:67:8A:D4:B4:7A:B5:12
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: RSA (2048)
Version: 3
```



Extensions:

```
#1: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: 7F C9 95 48 42 8D 68 91 BA 1E E6 5C 2C 6B FF 75 ...HB.h...\,k.u
0010: 5F 19 78 43 _.xC
]
```

3. Export and examine the self-signed certificate.

```
% keytool -export -alias duke -keystore keystore -rfc -file duke.cer
Enter keystore password: <password>
Certificate stored in file <duke.cer>
% cat duke.cer
----BEGIN CERTIFICATE----
MIIDdzCCAl+gAwIBAgIEIQzM/DANBgkqhkiG9w0BAQsFADBsMQswCQYDVQQGEwJV
UzELMAkGA1UECBMCQ0ExEjAQBqNVBAcTCVBhbG8qQWx0bzEVMBMGA1UEChMMT3Jh
Y2x1LCBJbmMuMRYwFAYDVQQLEw1KYXZhIFNvZnR3YXJ1MQ0wCwYDVQQDEwREdWt1
MB4XDTE2MDcyNTA1MDMyN1oXDTE2MDgwMTA1MDMyN1owbDELMAkGA1UEBhMCVVMx
CzajbgnvbagtaknbmriweaydvQQHewlQYWxvIEFsdG8xFTaTbgnvbaotdE9yYWNs
{\tt ZSwgSW5jLjEWMBQGA1UECxMNSmF2YSBTb2Z0d2FyZTENMAsGA1UEAxMERHVrZTCC}
ASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEBAJ7+Yeu6HDZgWwkGlG4iKH9w
vGKrxXVR57FaFyheMevrgjlovVnQVFhfdMvjPkjWmpqLg6rfTqU4bKbtoMWV6+Rn
uQrCw2w9xNC93hX9PxRa20UKrSRDKnUSvi1wjlaxfj0KUKuMwbbY9S8x/naYGeTL
lwbHiiMvkoFkP2kzhVgeqHjIwSz4HRN8vWHCwgIDFWX/ZlS+LbvB4TSZkS0ZcQUV
vJWTocOd8RB90W3bkibWkWq166XYGE1Nq1L4WIhrVJwbav6ual69yJsEpVcshVkx
E1WKzJq7dGb03to4aqbReb6+aoCUwb2vNUudNWasSrxoEFArVFGD/ZkPT0esfqEC
AwEAAaMhMB8wHQYDVR00BBYEFH/JlUhCjWiRuh7mXCxr/3VfGXhDMA0GCSqGSIb3
DOEBCwUAA4IBAOAmcTm2ahsIJLayajsvm8yPzOsHA7kIwWfPPHCoHmNbynG67oHB
fleaNvrgm/raTT3TrqQkq0525qI6Cqaoyy8JA2fAp3i+hmyoGHaIlo14bKazaiPS
RCCqk0J8vwY3CY9nVal1XlHJMEcYV7X1sxKbuAKFoAJ29E/p6ie0JdHtQe31M7X9
FNLYzt8EpJYUtWo13B9Oufz/Guuex9PQ7aC93rbO32MxtnnCGMxQHlaHLLPygc/x
\verb|cffGz5Xe5s+NEm78CY7thgN+drI7icBYmv4navsnr20QaD3AfnJ4WYSQyyUUCPxN||
zuk+B0fbLn7PCCcQspmqfgzIpgbEM9M1/yav
```

Alternatively, you could generate a Certificate Signing Request (CSR) with - certreq and send that to a Certificate Authority (CA) for signing, but that is beyond the scope of this example.

4. Import the certificate into a new truststore.

----END CERTIFICATE----



```
Extensions:
#1: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: 7F C9 95 48 42 8D 68 91 BA 1E E6 5C 2C 6B FF 75 ...HB.h...\,k.u
0010: 5F 19 78 43 __.xC
]
Trust this certificate? [no]: yes
Certificate was added to keystore
```

5. Examine the truststore. Note that the entry type is trustedCertEntry, which means that a private key is not available for this entry. It also means that this file is not suitable as a keystore of the KeyManager.

```
% keytool -list -v -keystore truststore
Enter keystore password: <password>
Keystore type: PKCS12
Keystore provider: SUN
Your keystore contains 1 entry
Alias name: dukecert
Creation date: Jul 25, 2016
Entry type: trustedCertEntry
Owner: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto, ST=CA, C=US
Issuer: CN=Duke, OU=Java Software, O="Oracle, Inc.", L=Palo Alto, ST=CA, C=US
Serial number: 210cccfc
Valid from: Mon Jul 25 10:33:27 IST 2016 until: Mon Aug 01 10:33:27 IST 2016
Certificate fingerprints:
    SHA1: 80:E5:8A:47:7E:4F:5A:70:83:97:DD:F4:DA:29:3D:15:6B:2A:45:1F
    SHA256: ED:3C:70:68:4E:86:35:9C:63:CC:B9:59:35:58:94:1F:7E:B8:B0:EE:D2:
4B:9D:80:31:67:8A:D4:B4:7A:B5:12
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: RSA (2048)
Version: 3
Extensions:
#1: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: 7F C9 95 48 42 8D 68 91 BA 1E E6 5C 2C 6B FF 75 ...HB.h....\,k.u
0010: 5F 19 78 43
                                                      _.xC
]
]
************
************
```

6. Now run your applications with the appropriate keystores. Because this example assumes that the default x509KeyManager and X509TrustManager are used, you select the keystores using the system properties described in Customizing JSSE.

```
% java -Djavax.net.ssl.keyStore=keystore -
Djavax.net.ssl.keyStorePassword=password Server

% java -Djavax.net.ssl.trustStore=truststore -
Djavax.net.ssl.trustStorePassword=trustword Client
```



This example authenticated the server only. For client authentication, provide a similar keystore for the client's keys and an appropriate truststore for the server.

Using the Server Name Indication (SNI) Extension

These examples illustrate how you can use the Server Name Indication (SNI) Extension for client-side and server-side applications, and how it can be applied to a virtual infrastructure.

For all examples in this section, to apply the parameters after you set them, call the setSSLParameters(SSLParameters) method on the corresponding SSLSocket, SSLEngine, or SSLServerSocket Object.

Typical Client-Side Usage Examples

The following is a list of use cases that require understanding of the SNI extension for developing a client application:

Case 1. The client wants to access www.example.com.

Set the host name explicitly:

```
SNIHostName serverName = new SNIHostName("www.example.com");
sslParameters.setServerNames(Collections.singletonList(serverName));
```

The client should always specify the host name explicitly.

Case 2. The client does not want to use SNI because the server does not support
it.

Disable SNI with an empty server name list:

```
sslParameters.setServerNames(Collections.emptyList());
```

Case 3. The client wants to access URL https://www.example.com.

Oracle providers will set the host name in the SNI extension by default, but third-party providers may not support the default server name indication. To keep your application provider-independent, always set the host name explicitly.

Case 4. The client wants to switch a socket from server mode to client mode.

First switch the mode with the following method: sslSocket.setUseClientMode(true). Then reset the server name indication parameters on the socket.



Typical Server-Side Usage Examples

The following is a list of use cases that require understanding of the SNI extension for developing a server application:

Case 1. The server wants to accept all server name indication types.

If you do not have any code dealing with the SNI extension, then the server ignores all server name indication types.

Case 2. The server wants to deny all server name indications of type host_name.

Set an invalid server name pattern for host_name:

```
SNIMatcher matcher = SNIHostName.createSNIMatcher("");
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);
sslParameters.setSNIMatchers(matchers);
```

Another way is to create an ${\tt SNIMatcher}$ subclass with a ${\tt matches}()$ method that always returns ${\tt false}$:

```
class DenialSNIMatcher extends SNIMatcher {
    DenialSNIMatcher() {
        super(StandardConstants.SNI_HOST_NAME);
    }

    @Override
    public boolean matches(SNIServerName serverName) {
        return false;
    }
}

SNIMatcher matcher = new DenialSNIMatcher();
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);
sslParameters.setSNIMatchers(matchers);
```

 Case 3. The server wants to accept connections to any host names in the example.com domain.

Set the recognizable server name for host_name as a pattern that includes all *.example.com addresses:

```
SNIMatcher matcher = SNIHostName.createSNIMatcher("(.*\\.)*example\\.com");
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);
sslParameters.setSNIMatchers(matchers);
```

Case 4. The server wants to switch a socket from client mode to server mode.

First switch the mode with the following method:

 ${\tt sslSocket.setUseClientMode(false)}. \ Then \ reset \ the \ server \ name \ indication \\ parameters \ on \ the \ socket.$

Working with Virtual Infrastructures

This section describes how to use the Server Name Indication (SNI) extension from within a virtual infrastructure. It illustrates how to create a parser for ClientHello messages from a socket, provides examples of virtual server dispatchers using

SSLSocket and SSLEngine, describes what happens when the SNI extension is not available, and demonstrates how to create a failover SSLContext.

Preparing the ClientHello Parser

Applications must implement an API to parse the ClientHello messages from a socket. The following examples illustrate the <code>SSLCapabilities</code> and <code>SSLExplorer</code> classes that can perform these functions.

SSLSocketClient.java encapsulates the SSL/TLS/DTLS security capabilities during handshaking (that is, the list of cipher suites to be accepted in an SSL/TLS/DTLS handshake, the record version, the hello version, and the server name indication). It can be retrieved by exploring the network data of an SSL/TLS/DTLS connection via the SSLExplorer.explore() method.

SSLExplorer.java explores the initial ClientHello message from a TLS client, but it does not initiate handshaking or consume network data. The SSLExplorer.explore() method parses the ClientHello message, and retrieves the security parameters into SSLCapabilities. The method must be called before handshaking occurs on any TLS connections.

Virtual Server Dispatcher Based on SSLSocket

This section describes the procedure for using a virtual server dispatcher based on SSLSocket.

1. Register the server name handler.

At this step, the application may create different SSLContext objects for different server name indications, or link a certain server name indication to a specified virtual machine or distributed system.

For example, if the server name is www.example.org, then the registered server name handler may be for a local virtual hosting web service. The local virtual hosting web service will use the specified SSLContext. If the server name is www.example.com, then the registered server name handler may be for a virtual machine hosting on 10.0.0.36. The handler may map this connection to the virtual machine.

2. Create a ServerSocket and accept the new connection.

```
ServerSocket serverSocket = new ServerSocket(serverPort);
Socket socket = serverSocket.accept();
```

3. Read and buffer bytes from the socket input stream, and then explore the buffered bytes.

```
InputStream ins = socket.getInputStream();
byte[] buffer = new byte[0xFF];
int position = 0;
SSLCapabilities capabilities = null;

// Read the header of TLS record
while (position < SSLExplorer.RECORD_HEADER_SIZE) {
   int count = SSLExplorer.RECORD_HEADER_SIZE - position;
   int n = ins.read(buffer, position, count);
   if (n < 0) {
      throw new Exception("unexpected end of stream!");
   }
   position += n;
}</pre>
```



```
// Get the required size to explore the SSL capabilities
int recordLength = SSLExplorer.getRequiredSize(buffer, 0, position);
if (buffer.length < recordLength) {
    buffer = Arrays.copyOf(buffer, recordLength);
}

while (position < recordLength) {
    int count = recordLength - position;
    int n = ins.read(buffer, position, count);
    if (n < 0) {
        throw new Exception("unexpected end of stream!");
    }
    position += n;
}

// Explore
capabilities = SSLExplorer.explore(buffer, 0, recordLength);
if (capabilities != null) {
    System.out.println("Record version: " + capabilities.getRecordVersion());
    System.out.println("Hello version: " + capabilities.getHelloVersion());
}</pre>
```

4. Get the requested server name from the explored capabilities.

List<SNIServerName> serverNames = capabilities.getServerNames();

5. Look for the registered server name handler for this server name indication.

If the service of the host name is resident in a virtual machine or another distributed system, then the application must forward the connection to the destination. The application will need to read and write the raw internet data, rather then the SSL application from the socket stream.

```
Socket destinationSocket = new Socket(serverName, 443);
// Forward buffered bytes and network data from the current socket to the
destinationSocket.
```

If the service of the host name is resident in the same process, and the host name service can use the SSLSocket directly, then the application will need to set the SSLSocket instance to the server:

```
// Get service context from registered handler
// or create the context
SSLContext serviceContext = ...

SSLSocketFactory serviceSocketFac = serviceContext.getSSLSocketFactory();

// wrap the buffered bytes
ByteArrayInputStream bais = new ByteArrayInputStream(buffer, 0, position);
SSLSocket serviceSocket = (SSLSocket)serviceSocketFac.createSocket(socket, bais, true);

// Now the service can use serviceSocket as usual.
```

Virtual Server Dispatcher Based on SSLEngine

This section describes the procedure for using a virtual server dispatcher based on SSLEngine.

Register the server name handler.

At this step, the application may create different SSLContext objects for different server name indications, or link a certain server name indication to a specified virtual machine or distributed system.

For example, if the server name is www.example.org, then the registered server name handler may be for a local virtual hosting web service. The local virtual hosting web service will use the specified SSLContext. If the server name is www.example.com, then the registered server name handler may be for a virtual machine hosting on 10.0.0.36. The handler may map this connection to the virtual machine.

2. Create a ServerSocket Or ServerSocketChannel and accept the new connection.

```
ServerSocketChannel serverSocketChannel = ServerSocketChannel.open();
serverSocketChannel.bind(...);
...
SocketChannel socketChannel = serverSocketChannel.accept();
```

3. Read and buffer bytes from the socket input stream, and then explore the buffered bytes.

```
ByteBuffer buffer = ByteBuffer.allocate(0xFF);
SSLCapabilities capabilities = null;
while (true) {
    // ensure the capacity
    if (buffer.remaining() == 0) {
       ByteBuffer oldBuffer = buffer;
       buffer = ByteBuffer.allocate(buffer.capacity() + 0xFF);
       buffer.put(oldBuffer);
    }
    int n = sc.read(buffer);
    if (n < 0) {
        throw new Exception("unexpected end of stream!");
    int position = buffer.position();
   buffer.flip();
    capabilities = explorer.explore(buffer);
   buffer.rewind();
    buffer.position(position);
    buffer.limit(buffer.capacity());
    if (capabilities != null) {
        System.out.println("Record version: " +
            capabilities.getRecordVersion());
        System.out.println("Hello version: " +
            capabilities.getHelloVersion());
       break;
    }
}
```

buffer.flip(); // reset the buffer position and limitation

4. Get the requested server name from the explored capabilities.

List<SNIServerName> serverNames = capabilities.getServerNames();

5. Look for the registered server name handler for this server name indication.

If the service of the host name is resident in a virtual machine or another distributed system, then the application must forward the connection to the destination. The application will need to read and write the raw internet data, rather then the SSL application from the socket stream.

```
Socket destinationSocket = new Socket(serverName, 443);
// Forward buffered bytes and network data from the current socket to the
destinationSocket.
```

If the service of the host name is resident in the same process, and the host name service can use the <code>sslengine</code> directly, then the application will simply feed the net data to the <code>sslengine</code> instance:

```
// Get service context from registered handler
// or create the context
SSLContext serviceContext = ...

SSLEngine serviceEngine = serviceContext.createSSLEngine();
// Now the service can use the buffered bytes and other byte buffer as usual.
```

No SNI Extension Available

If there is no server name indication in a ClientHello message, then there is no way to select the proper service according to SNI. For such cases, the application may need to specify a default service, so that the connection can be delegated to it if there is no server name indication.

Failover SSLContext

The <code>SSLExplorer.explore()</code> method does not check the validity of SSL/TLS/DTLS contents. If the record format does not comply with SSL/TLS/DTLS specification, or the <code>explore()</code> method is invoked after handshaking has started, then the method may throw an <code>IOException</code> and be unable to produce network data. In such cases, handle the exception thrown by <code>SSLExplorer.explore()</code> by using a failover <code>SSLContext</code>, which is not used to negotiate an <code>SSL/TLS/DTLS</code> connection, but to close the connection with the proper alert message. The following example illustrates a failover <code>SSLContext</code>. You can find an example of the <code>DenialSNIMatcher</code> class in Case 2 in Typical Server-Side Usage Examples.

```
byte[] buffer = ...
                         // buffered network data
boolean failed = true; // SSLExplorer.explore() throws an exception
SSLContext context = SSLContext.getInstance("TLS");
// the failover SSLContext
context.init(null, null, null);
SSLSocketFactory sslsf = context.getSocketFactory();
ByteArrayInputStream bais = new ByteArrayInputStream(buffer, 0, position);
SSLSocket sslSocket = (SSLSocket)sslsf.createSocket(socket, bais, true);
SNIMatcher matcher = new DenialSNIMatcher();
Collection<SNIMatcher> matchers = new ArrayList<>(1);
matchers.add(matcher);
SSLParameters params = sslSocket.getSSLParameters();
params.setSNIMatchers(matchers);
                                  // no recognizable server name
sslSocket.setSSLParameters(params);
try {
    InputStream sslIS = sslSocket.getInputStream();
    sslIS.read();
} catch (Exception e) {
    System.out.println("Server exception " + e);
} finally {
    sslSocket.close();
```



Standard Names

The JDK Security API requires and uses a set of standard names for algorithms, certificates and keystore types. See Java Security Standard Algorithm Names. Find specific provider information in JDK Providers Documentation.

Provider Pluggability

JSSE is fully pluggable and does not restrict the use of third-party JSSE providers in any way.

JSSE Cipher Suite Parameters

Table 8-14 contains a list of additional JSSE cipher suite names related parameters. See Java Security Standard Algorithm Names.

Table 8-14 JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
SSL_NULL_WITH_NULL_NULL	K_NULL	B_NULL	M_NULL
IANA: TLS_NULL_WITH_NULL_NULL			
SSL_RSA_WITH_NULL_MD5	RSA	B_NULL	MD5
IANA: TLS_RSA_WITH_NULL_MD5			
SSL_RSA_WITH_NULL_SHA	RSA	B_NULL	SHA-1
IANA: TLS_RSA_WITH_NULL_SHA			
SSL_RSA_EXPORT_WITH_RC4_40_MD5	RSA_EXPORT	RC4_40	MD5
IANA:			
TLS_RSA_EXPORT_WITH_RC4_40_MD5			
SSL_RSA_WITH_RC4_128_MD5	RSA	RC4	MD5
IANA: TLS_RSA_WITH_RC4_128_MD5			
SSL_RSA_WITH_RC4_128_SHA	RSA	RC4	SHA-1
IANA: TLS_RSA_WITH_RC4_128_SHA			
SSL_RSA_EXPORT_WITH_RC2_CBC_40_ MD5	RSA_EXPORT	RC2_CBC_40	MD5
IANA: TLS RSA EXPORT WITH RC2 CBC 40			
MD5			
SSL_RSA_WITH_IDEA_CBC_SHA	RSA	IDEA_CBC	SHA-1
IANA: TLS_RSA_WITH_IDEA_CBC_SHA			
SSL_RSA_EXPORT_WITH_DES40_CBC_S HA	RSA_EXPORT	DES40_CBC	SHA-1
IANA:			
TLS_RSA_EXPORT_WITH_DES40_CBC_S HA			



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

		_	
Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
SSL_RSA_WITH_DES_CBC_SHA	RSA	DES_CBC	SHA-1
IANA: TLS_RSA_WITH_DES_CBC_SHA			
SSL_RSA_WITH_3DES_EDE_CBC_SHA	RSA	3DES_EDE_CBC	SHA-1
IANA: TLS_RSA_WITH_3DES_EDE_CBC_SHA			
SSL_DH_DSS_EXPORT_WITH_DES40_C BC_SHA	DH_DSS	DES40_CBC	SHA-1
IANA: TLS_DH_DSS_EXPORT_WITH_DES40_CB C_SHA			
SSL_DH_DSS_WITH_DES_CBC_SHA	DH_DSS	DES_CBC	SHA-1
IANA: TLS_DH_DSS_WITH_DES_CBC_SHA			
SSL_DH_DSS_WITH_3DES_EDE_CBC_S HA	DH_DSS	3DES_EDE_CBC	SHA-1
IANA: TLS_DH_DSS_WITH_3DES_EDE_CBC_S HA			
SSL_DH_RSA_EXPORT_WITH_DES40_C BC_SHA	DH_RSA_EXPORT	DES40_CBC	SHA-1
IANA: TLS_DH_RSA_EXPORT_WITH_DES40_CB C_SHA			
SSL_DH_RSA_WITH_DES_CBC_SHA	DH_RSA	DES_CBC	SHA-1
IANA: TLS_DH_RSA_WITH_DES_CBC_SHA			
SSL_DH_RSA_WITH_3DES_EDE_CBC_S HA	DH_RSA	3DES_EDE_CBC	SHA-1
IANA: TLS_DH_RSA_WITH_3DES_EDE_CBC_S HA			
SSL_DHE_DSS_EXPORT_WITH_DES40_ CBC_SHA	DHE_DSS_EXPORT	DES40_CBC	SHA-1
IANA: TLS_DHE_DSS_EXPORT_WITH_DES40_C BC_SHA			
SSL_DHE_DSS_WITH_DES_CBC_SHA	DHE_DSS	DES_CBC	SHA-1
IANA: TLS_DHE_DSS_WITH_DES_CBC_SHA			
SSL_DHE_DSS_WITH_3DES_EDE_CBC_ SHA	DHE_DSS	3DES_EDE_CBC	SHA-1
IANA: TLS_DHE_DSS_WITH_3DES_EDE_CBC_S HA			



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
SSL_DHE_RSA_EXPORT_WITH_DES40_ CBC_SHA IANA: TLS_DHE_RSA_EXPORT_WITH_DES40_C BC_SHA	DHE_RSA_EXPORT	DES40_CBC	SHA-1
	DHE_RSA	DES_CBC	SHA-1
SSL_DHE_RSA_WITH_3DES_EDE_CBC_ SHA IANA: TLS_DHE_RSA_WITH_3DES_EDE_CBC_S HA	DHE_RSA	3DES_EDE_CBC	SHA-1
SSL_DH_anon_EXPORT_WITH_RC4_40_ MD5 IANA: TLS_DH_anon_EXPORT_WITH_RC4_40_ MD5	DH_anon_EXPORT	RC4_40	MD5
SSL_DH_anon_WITH_RC4_128_MD5 IANA: TLS_DH_anon_WITH_RC4_128_MD5	DH_anon	RC4	MD5
SSL_DH_anon_EXPORT_WITH_DES40_C BC_SHA IANA: TLS_DH_anon_EXPORT_WITH_DES40_C BC_SHA	DH_anon	DES40_CBC	SHA-1
SSL_DH_anon_WITH_DES_CBC_SHA IANA: TLS_DH_anon_WITH_DES_CBC_SHA	DH_anon	DES_CBC	SHA-1
SSL_DH_anon_WITH_3DES_EDE_CBC_S HA IANA: TLS_DH_anon_WITH_3DES_EDE_CBC_S HA	DH_anon	3DES_EDE_CBC	SHA-1
TLS_KRB5_WITH_DES_CBC_SHA	KRB5	DES_CBC	SHA-1
TLS_KRB5_WITH_3DES_EDE_CBC_SHA	KRB5	3DES_EDE_CBC	SHA-1
TLS_KRB5_WITH_RC4_128_SHA	KRB5	RC4	SHA-1
TLS_KRB5_WITH_IDEA_CBC_SHA	KRB5	IDEA_CBC	SHA-1
TLS_KRB5_WITH_DES_CBC_MD5	KRB5	DES_CBC	MD5
TLS_KRB5_WITH_3DES_EDE_CBC_MD5	KRB5	3DES_EDE_CBC	MD5
TLS_KRB5_WITH_RC4_128_MD5	KRB5	RC4	MD5
TLS_KRB5_WITH_IDEA_CBC_MD5	KRB5	IDEA_CBC	MD5



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_KRB5_EXPORT_WITH_DES_CBC_40 _SHA	KRB5_EXPORT	DES_CBC	SHA-1
TLS_KRB5_EXPORT_WITH_RC2_CBC_40 _SHA	KRB5_EXPORT	RC2_CBC_40	SHA-1
TLS_KRB5_EXPORT_WITH_RC4_40_SHA	KRB5_EXPORT	RC4_40	SHA-1
TLS_KRB5_EXPORT_WITH_DES_CBC_40 _MD5	KRB5_EXPORT	DES_CBC	MD5
TLS_KRB5_EXPORT_WITH_RC2_CBC_40 _MD5	KRB5_EXPORT	RC2_CBC_40	MD5
TLS_KRB5_EXPORT_WITH_RC4_40_MD5	KRB5_EXPORT	RC4_40	MD5
TLS_PSK_WITH_NULL_SHA	PSK	B_NULL	SHA-1
TLS_DHE_PSK_WITH_NULL_SHA	DHE_PSK	B_NULL	SHA-1
TLS_RSA_PSK_WITH_NULL_SHA	RSA_PSK	B_NULL	SHA-1
TLS_RSA_WITH_AES_128_CBC_SHA	RSA	AES_128_CBC	SHA-1
TLS_DH_DSS_WITH_AES_128_CBC_SHA	DH_DSS	AES_128_CBC	SHA-1
TLS_DH_RSA_WITH_AES_128_CBC_SHA	DH_RSA	AES_128_CBC	SHA-1
TLS_DHE_DSS_WITH_AES_128_CBC_SH A	DHE_DSS	AES_128_CBC	SHA-1
TLS_DHE_RSA_WITH_AES_128_CBC_SH A	DHE_RSA	AES_128_CBC	SHA-1
TLS_DH_anon_WITH_AES_128_CBC_SHA	DH_anon	AES_128_CBC	SHA-1
TLS_RSA_WITH_AES_256_CBC_SHA	RSA	AES_256_CBC	SHA-1
TLS_DH_DSS_WITH_AES_256_CBC_SHA	DH_DSS	AES_256_CBC	SHA-1
TLS_DH_RSA_WITH_AES_256_CBC_SHA	DH_RSA	AES_256_CBC	SHA-1
TLS_DHE_DSS_WITH_AES_256_CBC_SH A	DHE_DSS	AES_256_CBC	SHA-1
TLS_DHE_RSA_WITH_AES_256_CBC_SH A	DHE_RSA	AES_256_CBC	SHA-1
TLS_DH_anon_WITH_AES_256_CBC_SHA	DH_anon	AES_256_CBC	SHA-1
TLS_RSA_WITH_NULL_SHA256	RSA	B_NULL	SHA-1
TLS_RSA_WITH_AES_128_CBC_SHA256	RSA	AES_128_CBC	SHA-256
TLS_RSA_WITH_AES_256_CBC_SHA256	RSA	AES_256_CBC	SHA-256
TLS_DH_DSS_WITH_AES_128_CBC_SHA 256	DH_DSS	AES_128_CBC	SHA-256
TLS_DH_RSA_WITH_AES_128_CBC_SHA 256	DH_RSA	AES_128_CBC	SHA-256
TLS_DHE_DSS_WITH_AES_128_CBC_SH A256	DHE_DSS	AES_128_CBC	SHA-256



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange	Bulk Cipher	Message
Canada Name (ANA name ii amerene)	Algorithm	Algorithm	Authentication Code Algorithm
TLS_RSA_WITH_CAMELLIA_128_CBC_SH A	RSA	CAMELLIA_128_CB C	SHA-1
TLS_DH_DSS_WITH_CAMELLIA_128_CB C_SHA	DH_DSS	CAMELLIA_128_CB C	SHA-1
TLS_DH_RSA_WITH_CAMELLIA_128_CB C_SHA	DH_RSA	CAMELLIA_128_CB C	SHA-1
TLS_DHE_DSS_WITH_CAMELLIA_128_CB C_SHA	DHE_DSS	CAMELLIA_128_CB C	SHA-1
TLS_DHE_RSA_WITH_CAMELLIA_128_CB C_SHA	DHE_RSA	CAMELLIA_128_CB C	SHA-1
TLS_DH_anon_WITH_CAMELLIA_128_CB C_SHA	DH_anon	CAMELLIA_128_CB C	SHA-1
TLS_DHE_RSA_WITH_AES_128_CBC_SH A256	DHE_RSA	AES_128_CBC	SHA-256
TLS_DH_DSS_WITH_AES_256_CBC_SHA 256	DH_DSS	AES_256_CBC	SHA-256
TLS_DH_RSA_WITH_AES_256_CBC_SHA 256	DH_RSA	AES_256_CBC	SHA-256
TLS_DHE_DSS_WITH_AES_256_CBC_SH A256	DHE_DSS	AES_256_CBC	SHA-256
TLS_DHE_RSA_WITH_AES_256_CBC_SH A256	DHE_RSA	AES_256_CBC	SHA-256
TLS_DH_anon_WITH_AES_128_CBC_SHA 256	DH_anon	AES_128_CBC	SHA-256
TLS_DH_anon_WITH_AES_256_CBC_SHA 256	DH_anon	AES_256_CBC	SHA-256
TLS_RSA_WITH_CAMELLIA_256_CBC_SH A	RSA	CAMELLIA_256_CB C	SHA-1
TLS_DH_DSS_WITH_CAMELLIA_256_CB C_SHA	DH_DSS	CAMELLIA_256_CB C	SHA-1
TLS_DH_RSA_WITH_CAMELLIA_256_CB C_SHA	DH_RSA	CAMELLIA_256_CB C	SHA-1
TLS_DHE_DSS_WITH_CAMELLIA_256_CB C_SHA	DHE_DSS	CAMELLIA_256_CB C	SHA-1
TLS_DHE_RSA_WITH_CAMELLIA_256_CB C_SHA	DHE_RSA	CAMELLIA_256_CB C	SHA-1
TLS_DH_anon_WITH_CAMELLIA_256_CB C_SHA	DH_anon	CAMELLIA_256_CB C	SHA-1
TLS_PSK_WITH_RC4_128_SHA	PSK	RC4	SHA-1
TLS_PSK_WITH_3DES_EDE_CBC_SHA	PSK	3DES_EDE_CBC	SHA-1
TLS_PSK_WITH_AES_128_CBC_SHA	PSK	AES_128_CBC	SHA-1
TLS_PSK_WITH_AES_256_CBC_SHA	PSK	AES_256_CBC	SHA-1



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_DHE_PSK_WITH_RC4_128_SHA	DHE_PSK	RC4	SHA-1
TLS_DHE_PSK_WITH_3DES_EDE_CBC_S HA	DHE_PSK	3DES_EDE_CBC	SHA-1
TLS_DHE_PSK_WITH_AES_128_CBC_SH A	DHE_PSK	AES_128_CBC	SHA-1
TLS_DHE_PSK_WITH_AES_256_CBC_SH A	DHE_PSK	AES_256_CBC	SHA-1
TLS_RSA_PSK_WITH_RC4_128_SHA	RSA_PSK	RC4	SHA-1
TLS_RSA_PSK_WITH_3DES_EDE_CBC_S HA	RSA_PSK	3DES_EDE_CBC	SHA-1
TLS_RSA_PSK_WITH_AES_128_CBC_SH A	RSA_PSK	AES_128_CBC	SHA-1
TLS_RSA_PSK_WITH_AES_256_CBC_SH A	RSA_PSK	AES_256_CBC	SHA-1
TLS_RSA_WITH_SEED_CBC_SHA	RSA	SEED_CBC	SHA-1
TLS_DH_DSS_WITH_SEED_CBC_SHA	DH_DSS	SEED_CBC	SHA-1
TLS_DH_RSA_WITH_SEED_CBC_SHA	DH_RSA	SEED_CBC	SHA-1
TLS_DHE_DSS_WITH_SEED_CBC_SHA	DHE_DSS	SEED_CBC	SHA-1
TLS_DHE_RSA_WITH_SEED_CBC_SHA	DHE_RSA	SEED_CBC	SHA-1
TLS_DH_anon_WITH_SEED_CBC_SHA	DH_anon	SEED_CBC	SHA-1
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA	AES_128_GCM	SHA-256
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA	AES_256_GCM	SHA-384
TLS_DHE_RSA_WITH_AES_128_GCM_SH A256	DHE_RSA	AES_128_GCM	SHA-256
TLS_DHE_RSA_WITH_AES_256_GCM_SH A384	DHE_RSA	AES_256_GCM	SHA-384
TLS_DH_RSA_WITH_AES_128_GCM_SHA 256	DH_RSA	AES_128_GCM	SHA-256
TLS_DH_RSA_WITH_AES_256_GCM_SHA 384	DH_RSA	AES_256_GCM	SHA-384
TLS_DHE_DSS_WITH_AES_128_GCM_SH A256	DHE_DSS	AES_128_GCM	SHA-256
TLS_DHE_DSS_WITH_AES_256_GCM_SH A384	DHE_DSS	AES_256_GCM	SHA-384
TLS_DH_DSS_WITH_AES_128_GCM_SHA 256	DH_DSS	AES_128_GCM	SHA-256
TLS_DH_DSS_WITH_AES_256_GCM_SHA 384	DH_DSS	AES_256_GCM	SHA-384
TLS_DH_anon_WITH_AES_128_GCM_SH A256	DH_anon	AES_128_GCM	SHA-256



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_DH_anon_WITH_AES_256_GCM_SH A384	DH_anon	AES_256_GCM	SHA-384
TLS_PSK_WITH_AES_128_GCM_SHA256	PSK	AES_128_GCM	SHA-256
TLS_PSK_WITH_AES_256_GCM_SHA384	PSK	AES_256_GCM	SHA-384
TLS_DHE_PSK_WITH_AES_128_GCM_SH A256	DHE_PSK	AES_128_GCM	SHA-256
TLS_DHE_PSK_WITH_AES_256_GCM_SH A384	DHE_PSK	AES_256_GCM	SHA-384
TLS_RSA_PSK_WITH_AES_128_GCM_SH A256	RSA_PSK	AES_128_GCM	SHA-256
TLS_RSA_PSK_WITH_AES_256_GCM_SH A384	RSA_PSK	AES_256_GCM	SHA-384
TLS_PSK_WITH_AES_128_CBC_SHA256	PSK	AES_128_CBC	SHA-256
TLS_PSK_WITH_AES_256_CBC_SHA384	PSK	AES_256_CBC	SHA-384
TLS_PSK_WITH_NULL_SHA256	PSK	B_NULL	SHA-256
TLS_PSK_WITH_NULL_SHA384	PSK	B_NULL	SHA-384
TLS_DHE_PSK_WITH_AES_128_CBC_SH A256	DHE_PSK	AES_128_CBC	SHA-256
TLS_DHE_PSK_WITH_AES_256_CBC_SH A384	DHE_PSK	AES_256_CBC	SHA-384
TLS_DHE_PSK_WITH_NULL_SHA256	DHE_PSK	B_NULL	SHA-256
TLS_DHE_PSK_WITH_NULL_SHA384	DHE_PSK	B_NULL	SHA-384
TLS_RSA_PSK_WITH_AES_128_CBC_SH A256	RSA_PSK	AES_128_CBC	SHA-256
TLS_RSA_PSK_WITH_AES_256_CBC_SH A384	RSA_PSK	AES_256_CBC	SHA-384
TLS_RSA_PSK_WITH_NULL_SHA256	RSA_PSK	B_NULL	SHA-256
TLS_RSA_PSK_WITH_NULL_SHA384	RSA_PSK	B_NULL	SHA-384
TLS_RSA_WITH_CAMELLIA_128_CBC_SH A256	RSA	CAMELLIA_128_CB C	SHA-256
TLS_DH_DSS_WITH_CAMELLIA_128_CB C_SHA256	DH_DSS	CAMELLIA_128_CB C	SHA-256
TLS_DH_RSA_WITH_CAMELLIA_128_CB C_SHA256	DH_RSA	CAMELLIA_128_CB C	SHA-256
TLS_DHE_DSS_WITH_CAMELLIA_128_CB C_SHA256	DHE_DSS	CAMELLIA_128_CB C	SHA-256
TLS_DHE_RSA_WITH_CAMELLIA_128_CB C_SHA256	DHE_RSA	CAMELLIA_128_CB C	SHA-256
TLS_DH_anon_WITH_CAMELLIA_128_CB C_SHA256	DH_anon	CAMELLIA_128_CB C	SHA-256



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

		1	1
Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_RSA_WITH_CAMELLIA_256_CBC_SH A256	RSA	CAMELLIA_256_CB C	SHA-256
TLS_DH_DSS_WITH_CAMELLIA_256_CB C_SHA256	DH_DSS	CAMELLIA_256_CB C	SHA-256
TLS_DH_RSA_WITH_CAMELLIA_256_CB C_SHA256	DH_RSA	CAMELLIA_256_CB C	SHA-256
TLS_DHE_DSS_WITH_CAMELLIA_256_CB C_SHA256	DHE_DSS	CAMELLIA_256_CB C	SHA-256
TLS_DHE_RSA_WITH_CAMELLIA_256_CB C_SHA256	DHE_RSA	CAMELLIA_256_CB C	SHA-256
TLS_DH_anon_WITH_CAMELLIA_256_CB C_SHA256	DH_anon	CAMELLIA_256_CB C	SHA-256
TLS_EMPTY_RENEGOTIATION_INFO_SC SV	Not applicable	Not applicable	Not applicable
TLS_FALLBACK_SCSV	Not applicable	Not applicable	Not applicable
TLS_ECDH_ECDSA_WITH_NULL_SHA	ECDH_ECDSA	B_NULL	SHA-1
TLS_ECDH_ECDSA_WITH_RC4_128_SHA	ECDH_ECDSA	RC4	SHA-1
TLS_ECDH_ECDSA_WITH_3DES_EDE_C BC_SHA	ECDH_ECDSA	3DES_EDE_CBC	SHA-1
TLS_ECDH_ECDSA_WITH_AES_128_CBC _SHA	ECDH_ECDSA	AES_128_CBC	SHA-1
TLS_ECDH_ECDSA_WITH_AES_256_CBC _SHA	ECDH_ECDSA	AES_256_CBC	SHA-1
TLS_ECDHE_ECDSA_WITH_NULL_SHA	ECDHE_ECDSA	B_NULL	SHA-1
TLS_ECDHE_ECDSA_WITH_RC4_128_SH A	ECDHE_ECDSA	RC4	SHA-1
TLS_ECDHE_ECDSA_WITH_3DES_EDE_ CBC_SHA	ECDHE_ECDSA	3DES_EDE_CBC	SHA-1
TLS_ECDHE_ECDSA_WITH_AES_128_CB C_SHA	ECDHE_ECDSA	AES_128_CBC	SHA-1
TLS_ECDHE_ECDSA_WITH_AES_256_CB C_SHA	ECDHE_ECDSA	AES_256_CBC	SHA-1
TLS_ECDH_RSA_WITH_NULL_SHA	ECDH_RSA	B_NULL	SHA-1
TLS_ECDH_RSA_WITH_RC4_128_SHA	ECDH_RSA	RC4	SHA-1
TLS_ECDH_RSA_WITH_3DES_EDE_CBC _SHA	ECDH_RSA	3DES_EDE_CBC	SHA-1
TLS_ECDH_RSA_WITH_AES_128_CBC_S HA	ECDH_RSA	AES_128_CBC	SHA-1
TLS_ECDH_RSA_WITH_AES_256_CBC_S HA	ECDH_RSA	AES_256_CBC	SHA-1
TLS_ECDHE_RSA_WITH_NULL_SHA	ECDHE_RSA	B_NULL	SHA-1
			-



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_ECDHE_RSA_WITH_RC4_128_SHA	ECDHE_RSA	RC4	SHA-1
TLS_ECDHE_RSA_WITH_3DES_EDE_CB C_SHA	ECDHE_RSA	3DES_EDE_CBC	SHA-1
TLS_ECDHE_RSA_WITH_AES_128_CBC_ SHA	ECDHE_RSA	AES_128_CBC	SHA-1
TLS_ECDHE_RSA_WITH_AES_256_CBC_ SHA	ECDHE_RSA	AES_256_CBC	SHA-1
TLS_ECDH_anon_WITH_NULL_SHA	ECDH_anon	B_NULL	SHA-1
TLS_ECDH_anon_WITH_RC4_128_SHA	ECDH_anon	RC4	SHA-1
TLS_ECDH_anon_WITH_3DES_EDE_CBC _SHA	ECDH_anon	3DES_EDE_CBC	SHA-1
TLS_ECDH_anon_WITH_AES_128_CBC_S HA	ECDH_anon	AES_128_CBC	SHA-1
TLS_ECDH_anon_WITH_AES_256_CBC_S HA	ECDH_anon	AES_256_CBC	SHA-1
TLS_SRP_SHA_WITH_3DES_EDE_CBC_S HA	SRP_SHA	3DES_EDE_CBC	SHA-1
TLS_SRP_SHA_RSA_WITH_3DES_EDE_C BC_SHA	SRP_SHA	3DES_EDE_CBC	SHA-1
TLS_SRP_SHA_DSS_WITH_3DES_EDE_C BC_SHA	SRP_SHA	3DES_EDE_CBC	SHA-1
TLS_SRP_SHA_WITH_AES_128_CBC_SH A	SRP_SHA	AES_128_CBC	SHA-1
TLS_SRP_SHA_RSA_WITH_AES_128_CB C_SHA	SRP_SHA	AES_128_CBC	SHA-1
TLS_SRP_SHA_DSS_WITH_AES_128_CB C_SHA	SRP_SHA	AES_128_CBC	SHA-1
TLS_SRP_SHA_WITH_AES_256_CBC_SH A	SRP_SHA	AES_256_CBC	SHA-1
TLS_SRP_SHA_RSA_WITH_AES_256_CB C_SHA	SRP_SHA	AES_256_CBC	SHA-1
TLS_SRP_SHA_DSS_WITH_AES_256_CB C_SHA	SRP_SHA	AES_256_CBC	SHA-1
TLS_ECDHE_ECDSA_WITH_AES_128_CB C_SHA256	ECDHE_ECDSA	AES_128_CBC	SHA-256
TLS_ECDHE_ECDSA_WITH_AES_256_CB C_SHA384	ECDHE_ECDSA	AES_256_CBC	SHA-384
TLS_ECDH_ECDSA_WITH_AES_128_CBC _SHA256	ECDH_ECDSA	AES_128_CBC	SHA-256
TLS_ECDH_ECDSA_WITH_AES_256_CBC _SHA384	ECDH_ECDSA	AES_256_CBC	SHA-384



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_ECDHE_RSA_WITH_AES_128_CBC_ SHA256	ECDHE_RSA	AES_128_CBC	SHA-256
TLS_ECDHE_RSA_WITH_AES_256_CBC_ SHA384	ECDHE_RSA	AES_256_CBC	SHA-384
TLS_ECDH_RSA_WITH_AES_128_CBC_S HA256	ECDH_RSA	AES_128_CBC	SHA-256
TLS_ECDH_RSA_WITH_AES_256_CBC_S HA384	ECDH_RSA	AES_256_CBC	SHA-384
TLS_ECDHE_ECDSA_WITH_AES_128_GC M_SHA256	ECDHE_ECDSA	AES_128_GCM	SHA-256
TLS_ECDHE_ECDSA_WITH_AES_256_GC M_SHA384	ECDHE_ECDSA	AES_256_GCM	SHA-384
TLS_ECDH_ECDSA_WITH_AES_128_GC M_SHA256	ECDH_ECDSA	AES_128_GCM	SHA-256
TLS_ECDH_ECDSA_WITH_AES_256_GC M_SHA384	ECDH_ECDSA	AES_256_GCM	SHA-384
TLS_ECDHE_RSA_WITH_AES_128_GCM_ SHA256	ECDHE_RSA	AES_128_GCM	SHA-256
TLS_ECDHE_RSA_WITH_AES_256_GCM_ SHA384	ECDHE_RSA	AES_256_GCM	SHA-384
TLS_ECDH_RSA_WITH_AES_128_GCM_S HA256	ECDH_RSA	AES_128_GCM	SHA-256
TLS_ECDH_RSA_WITH_AES_256_GCM_S HA384	ECDH_RSA	AES_256_GCM	SHA-384
TLS_ECDHE_PSK_WITH_RC4_128_SHA	ECDHE_PSK	RC4	SHA-1
TLS_ECDHE_PSK_WITH_3DES_EDE_CB C_SHA	ECDHE_PSK	3DES_EDE_CBC	SHA-1
TLS_ECDHE_PSK_WITH_AES_128_CBC_ SHA	ECDHE_PSK	AES_128_CBC	SHA-1
TLS_ECDHE_PSK_WITH_AES_256_CBC_ SHA	ECDHE_PSK	AES_256_CBC	SHA-1
TLS_ECDHE_PSK_WITH_AES_128_CBC_ SHA256	ECDHE_PSK	AES_128_CBC	SHA-256
TLS_ECDHE_PSK_WITH_AES_256_CBC_ SHA384	ECDHE_PSK	AES_256_CBC	SHA-384
TLS_ECDHE_PSK_WITH_NULL_SHA	ECDHE_PSK	B_NULL	SHA-1
TLS_ECDHE_PSK_WITH_NULL_SHA256	ECDHE_PSK	B_NULL	SHA-256
TLS_ECDHE_PSK_WITH_NULL_SHA384	ECDHE_PSK	B_NULL	SHA-384
TLS_RSA_WITH_ARIA_128_CBC_SHA256	RSA	ARIA_128_CBC	SHA-256
TLS_RSA_WITH_ARIA_256_CBC_SHA384	RSA	ARIA_256_CBC	SHA-384



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_DH_DSS_WITH_ARIA_128_CBC_SH A256	DH_DSS	ARIA_128_CBC	SHA-256
TLS_DH_DSS_WITH_ARIA_256_CBC_SH A384	DH_DSS	ARIA_256_CBC	SHA-384
TLS_DH_RSA_WITH_ARIA_128_CBC_SH A256	DH_RSA	ARIA_128_CBC	SHA-256
TLS_DH_RSA_WITH_ARIA_256_CBC_SH A384	DH_RSA	ARIA_256_CBC	SHA-384
TLS_DHE_DSS_WITH_ARIA_128_CBC_S HA256	DHE_DSS	ARIA_128_CBC	SHA-256
TLS_DHE_DSS_WITH_ARIA_256_CBC_S HA384	DHE_DSS	ARIA_256_CBC	SHA-384
TLS_DHE_RSA_WITH_ARIA_128_CBC_S HA256	DHE_RSA	ARIA_128_CBC	SHA-256
TLS_DHE_RSA_WITH_ARIA_256_CBC_S HA384	DHE_RSA	ARIA_256_CBC	SHA-384
TLS_DH_anon_WITH_ARIA_128_CBC_SH A256	DH_anon	ARIA_128_CBC	SHA-256
TLS_DH_anon_WITH_ARIA_256_CBC_SH A384	DH_anon	ARIA_256_CBC	SHA-384
TLS_ECDHE_ECDSA_WITH_ARIA_128_C BC_SHA256	ECDHE_ECDSA	ARIA_128_CBC	SHA-256
TLS_ECDHE_ECDSA_WITH_ARIA_256_C BC_SHA384	ECDHE_ECDSA	ARIA_256_CBC	SHA-384
TLS_ECDH_ECDSA_WITH_ARIA_128_CB C_SHA256	ECDH_ECDSA	ARIA_128_CBC	SHA-256
TLS_ECDH_ECDSA_WITH_ARIA_256_CB C_SHA384	ECDH_ECDSA	ARIA_256_CBC	SHA-384
TLS_ECDHE_RSA_WITH_ARIA_128_CBC _SHA256	ECDHE_RSA	ARIA_128_CBC	SHA-256
TLS_ECDHE_RSA_WITH_ARIA_256_CBC _SHA384	ECDHE_RSA	ARIA_256_CBC	SHA-384
TLS_ECDH_RSA_WITH_ARIA_128_CBC_ SHA256	ECDH_RSA	ARIA_128_CBC	SHA-256
TLS_ECDH_RSA_WITH_ARIA_256_CBC_ SHA384	ECDH_RSA	ARIA_256_CBC	SHA-384
TLS_RSA_WITH_ARIA_128_GCM_SHA256	RSA	ARIA_128_GCM	SHA-256
TLS_RSA_WITH_ARIA_256_GCM_SHA384	RSA	ARIA_256_GCM	SHA-384
TLS_DHE_RSA_WITH_ARIA_128_GCM_S HA256	DHE_RSA	ARIA_128_GCM	SHA-256
TLS_DHE_RSA_WITH_ARIA_256_GCM_S HA384	DHE_RSA	ARIA_256_GCM	SHA-384



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_DH_RSA_WITH_ARIA_128_GCM_SH A256	DH_RSA	ARIA_128_GCM	SHA-256
TLS_DH_RSA_WITH_ARIA_256_GCM_SH A384	DH_RSA	ARIA_256_GCM	SHA-384
TLS_DHE_DSS_WITH_ARIA_128_GCM_S HA256	DHE_DSS	ARIA_128_GCM	SHA-256
TLS_DHE_DSS_WITH_ARIA_256_GCM_S HA384	DHE_DSS	ARIA_256_GCM	SHA-384
TLS_DH_DSS_WITH_ARIA_128_GCM_SH A256	DH_DSS	ARIA_128_GCM	SHA-256
TLS_DH_DSS_WITH_ARIA_256_GCM_SH A384	DH_DSS	ARIA_256_GCM	SHA-384
TLS_DH_anon_WITH_ARIA_128_GCM_SH A256	DH_anon	ARIA_128_GCM	SHA-256
TLS_DH_anon_WITH_ARIA_256_GCM_SH A384	DH_anon	ARIA_256_GCM	SHA-384
TLS_ECDHE_ECDSA_WITH_ARIA_128_G CM_SHA256	ECDHE_ECDSA	ARIA_128_GCM	SHA-256
TLS_ECDHE_ECDSA_WITH_ARIA_256_G CM_SHA384	ECDHE_ECDSA	ARIA_256_GCM	SHA-384
TLS_ECDH_ECDSA_WITH_ARIA_128_GC M_SHA256	ECDH_ECDSA	ARIA_128_GCM	SHA-256
TLS_ECDH_ECDSA_WITH_ARIA_256_GC M_SHA384	ECDH_ECDSA	ARIA_256_GCM	SHA-384
TLS_ECDHE_RSA_WITH_ARIA_128_GCM _SHA256	ECDHE_RSA	ARIA_128_GCM	SHA-256
TLS_ECDHE_RSA_WITH_ARIA_256_GCM _SHA384	ECDHE_RSA	ARIA_256_GCM	SHA-384
TLS_ECDH_RSA_WITH_ARIA_128_GCM_ SHA256	ECDH_RSA	ARIA_128_GCM	SHA-256
TLS_ECDH_RSA_WITH_ARIA_256_GCM_ SHA384	ECDH_RSA	ARIA_256_GCM	SHA-384
TLS_PSK_WITH_ARIA_128_CBC_SHA256	PSK	ARIA_128_CBC	SHA-256
TLS_PSK_WITH_ARIA_256_CBC_SHA384	PSK	ARIA_256_CBC	SHA-384
TLS_DHE_PSK_WITH_ARIA_128_CBC_SH A256	DHE_PSK	ARIA_128_CBC	SHA-256
TLS_DHE_PSK_WITH_ARIA_256_CBC_SH A384	DHE_PSK	ARIA_256_CBC	SHA-384
TLS_RSA_PSK_WITH_ARIA_128_CBC_SH A256	RSA_PSK	ARIA_128_CBC	SHA-256
TLS_RSA_PSK_WITH_ARIA_256_CBC_SH A384	RSA_PSK	ARIA_256_CBC	SHA-384



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_PSK_WITH_ARIA_128_GCM_SHA256	PSK	ARIA_128_GCM	SHA-256
TLS_PSK_WITH_ARIA_256_GCM_SHA384	PSK	ARIA_256_GCM	SHA-384
TLS_DHE_PSK_WITH_ARIA_128_GCM_S HA256	DHE_PSK	ARIA_128_GCM	SHA-256
TLS_DHE_PSK_WITH_ARIA_256_GCM_S HA384	DHE_PSK	ARIA_256_GCM	SHA-384
TLS_RSA_PSK_WITH_ARIA_128_GCM_S HA256	RSA_PSK	ARIA_128_GCM	SHA-256
TLS_RSA_PSK_WITH_ARIA_256_GCM_S HA384	RSA_PSK	ARIA_256_GCM	SHA-384
TLS_ECDHE_PSK_WITH_ARIA_128_CBC_ SHA256	ECDHE_PSK	ARIA_128_CBC	SHA-256
TLS_ECDHE_PSK_WITH_ARIA_256_CBC_ SHA384	ECDHE_PSK	ARIA_256_CBC	SHA-384
TLS_ECDHE_ECDSA_WITH_CAMELLIA_1 28_CBC_SHA256	ECDHE_ECDSA	CAMELLIA_128_CB C	SHA-256
TLS_ECDHE_ECDSA_WITH_CAMELLIA_2 56_CBC_SHA384	ECDHE_ECDSA	CAMELLIA_256_CB C	SHA-384
TLS_ECDH_ECDSA_WITH_CAMELLIA_12 8_CBC_SHA256	ECDH_ECDSA	CAMELLIA_128_CB C	SHA-256
TLS_ECDH_ECDSA_WITH_CAMELLIA_25 6_CBC_SHA384	ECDH_ECDSA	CAMELLIA_256_CB C	SHA-384
TLS_ECDHE_RSA_WITH_CAMELLIA_128_ CBC_SHA256	ECDHE_RSA	CAMELLIA_128_CB C	SHA-256
TLS_ECDHE_RSA_WITH_CAMELLIA_256_ CBC_SHA384	ECDHE_RSA	CAMELLIA_256_CB C	SHA-384
TLS_ECDH_RSA_WITH_CAMELLIA_128_ CBC_SHA256	ECDH_RSA	CAMELLIA_128_CB C	SHA-256
TLS_ECDH_RSA_WITH_CAMELLIA_256_ CBC_SHA384	ECDH_RSA	CAMELLIA_256_CB C	SHA-384
TLS_RSA_WITH_CAMELLIA_128_GCM_S HA256	RSA	CAMELLIA_128_GC M	SHA-256
TLS_RSA_WITH_CAMELLIA_256_GCM_S HA384	RSA	CAMELLIA_256_GC M	SHA-384
TLS_DHE_RSA_WITH_CAMELLIA_128_G CM_SHA256	DHE_RSA	CAMELLIA_128_GC M	SHA-256
TLS_DHE_RSA_WITH_CAMELLIA_256_G CM_SHA384	DHE_RSA	CAMELLIA_256_GC M	SHA-384
TLS_DH_RSA_WITH_CAMELLIA_128_GC M_SHA256	DH_RSA	CAMELLIA_128_GC M	SHA-256
TLS_DH_RSA_WITH_CAMELLIA_256_GC M_SHA384	DH_RSA	CAMELLIA_256_GC M	SHA-384



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_DHE_DSS_WITH_CAMELLIA_128_G CM_SHA256	DHE_DSS	CAMELLIA_128_GC M	SHA-256
TLS_DHE_DSS_WITH_CAMELLIA_256_G CM_SHA384	DHE_DSS	CAMELLIA_256_GC M	SHA-384
TLS_DH_DSS_WITH_CAMELLIA_128_GC M_SHA256	DH_DSS	CAMELLIA_128_GC M	SHA-256
TLS_DH_DSS_WITH_CAMELLIA_256_GC M_SHA384	DH_DSS	CAMELLIA_256_GC M	SHA-384
TLS_DH_anon_WITH_CAMELLIA_128_GC M_SHA256	DH_anon	CAMELLIA_128_GC M	SHA-256
TLS_DH_anon_WITH_CAMELLIA_256_GC M_SHA384	DH_anon	CAMELLIA_256_GC M	SHA-384
TLS_ECDHE_ECDSA_WITH_CAMELLIA_1 28_GCM_SHA256	ECDHE_ECDSA	CAMELLIA_128_GC M	SHA-256
TLS_ECDHE_ECDSA_WITH_CAMELLIA_2 56_GCM_SHA384	ECDHE_ECDSA	CAMELLIA_256_GC M	SHA-384
TLS_ECDH_ECDSA_WITH_CAMELLIA_12 8_GCM_SHA256	ECDH_ECDSA	CAMELLIA_128_GC M	SHA-256
TLS_ECDH_ECDSA_WITH_CAMELLIA_25 6_GCM_SHA384	ECDH_ECDSA	CAMELLIA_256_GC M	SHA-384
TLS_ECDHE_RSA_WITH_CAMELLIA_128_ GCM_SHA256	ECDHE_RSA	CAMELLIA_128_GC M	SHA-256
TLS_ECDHE_RSA_WITH_CAMELLIA_256_ GCM_SHA384	ECDHE_RSA	CAMELLIA_256_GC M	SHA-384
TLS_ECDH_RSA_WITH_CAMELLIA_128_ GCM_SHA256	ECDH_RSA	CAMELLIA_128_GC M	SHA-256
TLS_ECDH_RSA_WITH_CAMELLIA_256_ GCM_SHA384	ECDH_RSA	CAMELLIA_256_GC M	SHA-384
TLS_PSK_WITH_CAMELLIA_128_GCM_S HA256	PSK	CAMELLIA_128_GC M	SHA-256
TLS_PSK_WITH_CAMELLIA_256_GCM_S HA384	PSK	CAMELLIA_256_GC M	SHA-384
TLS_DHE_PSK_WITH_CAMELLIA_128_G CM_SHA256	DHE_PSK	CAMELLIA_128_GC M	SHA-256
TLS_DHE_PSK_WITH_CAMELLIA_256_G CM_SHA384	DHE_PSK	CAMELLIA_256_GC M	SHA-384
TLS_RSA_PSK_WITH_CAMELLIA_128_GC M_SHA256	RSA_PSK	CAMELLIA_128_GC M	SHA-256
TLS_RSA_PSK_WITH_CAMELLIA_256_GC M_SHA384	RSA_PSK	CAMELLIA_256_GC M	SHA-384
TLS_PSK_WITH_CAMELLIA_128_CBC_SH A256	PSK	CAMELLIA_128_CB C	SHA-256
TLS_PSK_WITH_CAMELLIA_256_GCM_S HA384 TLS_DHE_PSK_WITH_CAMELLIA_128_G CM_SHA256 TLS_DHE_PSK_WITH_CAMELLIA_256_G CM_SHA384 TLS_RSA_PSK_WITH_CAMELLIA_128_GC M_SHA256 TLS_RSA_PSK_WITH_CAMELLIA_256_GC M_SHA384 TLS_PSK_WITH_CAMELLIA_128_CBC_SH	DHE_PSK DHE_PSK RSA_PSK RSA_PSK	CAMELLIA_256_GC M CAMELLIA_128_GC M CAMELLIA_256_GC M CAMELLIA_128_GC M CAMELLIA_256_GC M CAMELLIA_256_GC M CAMELLIA_256_GC	SHA-256 SHA-384 SHA-256 SHA-384

Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_PSK_WITH_CAMELLIA_256_CBC_SH A384	PSK	CAMELLIA_256_CB C	SHA-384
TLS_DHE_PSK_WITH_CAMELLIA_128_CB C_SHA256	DHE_PSK	CAMELLIA_128_CB C	SHA-256
TLS_DHE_PSK_WITH_CAMELLIA_256_CB C_SHA384	DHE_PSK	CAMELLIA_256_CB C	SHA-384
TLS_RSA_PSK_WITH_CAMELLIA_128_CB C_SHA256	RSA_PSK	CAMELLIA_128_CB C	SHA-256
TLS_RSA_PSK_WITH_CAMELLIA_256_CB C_SHA384	RSA_PSK	CAMELLIA_256_CB C	SHA-384
TLS_ECDHE_PSK_WITH_CAMELLIA_128_ CBC_SHA256	ECDHE_PSK	CAMELLIA_128_CB C	SHA-256
TLS_ECDHE_PSK_WITH_CAMELLIA_256_ CBC_SHA384	ECDHE_PSK	CAMELLIA_256_CB C	SHA-384
TLS_RSA_WITH_AES_128_CCM	RSA	AES_128_CCM	ССМ
TLS_RSA_WITH_AES_256_CCM	RSA	AES_256_CCM	ССМ
TLS_DHE_RSA_WITH_AES_128_CCM	DHE_RSA	AES_128_CCM	ССМ
TLS_DHE_RSA_WITH_AES_256_CCM	DHE_RSA	AES_256_CCM	ССМ
TLS_RSA_WITH_AES_128_CCM_8	RSA	AES_128_CCM	CCM_8
TLS_RSA_WITH_AES_256_CCM_8	RSA	AES_256_CCM	CCM_8
TLS_DHE_RSA_WITH_AES_128_CCM_8	DHE_RSA	AES_128_CCM	CCM_8
TLS_DHE_RSA_WITH_AES_256_CCM_8	DHE_RSA	AES_256_CCM	CCM_8
TLS_PSK_WITH_AES_128_CCM	PSK	AES_128_CCM	ССМ
TLS_PSK_WITH_AES_256_CCM	PSK	AES_256_CCM	ССМ
TLS_DHE_PSK_WITH_AES_128_CCM	DHE_PSK	AES_128_CCM	ССМ
TLS_DHE_PSK_WITH_AES_256_CCM	DHE_PSK	AES_256_CCM	ССМ
TLS_PSK_WITH_AES_128_CCM_8	PSK	AES_128_CCM	CCM_8
TLS_PSK_WITH_AES_256_CCM_8	PSK	AES_256_CCM	CCM_8
TLS_DHE_PSK_WITH_AES_128_CCM_8	DHE_PSK	AES_128_CCM	CCM_8
TLS_DHE_PSK_WITH_AES_256_CCM_8	DHE_PSK	AES_256_CCM	CCM_8
TLS_ECDHE_ECDSA_WITH_AES_128_CC M	ECDHE_ECDSA	AES_128_CCM	ССМ
TLS_ECDHE_ECDSA_WITH_AES_256_CC M	ECDHE_ECDSA	AES_256_CCM	ССМ
TLS_ECDHE_ECDSA_WITH_AES_128_CC M_8	ECDHE_ECDSA	AES_128_CCM	CCM_8
TLS_ECDHE_ECDSA_WITH_AES_256_CC M_8	ECDHE_ECDSA	AES_256_CCM	CCM_8



Table 8-14 (Cont.) JSSE Cipher Suite Parameters

Standard Name (IANA name if different)	Key Exchange Algorithm	Bulk Cipher Algorithm	Message Authentication Code Algorithm
TLS_ECDHE_RSA_WITH_CHACHA20_PO LY1305_SHA256	ECDHE_RSA	AEAD_CHACHA20_ POLY1305	SHA-256
TLS_ECDHE_ECDSA_WITH_CHACHA20_ POLY1305_SHA256	ECDHE_ECDSA	AEAD_CHACHA20_ POLY1305	SHA-256
TLS_DHE_RSA_WITH_CHACHA20_POLY1 305_SHA256	DHE_RSA	AEAD_CHACHA20_ POLY1305	SHA-256
TLS_PSK_WITH_CHACHA20_POLY1305_ SHA256	PSK	AEAD_CHACHA20_ POLY1305	SHA-256
TLS_ECDHE_PSK_WITH_CHACHA20_PO LY1305_SHA256	ECDHE_PSK	AEAD_CHACHA20_ POLY1305	SHA-256
TLS_DHE_PSK_WITH_CHACHA20_POLY1 305_SHA256	DHE_PSK	AEAD_CHACHA20_ POLY1305	SHA-256
TLS_RSA_PSK_WITH_CHACHA20_POLY1 305_SHA256	RSA_PSK	AEAD_CHACHA20_ POLY1305	SHA-256



9

Java PKI Programmers Guide

The Java Certification Path API consists of classes and interfaces for handling certification paths, which are also called certification chains. If a certification path meets certain validation rules, it may be used to securely establish the mapping of a public key to a subject.

Topics

PKI Programmers Guide Overview

Core Classes and Interfaces

Implementing a Service Provider

Appendix A: Standard Names

Appendix B: CertPath Implementation in SUN Provider

Appendix C: OCSP Support

Appendix D: CertPath Implementation in JdkLDAP Provider

Appendix E: Disabling Cryptographic Algorithms

PKI Programmers Guide Overview

The Java Certification Path API defines interfaces and abstract classes for creating, building, and validating certification paths. Implementations may be plugged in using a provider-based interface.

This API is based on the Cryptographic Service Providers architecture, described in the *Java Cryptography Architecture Reference Guide*, and includes algorithm-specific classes for building and validating X.509 certification paths according to the PKIX standards. The PKIX standards were developed by the IETF PKIX working group.

This API was originally specified using the Java Community Process program as Java Specification Request (JSR) 000055. The API was included in the Java SDK, starting with Java SE Development Kit (JDK) 1.4. See JSR 55: Certification Path API.

Who Should Read This Document

This document is intended for two types of experienced developers:

- Those who want to design secure applications that build or validate certification paths.
- 2. Those who want to write a service provider implementation for building or validating certification paths.

This document assumes that you have already read Cryptographic Service Providers.

Introduction to Public Key Certificates

Users of public key applications and systems must be confident that a subject's public key is genuine, i.e., that the associated private key is owned by the subject. Public key certificates are used to establish this trust.

A **public key (or identity) certificate** is a binding of a public key to an identity, which is digitally signed by the private key of another entity, often called a **Certification Authority** (CA). For the remainder of this section, the term CA is used to refer to an entity that signs a certificate.

If the user does not have a trusted copy of the public key of the CA that signed the subject's public key certificate, then another public key certificate vouching for the signing CA is required. This logic can be applied recursively, until a chain of certificates (or a certification path) is discovered from a trust anchor or a most-trusted CA to the target subject (commonly referred to as the end-entity). The most-trusted CA is usually specified by a certificate issued to a CA that the user directly trusts. In general, a certification path is an ordered list of certificates, usually comprised of the end-entity's public key certificate and zero or more additional certificates. A certification path typically has one or more encodings, allowing it to be safely transmitted across networks and to different operating system architectures.

The following figure illustrates a certification path from a most-trusted CA's public key (CA 1) to the target subject (Alice). The certification path establishes trust in Alice's public key through an intermediate CA named CA2.

Figure 9-1 Certification Path from CA's Public Key (CA 1) to the Target Subject



A certification path must be validated before it can be relied on to establish trust in a subject's public key. Validation can consist of various checks on the certificates contained in the certification path, such as verifying the signatures and checking that each certificate has not been revoked. The PKIX standards define an algorithm for validating certification paths consisting of X.509 certificates.

Often a user may not have a certification path from a most-trusted CA to the subject. Providing services to build or discover certification paths is an important feature of public key enabled systems. RFC 2587 defines an LDAP (Lightweight Directory Access Protocol) schema definition that facilitates the discovery of X.509 certification paths using the LDAP directory service protocol.

Building and validating certification paths is an important part of many standard security protocols such as SSL/TLS/DTLS, S/MIME, and IPsec. The Java Certification Path API provides a set of classes and interfaces for developers who need to integrate this functionality into their applications. This API benefits two types of developers: those who need to write service provider implementations for a specific certification path building or validation algorithm; and those who need to access standard



algorithms for creating, building, and validating certification paths in an implementation-independent manner.

X.509 Certificates and Certificate Revocation Lists (CRLs)

A public-key certificate is a digitally signed statement from one entity saying that the public key and some other information of another entity has some specific value.

The following table defines some of the key terms:

Public Keys

These are numbers associated with a particular entity, and are intended to be known to everyone who needs to have trusted interactions with that entity. Public keys are used to verify signatures.

Digitally Signed

If some data is *digitally signed* it has been stored with the "identity" of an entity, and a signature that proves that entity knows about the data. The data is rendered unforgeable by signing with the entitys' private key.

Identity

A known way of addressing an entity. In some systems the identity is the public key, in others it can be anything from a UNIX UID to an Email address to an X.509 Distinguished Name.

Signature

A signature is computed over some data using the private key of an entity (the signer).

Private Keys

These are numbers, each of which is supposed to be known only to the particular entity whose private key it is (that is, it's supposed to be kept secret). Private and public keys exist in pairs in all public key cryptography systems (also referred to as "public key crypto systems"). In a typical public key crypto system, such as DSA, a private key corresponds to exactly one public key. Private keys are used to compute signatures.

Entity

An entity is a person, organization, program, computer, business, bank, or something else you are trusting to some degree.

Basically, public key cryptography requires access to users' public keys. In a large-scale networked environment it is impossible to guarantee that prior relationships between communicating entities have been established or that a trusted repository exists with all used public keys. Certificates were invented as a solution to this public key distribution problem. Now a *Certification Authority* (CA) can act as a *Trusted Third Party*. CAs are entities (e.g., businesses) that are trusted to sign (issue) certificates for other entities. It is assumed that CAs will only create valid and reliable certificates as they are bound by legal agreements. There are many public Certification Authorities, such as VeriSign, Thawte, Entrust, and so on. You can also run your own Certification Authority using products such as the Netscape/Microsoft Certificate Servers or the Entrust CA product for your organization.



What Applications use Certificates?

Probably the most widely visible application of X.509 certificates today is in web browsers (such as Mozilla Firefox and Microsoft Internet Explorer) that support the TLS protocol. TLS (Transport Layer Security) is a security protocol that provides privacy and authentication for your network traffic. These browsers can only use this protocol with web servers that support TLS.

Other technologies that rely on X.509 certificates include:

- Various code-signing schemes, such as signed Java ARchives, and Microsoft Authenticode.
- Various secure E-Mail standards, such as PEM and S/MIME.

How do I Get a Certificate?

There are two basic techniques used to get certificates:

- You can create one yourself (using the right tools, such as keytool).
- You can ask a Certification Authority to issue you one (either directly or using a tool such as **keytool** to generate the request).

The main inputs to the certificate creation process are:

- Matched public and private keys, generated using some special tools (such as keytool), or a browser. Only the public key is ever shown to anyone else. The private key is used to sign data; if someone knows your private key, they can masquerade as you ... perhaps forging legal documents attributed to you!
- You need to provide information about the entity being certified (e.g., you). This
 normally includes information such as your name and organizational address. If
 you ask a CA to issue a certificate for you, you will normally need to provide proof
 to show correctness of the information.

If you are asking a CA to issue you a certificate, you provide your public key and some information about you. You'll use a tool (such as keytool or a browser that supports Certificate Signing Request generation). to digitally sign this information, and send it to the CA. The CA will then generate the certificate and return it.

If you're generating the certificate yourself, you'll take that same information, add a little more (dates during which the certificate is valid, a serial number), and just create the certificate using some tool (such as keytool). Not everyone will accept self-signed certificates; one part of the value provided by a CA is to serve as a neutral and trusted introduction service, based in part on their verification requirements, which are openly published in their Certification Service Practices (CSP).

What's Inside an X.509 Certificate?

The X.509 standard defines what information can go into a certificate, and describes how to write it down (the data format). All X.509 certificates have the following data, in addition to the signature:

Version

This identifies which version of the X.509 standard applies to this certificate, which affects what information can be specified in it. Thus far, three versions are defined.



Serial Number

The entity that created the certificate is responsible for assigning it a serial number to distinguish it from other certificates it issues. This information is used in numerous ways, for example when a certificate is revoked its serial number is placed in a Certificate Revocation List (CRL).

Signature Algorithm Identifier

This identifies the algorithm used by the CA to sign the certificate.

Issuer Name

The X.500 name of the entity that signed the certificate. This is normally a CA. Using this certificate implies trusting the entity that signed this certificate. (Note that in some cases, such as *root or top-level* CA certificates, the issuer signs its own certificate.)

Validity Period

Each certificate is valid only for a limited amount of time. This period is described by a start date and time and an end date and time, and can be as short as a few seconds or almost as long as a century. The validity period chosen depends on a number of factors, such as the strength of the private key used to sign the certificate or the amount one is willing to pay for a certificate. This is the expected period that entities can rely on the public value, if the associated private key has not been compromised.

Subject Name

The name of the entity whose public key the certificate identifies. This name uses the X.500 standard, so it is intended to be unique across the Internet. This is the Distinguished Name (DN) of the entity, for example,

CN=Java Duke, OU=Java Software Division, O=Sun Microsystems Inc, C=US

(These refer to the subject's Common Name, Organizational Unit, Organization, and Country.)

Subject Public Key Information

This is the public key of the entity being named, together with an algorithm identifier which specifies which public key crypto system this key belongs to and any associated key parameters.

X.509 Version 1 has been available since 1988, is widely deployed, and is the most generic.

X.509 Version 2 introduced the concept of subject and issuer unique identifiers to handle the possibility of reuse of subject and/or issuer names over time. Most certificate profile documents strongly recommend that names not be reused, and that certificates should not make use of unique identifiers. Version 2 certificates are not widely used.

X.509 Version 3 is the most recent (1996) and supports the notion of extensions, whereby anyone can define an extension and include it in the certificate. Some common extensions in use today are: KeyUsage (limits the use of the keys to particular purposes such as "signing-only") and AlternativeNames (allows other identities to also be associated with this public key, e.g. DNS names, Email addresses, IP addresses). Extensions can be marked critical to indicate that the extension should be checked and enforced/used. For example, if a certificate has the KeyUsage extension marked critical and set to "keyCertSign" then if this certificate is presented during SSL communication, it should be rejected, as the certificate extension indicates



that the associated private key should only be used for signing certificates and not for SSL use.

All the data in a certificate is encoded using two related standards called ASN.1/DER. Abstract Syntax Notation 1 describes data. The Distinguished Encoding Rules describe a single way to store and transfer that data.

What Java API Can Be Used to Access and Manage Certificates?

The Certificate API, found in the java.security.cert package, includes the following:

- CertificateFactory class defines the functionality of a certificate factory, which is
 used to generate certificate, certificate revocation list (CRL), and certification path
 objects from their encoding.
- Certificate class is an abstract class for managing a variety of certificates. It is an
 abstraction for certificates that have different formats but important common uses.
 For example, different types of certificates, such as X.509 and PGP, share general
 certificate functionality (like encoding and verifying) and some types of information
 like public key.
- CRL class is an abstract class for managing a variety of Certificate Revocation Lists (CRLs).
- X509Certificate class is an abstract class for X.509 Certificates. It provides a standard way to access all the attributes of an X.509 certificate.
- X509Extension interface is an interface for an X.509 extension. The extensions
 defined for X.509 v3 certificates and v2 CRLs (Certificate Revocation Lists)
 provide mechanisms for associating additional attributes with users or public keys,
 such as for managing the certification hierarchy, and for managing CRL
 distribution.
- X509CRL class is an abstract class for an X.509 Certificate Revocation List (CRL).
 A CRL is a time-stamped list identifying revoked certificates. It is signed by a Certification Authority (CA) and made freely available in a public repository.
- X509CRLEntry class is an abstract class for a CRL entry.

What Java Tool Can Generate, Display, Import, and Export X.509 Certificates?

There is a tool named keytool that can be used to create public/private key pairs and X.509 v3 certificates, and to manage keystores. Keys and certificates are used to digitally sign your Java applications and applets (see jarsigner).

A *keystore* is a protected database that holds keys and certificates. Access to a keystore is guarded by a password (defined at the time the keystore is created, by the person who creates the keystore, and changeable only when providing the current password). In addition, each private key in a keystore can be guarded by its own password.

Using **keytool**, it is possible to display, import, and export X.509 v1, v2, and v3 certificates stored as files, and to generate new v3 certificates. For examples, see the "EXAMPLES" section for keytool in the *Java Platform, Standard Edition Tools Reference*.



Core Classes and Interfaces

The core classes of the Java Certification Path API consist of interfaces and classes that support certification path functionality in an algorithm and implementation-independent manner.

The API builds on and extends the existing <code>java.security.cert</code> package for handling certificates. The core classes can be broken up into 4 class categories: Basic, Validation, Building, and Storage:

- Basic Certification Path Classes
 - CertPath, CertificateFactory, and CertPathParameters
- Certification Path Validation Classes
 - CertPathValidator, CertPathValidatorResult, and CertPathChecker
- · Certification Path Building Classes
 - CertPathBuilder, and CertPathBuilderResult
- Certificate/CRL Storage Classes
 - CertStore, CertStoreParameters, CertSelector, and CRLSelector

The Java Certification Path API also includes a set of algorithm-specific classes modeled for use with the PKIX certification path validation algorithm defined in RFC 5280: Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. The PKIX Classes are:

- TrustAnchor
- PKIXParameters
- PKIXCertPathValidatorResult
- PKIXBuilderParameters
- PKIXCertPathBuilderResult
- PKIXCertPathChecker
- PKIXRevocationChecker

The complete reference documentation for the relevant Certification Path API classes can be found in java.security.cert .

Most of the classes and interfaces in the CertPath API are not thread-safe. However, there are some exceptions, which will be noted in this guide and in the API specification. Multiple threads that need to access a single non-thread-safe object concurrently should synchronize amongst themselves and provide the necessary locking. Multiple threads each manipulating separate objects need not synchronize.

Topics

Basic Certification Path Classes

Certification Path Validation Classes

Certification Path Building Classes

Certificate/CRL Storage Classes



PKIX Classes

Basic Certification Path Classes

The basic certification path classes provide fundamental functionality for encoding and representing certification paths. The key basic class in the Java Certification Path API is CertPath, which encapsulates the universal aspects shared by all types of certification paths. An application uses an instance of the CertificateFactory class to create a CertPath object.

Topics

The CertPath Class

The CertificateFactory Class

The CertPathParameters Interface

The CertPath Class

The CertPath class is an abstract class for certification paths. It defines the functionality shared by all certification path objects. Various certification path types can be implemented by subclassing the CertPath class, even though they may have different contents and ordering schemes.

All CertPath objects are serializable, immutable and thread-safe and share the following characteristics:

A type

This corresponds to the type of the certificates in the certification path, for example: X.509. The type of a CertPath is obtained using the method:

```
public String getType()
```

For standard certificate types, see CertificateFactory Types.

A list of certificates

The getCertificates method returns the list of certificates in the certification path:

```
public abstract List<? extends Certificate> getCertificates()
```

This method returns a List of zero or more java.security.cert.Certificate objects. The returned List and the Certificates contained within it are immutable, in order to protect the contents of the CertPath object. The ordering of the certificates returned depends on the type. By convention, the certificates in a CertPath object of type X.509 are ordered starting with the target certificate and ending with a certificate issued by the trust anchor. That is, the issuer of one certificate is the subject of the following one. The certificate representing the TrustAnchor should not be included in the certification path. Unvalidated X.509 CertPaths may not follow this convention. PKIX CertPathValidators will detect any departure from these conventions that cause the certification path to be invalid and throw a CertPathValidatorException.

One or more encodings



Each CertPath object supports one or more encodings. These are external encoded forms for the certification path, used when a standard representation of the path is needed outside the Java Virtual Machine (as when transmitting the path over a network to some other party). Each path can be encoded in a default format, the bytes of which are returned using the method:

```
public abstract byte[] getEncoded()
```

Alternatively, the <code>getEncoded(String)</code> method returns a specific supported encoding by specifying the encoding format as a <code>string</code> (ex: "PKCS7"). For standard encoding formats, see <code>CertPath Encodings</code>.

```
public abstract byte[] getEncoded(String encoding)
```

Also, the <code>getEncodings</code> method returns an iterator over the supported encoding format <code>strings</code> (the default encoding format is returned first):

```
public abstract Iterator<String> getEncodings()
```

All CertPath objects are also Serializable. CertPath objects are resolved into an alternate CertPath.CertPathRep object during serialization. This allows a CertPath object to be serialized into an equivalent representation regardless of its underlying implementation.

CertPath objects are generated from an encoded byte array or list of Certificates using a CertificateFactory. Alternatively, a CertPathBuilder may be used to try to find a CertPath from a most-trusted CA to a particular subject. Once a CertPath object has been created, it may be validated by passing it to the validate method of CertPathValidator. Each of these concepts are explained in more detail in subsequent sections.

The CertificateFactory Class

The CertificateFactory class is an engine class that defines the functionality of a certificate factory. It is used to generate Certificate, CRL, and CertPath objects.

A CertificateFactory should not be confused with a CertPathBuilder. A CertPathBuilder (discussed later) is used to discover or find a certification path when one does not exist. In contrast, a CertificateFactory is used when a certification path has already been discovered and the caller needs to instantiate a CertPath object from its contents, which exist in a different form such as an encoded byte array or an array of CertificateS.

Creating a CertificateFactory Object

See the CertificateFactory section in the *Java Cryptography Architecture Reference Guide* for the details of creating a CertificateFactory Object.

Generating CertPath Objects

A CertificateFactory instance generates CertPath objects from a List of Certificate objects or from an InputStream that contains the encoded form of a CertPath. Just like a CertPath, each CertificateFactory supports a default encoding format for certification paths (ex: PKCS#7). To generate a CertPath object and initialize it with the data read



from an input stream (in the default encoding format), use the ${\tt generateCertPath}$ method:

```
public final CertPath generateCertPath(InputStream inStream)
```

or from a particular encoding format:

To find out what encoding formats are supported, use the <code>getCertPathEncodings</code> method (the default encoding is returned first):

```
public final Iterator<String> getCertPathEncodings()
```

To generate a certification path object from a List of Certificate objects, use the following method:

```
public final CertPath generateCertPath(List<? extends Certificate> certificates)
```

A CertificateFactory always returns CertPath objects that consist of Certificates that are of the same type as the factory. For example, a CertificateFactory of type X.509 returns CertPath objects consisting of certificates that are an instance of java.security.cert.X509Certificate.

The following code sample illustrates generating a certification path from a PKCS#7 encoded certificate reply stored in a file:

```
// open an input stream to the file
FileInputStream fis = new FileInputStream(filename);
// instantiate a CertificateFactory for X.509
CertificateFactory cf = CertificateFactory.getInstance("X.509");
// extract the certification path from
// the PKCS7 SignedData structure
CertPath cp = cf.generateCertPath(fis, "PKCS7");
// print each certificate in the path
List<Certificate> certs = cp.getCertificates();
for (Certificate cert : certs) {
    System.out.println(cert);
}
```

Here's another code sample that fetches a certificate chain from a Keystore and converts it to a CertPath using a CertificateFactory:

Note that there is an existing method in CertificateFactory named generateCertificates that parses a sequence of Certificates. For encodings

consisting of multiple certificates, use <code>generateCertificates</code> when you want to parse a collection of possibly unrelated certificates. Otherwise, use <code>generateCertPath</code> when you want to <code>generate</code> a <code>CertPath</code> and subsequently validate it with a <code>CertPathValidator</code> (discussed later).

The CertPathParameters Interface

The CertPathParameters interface is a transparent representation of the set of parameters used with a particular certification path builder or validation algorithm.

Its main purpose is to group (and provide type safety for) all certification path parameter specifications. The <code>CertPathParameters</code> interface extends the <code>Cloneable</code> interface and defines a <code>clone()</code> method that does not throw an exception. All concrete implementations of this interface should implement and override the <code>Object.clone()</code> method, if necessary. This allows applications to clone any <code>CertPathParameters</code> object.

Objects implementing the <code>CertPathParameters</code> interface are passed as arguments to methods of the <code>CertPathValidator</code> and <code>CertPathBuilder</code> classes. Typically, a concrete implementation of the <code>CertPathParameters</code> interface will hold a set of input parameters specific to a particular certification path build or validation algorithm. For example, the <code>PKIXParameters</code> class is an implementation of the <code>CertPathParameters</code> interface that holds a set of input parameters for the PKIX certification path validation algorithm. One such parameter is the set of most-trusted CAs that the caller trusts for anchoring the validation process. This parameter among others is discussed in more detail in the section discussing the <code>PKIXParameters</code> class.

Certification Path Validation Classes

The Java Certification Path API includes classes and interfaces for validating certification paths. An application uses an instance of the CertPathValidator class to validate a CertPath object. If successful, the result of the validation algorithm is returned in an object implementing the CertPathValidatorResult interface.

Topics

The CertPathValidator Class

The CertPathValidatorResult Interface

The CertPathValidator Class

The CertPathValidator class is an engine class used to validate a certification path.

Creating a CertPathValidator Object

As with all engine classes, the way to get a <code>CertPathValidator</code> object for a particular validation algorithm is to call one of the <code>getInstance</code> static factory methods on the <code>CertPathValidator</code> class:



The algorithm parameter is the name of a certification path validation algorithm (for example, "PKIX"). Standard CertPathValidator algorithm names are listed in the Java Security Standard Algorithm Names Specification.

Validating a Certification Path

Once a <code>CertPathValidator</code> object is created, paths can be validated by calling the <code>validate</code> method, passing it the certification path to be validated and a set of algorithm-specific parameters:

If the validation algorithm is successful, the result is returned in an object implementing the CertPathValidatorResult interface. Otherwise, a CertPathValidatorException is thrown. The CertPathValidatorException contains methods that return the CertPath, and if relevant, the index of the certificate that caused the algorithm to fail and the root exception or cause of the failure.

Note that the CertPath and CertPathParameters passed to the validate method must be of a type that is supported by the validation algorithm. Otherwise, an InvalidAlgorithmParameterException is thrown. For example, a CertPathValidator instance that implements the PKIX algorithm validates CertPath objects of type X.509 and CertPathParameters that are an instance of PKIXParameters.

The CertPathValidatorResult Interface

The CertPathValidatorResult interface is a transparent representation of the successful result or output of a certification path validation algorithm.

The main purpose of this interface is to group and provide type safety for all validation results. Similar to the <code>CertPathParameters</code> interface, <code>CertPathValidatorResult</code> extends <code>Cloneable</code> and defines a <code>clone()</code> method that does not throw an exception. This allows applications to clone any <code>CertPathValidatorResult</code> object.

Objects implementing the CertPathValidatorResult interface are returned by the validate method of CertPathValidatorResult interface when successful. If not successful, a CertPathValidatorException is thrown with a description of the failure. Typically, a concrete implementation of the CertPathValidatorResult interface will hold a set of output parameters specific to a particular certification path validation algorithm. For example, the PKIXCertPathValidatorResult class is an implementation of the CertPathValidatorResult interface, which contains methods to get the output parameters of the PKIX certification path validation algorithm. One such parameter is the valid policy tree. This parameter among others is discussed in more detail in the section discussing the PKIXCertPathValidatorResult class.

The following code sample shows how to create a <code>CertPathValidator</code> and use it to validate a certification path. The sample assumes that the <code>CertPath</code> and <code>CertPathParameters</code> objects which are passed to the <code>validate</code> method have been previously created; a more complete example will be illustrated in the section describing the PKIX classes.

```
// create CertPathValidator that implements the "PKIX" algorithm
CertPathValidator cpv = null;
```



```
try {
   cpv = CertPathValidator.getInstance("PKIX");
} catch (NoSuchAlgorithmException nsae) {
   System.err.println(nsae);
   System.exit(1);
// validate certification path ("cp") with specified parameters ("params")
try {
    CertPathValidatorResult cpvResult = cpv.validate(cp, params);
} catch (InvalidAlgorithmParameterException iape) {
    System.err.println("validation failed: " + iape);
    System.exit(1);
} catch (CertPathValidatorException cpve) {
    System.err.println("validation failed: " + cpve);
    System.err.println("index of certificate that caused exception: "
            + cpve.getIndex());
    System.exit(1);
```

Certification Path Building Classes

The Java Certification Path API includes classes for building (or discovering) certification paths. An application uses an instance of the <code>CertPathBuilder</code> class to build a <code>CertPath</code> object. If successful, the result of the build is returned in an object implementing the <code>CertPathBuilderResult</code> interface.

Topics

The CertPathBuilder Class

The CertPathBuilderResult Interface

The CertPathBuilder Class

The CertPathBuilder class is an engine class used to build a certification path.

Creating a CertPathBuilder Object

As with all engine classes, the way to get a <code>CertPathBuilder</code> object for a particular build algorithm is to call one of the <code>getInstance</code> static factory method on the <code>CertPathBuilder</code> class:

The algorithm parameter is the name of a certification path builder algorithm (for example, "PKIX"). Standard CertPathBuilder algorithm names are listed in Java Security Standard Algorithm Names Specification.

Building a Certification Path

Once a CertPathBuilder object is created, paths can be constructed by calling the build method, passing it an algorithm-specific parameter specification:



If the build algorithm is successful, the result is returned in an object implementing the CertPathBuilderResult interface. Otherwise, a CertPathBuilderException is thrown containing information about the failure; for example, the underlying exception (if any) and an error message.

Note that the CertPathParameters passed to the build method must be of a type that is supported by the build algorithm. Otherwise, an InvalidAlgorithmParameterException is thrown.

The CertPathBuilderResult Interface

The CertPathBuilderResult interface is a transparent representation of the result or output of a certification path builder algorithm.

This interface contains a method to return the certification path that has been successfully built:

```
public CertPath getCertPath()
```

The purpose of the CertPathBuilderResult interface is to group (and provide type safety for) all build results. Like the CertPathValidatorResult interface, CertPathBuilderResult extends Cloneable and defines a clone() method that does not throw an exception. This allows applications to clone any CertPathBuilderResult object.

Objects implementing the CertPathBuilderResult interface are returned by the build method of CertPathBuilder.

The following code sample shows how to create a <code>CertPathBuilder</code> and use it to build a certification path. The sample assumes that the <code>CertPathParameters</code> object which is passed to the <code>build</code> method has been previously created; a more complete example will be illustrated in the section describing the PKIX classes.

```
// create CertPathBuilder that implements the "PKIX" algorithm
CertPathBuilder cpb = null;
try {
   cpb = CertPathBuilder.getInstance("PKIX");
} catch (NoSuchAlgorithmException nsae) {
   System.err.println(nsae);
    System.exit(1);
// build certification path using specified parameters ("params")
try {
    CertPathBuilderResult cpbResult = cpb.build(params);
    CertPath cp = cpbResult.getCertPath();
    System.out.println("build passed, path contents: " + cp);
} catch (InvalidAlgorithmParameterException iape) {
    System.err.println("build failed: " + iape);
    System.exit(1);
} catch (CertPathBuilderException cpbe) {
    System.err.println("build failed: " + cpbe);
    System.exit(1);
```



Certificate/CRL Storage Classes

The Java Certification Path API includes the Certstore class for retrieving certificates and CRLs from a repository.

This class enables a caller to specify the repository a CertPathValidator or CertPathBuilder implementation should use to find certificates and CRLs. See the addCertStores method of the PKIXParameters class.

A CertPathValidator implementation may use the CertStore object that the caller specifies as a callback mechanism to fetch CRLs for performing revocation checks. Similarly, a CertPathBuilder may use the CertStore as a callback mechanism to fetch certificates and, if performing revocation checks, CRLs.

Topics

The CertStore Class

The CertStoreParameters Interface

The CertSelector and CRLSelector Interfaces

The CertStore Class

The CertStore class is an engine class used to provide the functionality of a certificate and certificate revocation list (CRL) repository.

This class can be used by CertPathBuilder and CertPathValidator implementations to find certificates and CRLs, or as a general purpose certificate and CRL retrieval mechanism.

Unlike the <code>java.security.KeyStore</code> class, which provides access to a cache of private keys and trusted certificates, a <code>CertStore</code> is designed to provide access to a potentially vast repository of untrusted certificates and CRLs. For example, an LDAP implementation of <code>CertStore</code> provides access to certificates and CRLs stored in one or more directories using the LDAP protocol.

All public methods of <code>certstore</code> objects are thread-safe. That is, multiple threads may concurrently invoke these methods on a single <code>certstore</code> object (or more than one) with no ill effects. This allows a <code>certPathBuilder</code> to search for a CRL while simultaneously searching for further certificates, for instance.

Creating a CertStore Object

As with all engine classes, the way to get a CertStore object for a particular repository type is to call one of the getInstance static factory methods on the CertStore class:



The type parameter is the name of a certificate repository type (for example, "LDAP"). Standard CertStore types are listed in Java Security Standard Algorithm Names Specification.

The initialization parameters (params) are specific to the repository type. For example, the initialization parameters for a server-based repository may include the hostname and the port of the server. An InvalidAlgorithmParameterException is thrown if the parameters are invalid for this CertStore type. The getCertStoreParameters method returns the CertStoreParameters that were used to initialize a CertStore:

public final CertStoreParameters getCertStoreParameters()

Retrieving Certificates

After you have created a <code>certStore</code> object, you can retrieve certificates from the repository using the <code>getCertificates</code> method. This method takes a <code>CertSelector</code> (discussed in more detail later) object as an argument, which specifies a set of selection criteria for determining which certificates should be returned:

 $\verb"public final Collection"<? extends Certificate> \verb"getCertificates"(CertSelector selector)"$

throws CertStoreException

This method returns a <code>Collection</code> of <code>java.security.cert.Certificate</code> objects that satisfy the selection criteria. An empty <code>Collection</code> is returned if there are no matches. A <code>CertStoreException</code> is usually thrown if an unexpected error condition is encountered, such as a communications failure with a remote repository.

For some Certstore implementations, it may not be feasible to search the entire repository for certificates or CRLs that match the specified selection criteria. In these instances, the Certstore implementation may use information that is specified in the selectors to locate certificates and CRLs. For instance, an LDAP Certstore may not search all entries in the directory. Instead, it may just search entries that are likely to contain the certificates it is looking for. If the Certselector provided does not provide enough information for the LDAP Certstore to determine which entries it should look in, the LDAP Certstore may throw a CertstoreException.

Retrieving CRLs

You can also retrieve CRLs from the repository using the <code>getcrls</code> method. This method takes a <code>CRLSelector</code> (discussed in more detail later) object as an argument, which specifies a set of selection criteria for determining which CRLs should be returned:

This method returns a Collection of java.security.cert.CRL objects that satisfy the selection criteria. An empty Collection is returned if there are no matches.

The CertStoreParameters Interface

The CertStoreParameters interface is a transparent representation of the set of parameters used with a particular CertStore.



The main purpose of this interface is to group and provide type safety for all certificate storage parameter specifications. The <code>CertStoreParameters</code> interface extends the <code>Cloneable</code> interface and defines a <code>clone</code> method that does not throw an exception. Implementations of this interface should implement and override the <code>Object.clone()</code> method, if necessary. This allows applications to clone any <code>CertStoreParameters</code> object.

Objects implementing the <code>CertStoreParameters</code> interface are passed as arguments to the <code>getInstance</code> method of the <code>CertStore</code> class. Two classes implementing the <code>CertStoreParameters</code> interface are defined in this API: the <code>LDAPCertStoreParameters</code> class and the <code>CollectionCertStoreParameters</code> class.

The LDAPCertStoreParameters Class

The LDAPCertStoreParameters class is an implementation of the CertStoreParameters interface and holds a set of minimal initialization parameters (host and port number of the directory server) for retrieving certificates and CRLs from a CertStore of type LDAP.

See LDAPCertStoreParameters.

The CollectionCertStoreParameters Class

The <code>CollectionCertStoreParameters</code> class is an implementation of the <code>CertStoreParameters</code> interface and holds a set of initialization parameters for retrieving certificates and <code>CRLs</code> from a <code>CertStore</code> of type <code>Collection</code>.

See CollectionCertStoreParameters.

The CertSelector and CRLSelector Interfaces

The CertSelector and CRLSelector interfaces are a specification of the set of criteria for selecting certificates and CRLs from a collection or large group of certificates and CRLs.

The interfaces group and provide type safety for all selector specifications. Each selector interface extends <code>cloneable</code> and defines a <code>clone()</code> method that does not throw an exception. This allows applications to clone any <code>CertSelector</code> or <code>CRLSelector</code> object.

The CertSelector and CRLSelector interfaces each define a method named match. The match method takes a Certificate or CRL object as an argument and returns true if the object satisfies the selection criteria. Otherwise, it returns false. The match method for the CertSelector interface is defined as follows:

public boolean match(Certificate cert)

and for the CRLSelector interface:

public boolean match(CRL crl)

Typically, objects implementing these interfaces are passed as parameters to the getCertificates and getCRLs methods of the CertStore class. These methods return a Collection of Certificates or CRLs from the CertStore repository that match the specified selection criteria. CertSelectors may also be used to specify the validation constraints on a target or end-entity certificate in a certification path (see for example, the PKIXParameters.setTargetCertConstraints method.)



The X509CertSelector Class

The X509CertSelector class is an implementation of the CertSelector interface that defines a set of criteria for selecting X.509 certificates.

An x509certificate object must match *all* of the specified criteria to be selected by the match method. The selection criteria are designed to be used by a CertPathBuilder implementation to discover potential certificates as it builds an X.509 certification path.

For example, the setSubject method of X509CertSelector allows a PKIX CertPathBuilder to filter out X509Certificates that do not match the issuer name of the preceding X509Certificate in a partially completed chain. By setting this and other criteria in an X509CertSelector object, a CertPathBuilder is able to discard irrelevant certificates and more easily find an X.509 certification path that meets the requirements specified in the CertPathParameters object.

See RFC 5280 for definitions of the X.509 certificate extensions mentioned in this section.

Creating an X509CertSelector Object

An X509CertSelector object is created by calling the default constructor:

```
public X509CertSelector()
```

No criteria are initially set (any X509Certificate will match).

Setting Selection Criteria

The selection criteria allow a caller to match on different components of an X.509 certificate. A few of the methods for setting selection criteria are described here. See X509CertSelector.

The setIssuer methods set the issuer criterion:

```
public void setIssuer(X500Principal issuer)
public void setIssuer(String issuerDN)
public void setIssuer(byte[] issuerDN)
```

The specified distinguished name (in x500Principal, RFC 2253 String or ASN.1 DER encoded form) must match the issuer distinguished name in the certificate. If null, any issuer distinguished name will do. Note that use of an x500Principal to represent a distinguished name is preferred because it is more efficient and suitably typed.

Similarly, the setSubject methods set the subject criterion:

```
public void setSubject(X500Principal subject)
public void setSubject(String subjectDN)
public void setSubject(byte[] subjectDN)
```

The specified distinguished name (in x500Principal, RFC 2253 String or ASN.1 DER encoded form) must match the subject distinguished name in the certificate. If null, any subject distinguished name will do.

The setSerialNumber method sets the serialNumber criterion:



```
public void setSerialNumber(BigInteger serial)
```

The specified serial number must match the certificate serial number in the certificate. If null, any certificate serial number will do.

The setAuthorityKeyIdentifier method sets the authorityKeyIdentifier criterion:

```
public void setAuthorityKeyIdentifier(byte[] authorityKeyID)
```

The certificate must contain an Authority Key Identifier extension matching the specified value. If null, no check will be done on the authorityKeyIdentifier criterion.

The setCertificateValid method sets the certificateValid criterion:

```
public void setCertificateValid(Date certValid)
```

The specified date must fall within the certificate validity period for the certificate. If null, any date is valid.

The setKeyUsage method sets the keyUsage criterion:

```
public void setKeyUsage(boolean[] keyUsage)
```

The certificate's Key Usage Extension must allow the specified key usage values (those which are set to true). If null, no keyUsage check will be done.

Getting Selection Criteria

The current values for each of the selection criteria can be retrieved using an appropriate get method. See x509CertSelector.

Here is an example of retrieving X.509 certificates from an LDAP <code>CertStore</code> with the X509CertSelector class.

First, we create the LDAPCertStoreParameters object that we will use to initialize the CertStore object with the hostname and port of the LDAP server:

```
LDAPCertStoreParameters lcsp = new
    LDAPCertStoreParameters("ldap.sun.com", 389);
```

Next, create the CertStore object, and pass it the LDAPCertStoreParameters object, as in the following statement:

```
CertStore cs = CertStore.getInstance("LDAP", lcsp);
```

This call creates a CertStore object that retrieves certificates and CRLs from an LDAP repository using the schema defined in RFC 2587.

The following block of code establishes an x509CertSelector to retrieve all unexpired (as of the current date and time) end-entity certificates issued to a particular subject with 1) a key usage that allows digital signatures, and 2) a subject alternative name with a specific email address:

```
X509CertSelector xcs = new X509CertSelector();
```



```
// select only unexpired certificates
xcs.setCertificateValid(new Date());
// select only certificates issued to
// 'CN=alice, O=xyz, C=us'
xcs.setSubject(new X500Principal("CN=alice, O=xyz, C=us"));
// select only end-entity certificates
xcs.setBasicConstraints(-2);
// select only certificates with a digitalSignature
// keyUsage bit set (set the first entry in the
// boolean array to true)
boolean[] keyUsage = {true};
xcs.setKeyUsage(keyUsage);
// select only certificates with a subjectAltName of
// 'alice@xyz.example.com' (1 is the integer value of
// an RFC822Name)
xcs.addSubjectAlternativeName(1, "alice@xyz.example.com");
```

Then we pass the selector to the <code>getCertificates</code> method of our <code>CertStore</code> object that we previously created:

```
Collection<Certificate> certs = cs.getCertificates(xcs);
```

A PKIX CertPathBuilder may use similar code to help discover and sort through potential certificates by discarding those that do not meet validation constraints or other criteria.

The X509CRLSelector Class

The X509CRLSelector class is an implementation of the CRLSelector interface that defines a set of criteria for selecting X.509 CRLs.

An x509cRL object must match *all* of the specified criteria to be selected by the match method. The selection criteria are designed to be useful to a CertPathValidator or CertPathBuilder implementation that must retrieve CRLs from a repository to check the revocation status of certificates in an X.509 certification path.

For example, the setDateAndTime method of x509CRLSelector allows a PKIX CertPathValidator to filter out x509CRLs that have been issued after or expire before the time indicated. By setting this and other criteria in an x509CRLSelector object, it allows the CertPathValidator to discard irrelevant CRLs and more easily check if a certificate has been revoked.

Please refer to RFC 5280 for definitions of the X.509 CRL fields and extensions mentioned in this section.

Creating an X509CRLSelector Object

An X509CRLSelector object is created by calling the default constructor:

```
public X509CRLSelector()
```

No criteria are initially set (any X509CRL will match).



Setting Selection Criteria

The selection criteria allow a caller to match on different components of an X.509 CRL. Most of the methods for setting selection criteria are described here. Please refer to the x509CRLSelector class API documentation for details on the remaining methods.

The setIssuers and setIssuerNames methods set the issuerNames criterion:

```
public void setIssuers(Collection<X500Principal> issuers)
public void setIssuerNames(Collection<?> names)
```

The issuer distinguished name in the CRL must match at least one of the specified distinguished names. The setIssuers method is preferred as the use of x500Principals to represent distinguished names is more efficient and suitably typed. For the setIssuerNames method, each entry of the names argument is either a string or a byte array (representing the name, in RFC 2253 or ASN.1 DER encoded form, respectively). If null, any issuer distinguished name will do.

The setMinCRLNumber and setMaxCRLNumber methods set the minCRLNumber and maxCRLNumber criterion:

```
public void setMinCRLNumber(BigInteger minCRL)
public void setMaxCRLNumber(BigInteger maxCRL)
```

The CRL must have a CRL Number extension whose value is greater than or equal to the specified value if the <code>setMinCRLNumber</code> method is called, and less than or equal to the specified value if the <code>setMaxCRLNumber</code> method is called. If the value passed to one of these methods is null, the corresponding check is not done.

The setDateAndTime method sets the dateAndTime criterion:

```
public void setDateAndTime(Date dateAndTime)
```

The specified date must be equal to or later than the value of the thisUpdate component of the CRL and earlier than the value of the nextUpdate component. If null, no dateAndTime check will be done.

The ${\tt setCertificateChecking}$ method sets the certificate whose revocation status is being checked:

```
public void setCertificateChecking(X509Certificate cert)
```

This is not a criterion. Rather, it is optional information that may help a certstore find CRLs that would be relevant when checking revocation for the specified certificate. If null is specified, then no such optional information is provided. An application should always call this method when checking revocation for a particular certificate, as it may provide the certstore with more information for finding the correct CRLs and filtering out irrelevant ones.

Getting Selection Criteria

The current values for each of the selection criteria can be retrieved using an appropriate get method. Please refer to the x509CRLSelector Class API documentation for further details on these methods.



Creating an x509CRLSelector to retrieve CRLs from an LDAP repository is similar to the x509CertSelector example. Suppose we want to retrieve all current (as of the current date and time) CRLs issued by a specific CA and with a minimum CRL number. First, we create an x509CRLSelector object and call the appropriate methods to set the selection criteria:

```
X509CRLSelector xcrls = new X509CRLSelector();
// select CRLs satisfying current date and time
xcrls.setDateAndTime(new Date());
// select CRLs issued by 'O=xyz, C=us'
xcrls.addIssuerName("O=xyz, C=us");
// select only CRLs with a CRL number at least '2'
xcrls.setMinCRLNumber(new BigInteger("2"));
```

Then we pass the selector to the getCRLs method of our CertStore object (created in the X509CertSelector example):

```
Collection<CRL> crls = cs.getCRLs(xcrls);
```

PKIX Classes

The Java Certification Path API includes a set of algorithm-specific classes modeled for use with the PKIX certification path validation algorithm.

The PKIX certification path validation algorithm is defined in RFC 5280: *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*.

Topics

The TrustAnchor Class

The PKIXParameters Class

The CertPathValidatorResult Interface

The PolicyNode Interface and PolicyQualifierInfo Class

The PKIXBuilderParameters Class

The PKIXCertPathBuilderResult Class

The PKIXCertPathChecker Class

Using PKIXCertPathChecker in Certificate Path Validation

The TrustAnchor Class

The TrustAnchor class represents a "most-trusted CA", which is used as a trust anchor for validating X.509 certification paths.

A TrustAnchor includes the public key of the CA, the CA's name, and any constraints on the set of paths that can be validated using this key. These parameters can be specified in the form of a trusted X509Certificate or as individual parameters.

All TrustAnchor objects are immutable and thread-safe. That is, multiple threads may concurrently invoke the methods defined in this class on a single TrustAnchor object (or more than one) with no ill effects. Requiring TrustAnchor objects to be immutable and

thread-safe allows them to be passed around to various pieces of code without worrying about coordinating access.



Although this class is described as a PKIX class it may be used with other X. 509 certification path validation algorithms.

Creating a TrustAnchor Object

To instantiate a TrustAnchor object, a caller must specify "the most-trusted CA" as a trusted x509Certificate or public key and distinguished name pair. The caller may also optionally specify name constraints that are applied to the trust anchor by the validation algorithm during initialization. Note that support for name constraints on trust anchors is not required by the PKIX algorithm, therefore a PKIX CertPathValidator or CertPathBuilder may choose not to support this parameter and instead throw an exception. Use one of the following constructors to create a TrustAnchor object:

The nameConstraints parameter is specified as a byte array containing the ASN.1 DER encoding of a NameConstraints extension. An IllegalArgumentException is thrown if the name constraints cannot be decoded (are not formatted correctly).

Getting Parameter Values

Each of the parameters can be retrieved using a corresponding get method:

```
public final X509Certificate getTrustedCert()
public final X500Principal getCA()
public final String getCAName()
public final PublicKey getCAPublicKey()
public final byte[] getNameConstraints()
```



The getTrustedCert method returns null if the trust anchor was specified as a public key and name pair. Likewise, the getCA, getCAName and getCAPublicKey methods return null if the trust anchor was specified as an X509Certificate.

The PKIXParameters Class

The PKIXParametersClass class specifies the set of input parameters defined by the PKIX certification path validation algorithm. It also includes a few additional useful parameters.

This class implements the CertPathParameters interface.



An X.509 CertPath object and a PKIXParameters object are passed as arguments to the validate method of a CertPathValidator instance implementing the PKIX algorithm. The CertPathValidator uses the parameters to initialize the PKIX certification path validation algorithm.

Creating a PKIXParameters Object

To instantiate a PKIXParameters object, a caller must specify "the most-trusted CA(s)" as defined by the PKIX validation algorithm. The most-trusted CAs can be specified using one of two constructors:

```
public PKIXParameters(Set<TrustAnchor> trustAnchors)
    throws InvalidAlgorithmParameterException
public PKIXParameters(KeyStore keystore)
    throws KeyStoreException, InvalidAlgorithmParameterException
```

The first constructor allows the caller to specify the most-trusted CAs as a set of TrustAnchor objects. Alternatively, a caller can use the second constructor and specify a KeyStore instance containing trusted certificate entries, each of which will be considered as a most-trusted CA.

Setting Parameter Values

After a PKIXParameters object has been created, a caller can set (or replace the current value of) various parameters. A few of the methods for setting parameters are described here. Please refer to the PKIXParameters API documentation for details on the other methods.

The setInitialPolicies method sets the initial policy identifiers, as specified by the PKIX validation algorithm. The elements of the set are object identifiers (OIDs) represented as a String. If the initialPolicies parameter is null or not set, any policy is acceptable:

```
public void setInitialPolicies(Set<String> initialPolicies)
```

The setDate method sets the time for which the validity of the path should be determined. If the date parameter is not set or is null, the current date is used:

```
public void setDate(Date date)
```

The setPolicyMappingInhibited method sets the value of the policy mapping inhibited flag. The default value for the flag, if not specified, is false:

```
public void setPolicyMappingInhibited(boolean val)
```

The setExplicitPolicyRequired method sets the value of the explicit policy required flag. The default value for the flag, if not specified, is false:

```
public void setExplicitPolicyRequired(boolean val)
```

The setAnyPolicyInhibited method sets the value of the any policy inhibited flag. The default value for the flag, if not specified, is false:

public void setAnyPolicyInhibited(boolean val)



The setTargetCertConstraints method allows the caller to set constraints on the target or end-entity certificate. For example, the caller can specify that the target certificate must contain a specific subject name. The constraints are specified as a CertSelector object. If the selector parameter is null or not set, no constraints are defined on the target certificate:

public void setTargetCertConstraints(CertSelector selector)

The setCertStores method allows a caller to specify a List of CertStore objects that will be used by a PKIX implementation of CertPathValidator to find CRLs for path validation. This provides an extensible mechanism for specifying where to locate CRLs. The setCertStores method takes a List of CertStore objects as a parameter. The first CertStores in the list may be preferred to those that appear later.

public void setCertStores(List<CertStore> stores)

The setCertPathCheckers method allows a caller to extend the PKIX validation algorithm by creating implementation-specific certification path checkers. For example, this mechanism can be used to process private certificate extensions. The setCertPathCheckers method takes a list of PKIXCertPathChecker (discussed later) objects as a parameter:

public void setCertPathCheckers(List<PKIXCertPathChecker> checkers)

The setRevocationEnabled method allows a caller to disable revocation checking. Revocation checking is enabled by default, since it is a required check of the PKIX validation algorithm. However, PKIX does not define how revocation should be checked. An implementation may use CRLs or OCSP, for example. This method allows the caller to disable the implementation's default revocation checking mechanism if it is not appropriate. A different revocation checking mechanism can then be specified by calling the setCertPathCheckers method, and passing it a PKIXCertPathChecker that implements the alternate mechanism.

public void setRevocationEnabled(boolean val)

The setPolicyQualifiersRejected method allows a caller to enable or disable policy qualifier processing. When a PKIXParameters object is created, this flag is set to true. This setting reflects the most common (and simplest) strategy for processing policy qualifiers. Applications that want to use a more sophisticated policy must set this flag to false.

public void setPolicyQualifiersRejected(boolean qualifiersRejected)

Getting Parameter Values

The current values for each of the parameters can be retrieved using an appropriate get method. Please refer to the Class PKIXParameters API documentation for further details on these methods.

The PKIXCertPathValidatorResult Class

The PKIXCertPathValidatorResult class represents the result of the PKIX certification path validation algorithm.



This class implements the CertPathValidatorResult interface. It holds the valid policy tree and subject public key resulting from the validation algorithm, and includes methods (getPolicyTree() and getPublicKey()) for returning them. Instances of PKIXCertPathValidatorResult are returned by the validate method of CertPathValidator objects implementing the PKIX algorithm.

Please refer to the PKIXCertPathValidatorResult API documentation for more detailed information on this class.

The PolicyNode Interface and PolicyQualifierInfo Class

The PKIX validation algorithm defines several outputs related to certificate policy processing. Most applications will not need to use these outputs, but all providers that implement the PKIX validation or building algorithm must support them.

The PolicyNode interface represents a node of a valid policy tree resulting from a successful execution of the PKIX certification path validation. An application can obtain the root of a valid policy tree using the <code>getPolicyTree</code> method of <code>PKIXCertPathValidatorResult</code>. Policy Trees are discussed in more detail in the RFC 5280.

The getPolicyQualifiers method of PolicyNode returns a Set of PolicyQualifierInfo objects, each of which represents a policy qualifier contained in the Certificate Policies extension of the relevant certificate that this policy applies to.

Most applications will not need to examine the valid policy tree and policy qualifiers. They can achieve their policy processing goals by setting the policy-related parameters in PKIXParameters. However, the valid policy tree is available for more sophisticated applications, especially those that process policy qualifiers.

Please refer to the Interface PolicyNode and PolicyQualifierInfo API documentation for more detailed information on these classes.

Example 9-1 Example of Validating a Certification Path using the PKIX algorithm

This is an example of validating a certification path with the PKIX validation algorithm. The example ignores most of the exception handling and assumes that the certification path and public key of the trust anchor have already been created.

First, create the CertPathValidator, as in the following line:

```
CertPathValidator cpv = CertPathValidator.getInstance("PKIX");
```

The next step is to create a TrustAnchor object. This will be used as an anchor for validating the certification path. In this example, the most-trusted CA is specified as a public key and name (name constraints are not applied and are specified as null):

```
TrustAnchor anchor = new TrustAnchor("0=xyz,C=us", pubkey, null);
```

The next step is to create a PKIXParameters object. This will be used to populate the parameters used by the PKIX algorithm. In this example, we pass to the constructor a set containing a single element - the TrustAnchor we created in the previous step:

```
PKIXParameters params = new PKIXParameters(Collections.singleton(anchor));
```



Next, we populate the parameters object with constraints or other parameters used by the validation algorithm. In this example, we enable the explicitPolicyRequired flag and specify a set of initial policy OIDs (the contents of the set are not shown):

```
// set other PKIX parameters here
params.setExplicitPolicyRequired(true);
params.setInitialPolicies(policyIds);
```

The final step is to validate the certification path using the input parameter set we have created:

If the validation algorithm is successful, the policy tree and subject public key resulting from the validation algorithm are obtained using the <code>getPolicyTree</code> and <code>getPublicKey</code> methods of <code>PKIXCertPathValidatorResult</code>.

Otherwise, a CertPathValidatorException is thrown and the caller can catch the exception and print some details about the failure, such as the error message and the index of the certificate that caused the failure.

The PKIXBuilderParameters Class

The PKIXBuilderParameters class specifies the set of parameters to be used with CertPathBuilder Class.

This class (which extends the PKIXParameters class) specifies the set of parameters to be used with CertPathBuilder class that build certification paths validated against the PKIX certification path validation algorithm.

A PKIXBuilderParameters object is passed as an argument to the build method of a CertPathBuilder instance implementing the PKIX algorithm. All PKIX CertPathBuilders must return certification paths which have been validated according to the PKIX certification path validation algorithm.

Please note that the mechanism that a PKIX <code>certPathBuilder</code> uses to validate a constructed path is an implementation detail. For example, an implementation might attempt to first build a path with minimal validation and then fully validate it using an instance of a <code>PKIX CertPathValidator</code>, whereas a more efficient implementation may validate more of the path as it is building it, and backtrack to previous stages if it encounters validation failures or dead-ends.

Creating a PKIXBuilderParameters Object

Creating a PKIXBuilderParameters object is similar to creating a PKIXParameters object. However, a caller *must* specify constraints on the target or end-entity certificate when creating a PKIXBuilderParameters object. These constraints should provide the CertPathBuilder with enough information to find the target certificate. The constraints



are specified as a CertSelector object. Use one of the following constructors to create a PKIXBuilderParameters object:

Getting/Setting Parameter Values

The PKIXBuilderParameters class inherits all of the parameters that can be set in the PKIXParameters class. In addition, the setMaxPathLength method can be called to place a limit on the maximum number of certificates in a certification path:

```
public void setMaxPathLength(int maxPathLength)
```

The maxPathLength parameter specifies the maximum number of non-self-issued intermediate certificates that may exist in a certification path. A CertPathBuilder instance implementing the PKIX algorithm must not build paths longer than the length specified. If the value is 0, the path can only contain a single certificate. If the value is -1, the path length is unconstrained (i.e., there is no maximum). The default maximum path length, if not specified, is 5. This method is useful to prevent the CertPathBuilder from spending resources and time constructing long paths that may or may not meet the caller's requirements.

If any of the CA certificates in the path contain a Basic Constraints extension, the value of the pathLenConstraint component of the extension overrides the value of the maxPathLength parameter whenever the result is a certification path of smaller length. There is also a corresponding getMaxPathLength method for retrieving this parameter:

```
public int getMaxPathLength()
```

Also, the setCertStores method (inherited from the PKIXParameters class) is typically used by a PKIX implementation of CertPathBuilder to find Certificates for path construction as well as finding CRLs for path validation. This provides an extensible mechanism for specifying where to locate Certificates and CRLs.

The PKIXCertPathBuilderResult Class

The $\tt PKIXCertPathBuilderResult$ class represents the successful result of the PKIX certification path construction algorithm.

This class extends the PKIXCertPathValidatorResult class and implements the CertPathBuilder interface. Instances of PKIXCertPathBuilderResult are returned by the build method of CertPathBuilder objects implementing the PKIX algorithm.

The getCertPath method of a PKIXCertPathBuilderResult instance always returns a CertPath object validated using the PKIX certification path validation algorithm. The returned CertPath object does not include the most-trusted CA certificate that may have been used to anchor the path. Instead, use the getTrustAnchor method to get the Certificate of the most-trusted CA.



See the ${\tt PKIXCertPathBuilderResult}$ API documentation for more detailed information on this class.

Example 9-2 Example of Building a Certification Path using the PKIX algorithm

This is an example of building a certification path validated against the PKIX algorithm. Some details have been left out, such as exception handling, and the creation of the trust anchors and certificates for populating the certstore.

First, create the CertPathBuilder, as in the following example:

```
CertPathBuilder cpb = CertPathBuilder.getInstance("PKIX");
```

This call creates a CertPathBuilder object that returns paths validated against the PKIX algorithm.

The next step is to create a PKIXBuilderParameters object. This will be used to populate the PKIX parameters used by the CertPathBuilder:

```
// Create parameters object, passing it a Set of
// trust anchors for anchoring the path
// and a target subject DN.
X509CertSelector targetConstraints = new X509CertSelector();
targetConstraints.setSubject("CN=alice,O=xyz,C=us");
PKIXBuilderParameters params =
    new PKIXBuilderParameters(trustAnchors, targetConstraints);
```

The next step is to specify the <code>CertStore</code> that the <code>CertPathBuilder</code> will use to look for certificates and CRLs. For this example, we will populate a Collection <code>CertStore</code> with the certificates and CRLs:

```
CollectionCertStoreParameters ccsp =
   new CollectionCertStoreParameters(certsAndCrls);
CertStore store = CertStore.getInstance("Collection", ccsp);
params.addCertStore(store);
```

The next step is to build the certification path using the input parameter set we have created:

```
try {
    PKIXCertPathBuilderResult result =
          (PKIXCertPathBuilderResult) cpb.build(params);
    CertPath cp = result.getCertPath();
} catch (CertPathBuilderException cpbe) {
    System.out.println("build failed: " + cpbe.getMessage());
}
```

If the <code>CertPathBuilder</code> cannot build a path that meets the supplied parameters it will throw a <code>CertPathBuilderException</code>. Otherwise, the validated certification path can be obtained from the <code>PKIXCertPathBuilderResult</code> using the <code>getCertPath</code> method.

The PKIXCertPathChecker Class

The PKIXCertPathChecker class allows a user to extend a PKIX CertPathValidator Or CertPathBuilder implementation. This is an advanced feature that most users will not

need to understand. However, anyone implementing a PKIX service provider should read this section

The PKIXCertPathChecker class is an abstract class that executes one or more checks on an X.509 certificate. Developers should create concrete implementations of the PKIXCertPathChecker class when it is necessary to dynamically extend a PKIX CertPathValidator Or CertPathBuilder implementation at runtime. The following are a few examples of when a PKIXCertPathChecker implementation is useful:

- If the revocation mechanism supplied by a PKIX CertPathValidator or CertPathBuilder implementation is not adequate: For example, you can use the PKIXRevocationChecker (introduced in JDK 8; see Check Revocation Status of Certificates with PKIXRevocationChecker Class) to have more control over the revocation mechanism, or you can implement your own PKIXCertPathChecker to check that certificates have not been revoked.
- If the user wants to recognize certificates containing a critical private extension. Since the extension is private, it will not be recognized by the PKIX CertPathValidator Or CertPathBuilder implementation and a CertPathValidatorException will be thrown. In this case, a developer can implement a PKIXCertPathChecker that recognizes and processes the critical private extension.
- If the developer wants to record information about each certificate processed for debugging or display purposes.
- If the user wants to reject certificates with certain policy qualifiers.

The setCertPathCheckers method of the PKIXParameters class allows a user to pass a List of PKIXCertPathChecker Objects to a PKIX CertPathValidator Or CertPathBuilder implementation. Each of the PKIXCertPathChecker Objects will be called in turn, for each certificate processed by the PKIX CertPathValidator Or CertPathBuilder implementation.

Creating and using a PKIXCertPathChecker Object

The PKIXCertPathChecker class does not have a public constructor. This is intentional, since creating an instance of PKIXCertPathChecker is an implementation-specific issue. For example, the constructor for a PKIXCertPathChecker implementation that uses OCSP to check a certificate's revocation status may require the hostname and port of the OCSP server:

```
PKIXCertPathChecker checker = new OCSPChecker("ocsp.sun.com", 1321);
```

Once the checker has been instantiated, it can be added as a parameter using the addCertPathChecker method of the PKIXParameters class:

```
params.addCertPathChecker(checker);
```

Alternatively, a List of checkers can be added using the setCertPathCheckers method of the PKIXParameters class.

Implementing a PKIXCertPathChecker Object

The PKIXCertPathChecker class is abstract. It has four methods (check, getSupportedExtensions, init, and isForwardCheckingSupported) that all concrete subclasses must implement.



Implementing a PKIXCertPathChecker may be trivial or complex. A PKIXCertPathChecker implementation can be stateless or stateful. A stateless implementation does not maintain state between successive calls of the check method. For example, a PKIXCertPathChecker that checks that each certificate contains a particular policy qualifier is stateless. In contrast, a stateful implementation does maintain state between successive calls of the check method. The check method of a stateful implementation usually depends on the contents of prior certificates in the certification path. For example, a PKIXCertPathChecker that processes the NameConstraints extension is stateful.

Also, the order in which the certificates processed by a service provider implementation are presented (passed) to a PKIXCertPathChecker is very important, especially if the implementation is stateful. Depending on the algorithm used by the service provider, the certificates may be presented in *reverse* or *forward* order. A reverse ordering means that the certificates are ordered from the most trusted CA (if present) to the target subject, whereas a forward ordering means that the certificates are ordered from the target subject to the most trusted CA. The order must be made known to the PKIXCertPathChecker implementation, so that it knows how to process consecutive certificates.

Initializing a PKIXCertPathChecker Object

The init method initializes the internal state of the checker:

```
public abstract void init(boolean forward)
```

All stateful implementations should clear or initialize any internal state in the checker. This prevents a service provider implementation from calling a checker that is in an uninitialized state. It also allows stateful checkers to be reused in subsequent operations without reinstantiating them. The forward parameter indicates the order of the certificates presented to the PKIXCertPathChecker. If forward is true, the certificates are presented from target to trust anchor; if false, from trust anchor to target.

Forward Checking

The isForwardCheckingSupported method returns a boolean that indicates if the PKIXCertPathChecker supports forward checking:

```
public abstract boolean isForwardCheckingSupported()
```

All PKIXCertPathChecker implementations *must*support reverse checking. A PKIXCertPathChecker implementation *maysupport* forward checking.

Supporting forward checking improves the efficiency of <code>CertPathBuilders</code> that build forward, since it allows paths to be checked as they are built. However, some stateful <code>PKIXCertPathCheckers</code> may find it difficult or impossible to support forward checking.

Supported Extensions

The getSupportedExtensions method returns an immutable set of OID Strings for the X. 509 extensions that the PKIXCertPathChecker implementation supports (i.e., recognizes, is able to process):

public abstract Set<String> getSupportedExtensions()



The method should return <code>null</code> if no extensions are processed. All implementations should return the <code>set</code> of OID <code>Strings</code> that the <code>check</code> method may process.

A CertPathBuilder can use this information to identify certificates with unrecognized critical extensions, even when performing a forward build with a PKIXCertPathChecker that does not support forward checking.

Executing the Check

The following method executes a check on the certificate:

The unresolvedCritExts parameter contains a collection of OIDs as strings. These OIDs represent the set of critical extensions in the certificate that have not yet been resolved by the certification path validation algorithm. Concrete implementations of the check method should remove any critical extensions that it processes from the unresolvedCritExts parameter.

If the certificate does not pass the check(s), a CertPathValidatorException should be thrown.

Cloning a PKIXCertPathChecker

The PKIXCertPathChecker class implements the Cloneable interface. All stateful PKIXCertPathChecker implementations must override the clone method if necessary. The default implementation of the clone method calls the object.clone method, which performs a simple clone by copying all fields of the original object to the new object. A stateless implementation should not override the clone method. However, all stateful implementations must ensure that the default clone method is correct, and override it if necessary. For example, a PKIXCertPathChecker that stores state in an array must override the clone method to make a copy of the array, rather than just a reference to the array.

The reason that PKIXCertPathChecker objects are cloneable is to allow a PKIX CertPathBuilder implementation to efficiently backtrack and try another path when a potential certification path reaches a dead end or point of failure. In this case, the implementation is able to restore prior path validation states by restoring the cloned objects.

Example 9-3 Sample Code to Check for a Private Extension

This is an example of a stateless PKIXCertPathChecker implementation. It checks if a private extension exists in a certificate and processes it according to some rules.



```
* Initialize checker
public void init(boolean forward)
    throws CertPathValidatorException {
    // nothing to initialize
public Set getSupportedExtensions() {
   return supportedExtensions;
public boolean isForwardCheckingSupported() {
   return true;
 * Check certificate for presence of Netscape's
 * private extension
 * with OID "2.16.840.1.113730.1.1"
public void check(Certificate cert,
                  Collection unresolvedCritExts)
    throws CertPathValidatorException
    X509Certificate xcert = (X509Certificate) cert;
    byte[] ext =
        xcert.getExtensionValue("2.16.840.1.113730.1.1");
    if (ext == null)
        return;
    // process private extension according to some
    // rules - if check fails, throw a
    // CertPathValidatorException ...
    // {insert code here}
    // remove extension from collection of unresolved
    // extensions (if it exists)
    if (unresolvedCritExts != null)
        unresolvedCritExts.remove("2.16.840.1.113730.1.1");
```

How a PKIX Service Provider implementation should use a PKIXCertPathChecker

Each PKIXCertPathChecker object must be initialized by a service provider implementation before commencing the build or validation algorithm, for example:

```
List<PKIXCertPathChecker> checkers = params.getCertPathCheckers();
for (PKIXCertPathChecker checker : checkers) {
    checker.init(false);
}
```

For each certificate that it validates, the service provider implementation must call the check method of each PKIXCertPathChecker object in turn, passing it the certificate and any remaining unresolved critical extensions:

```
for (PKIXCertPathChecker checker : checkers) {
   checker.check(cert, unresolvedCritExts);
}
```

If any of the <code>checks</code> throw a <code>CertPathValidatorException</code>, a <code>CertPathValidator</code> implementation should terminate the validation procedure. However, a <code>CertPathBuilder</code> implementation may simply log the failure and continue to search for other potential paths. If all of the <code>checks</code> are successful, the service provider implementation should check that all critical extensions have been resolved and if not, consider the validation to have failed. For example:

As discussed in the previous section, a <code>CertPathBuilder</code> implementation may need to backtrack when a potential certification path reaches a dead end or point of failure. Backtracking in this context implies returning to the previous certificate in the path and checking for other potential paths. If the <code>CertPathBuilder</code> implementation is validating the path as it is building it, it will need to restore the previous state of each <code>PKIXCertPathChecker</code>. It can do this by making clones of the <code>PKIXCertPathChecker</code> objects <code>before</code> each certificate is processed, for example:

```
/* clone checkers */
List newList = new ArrayList(checkers);
ListIterator li = newList.listIterator();
while (li.hasNext()) {
    PKIXCertPathChecker checker = (PKIXCertPathChecker) li.next();
    li.set(checker.clone());
}
```

Using PKIXCertPathChecker in Certificate Path Validation

Using a PKIXCertPathChecker to customize certificate path validation is relatively straightforward.

Basic Certification Path Validation

First, consider code that validates a certificate path:



If the validation fails, the validate() method throws an exception.

The fundamental steps are as follows:

- 1. Obtain the CA root certificates and the certification path to be validated.
- 2. Create a PKIXParameters with the trust anchors.
- 3. Use a CertPathValidator to validate the certificate path.

In this example, <code>getTrustAnchors()</code> and <code>getCertPath()</code> are the methods that obtain CA root certificates and the certification path.

The <code>getTrustAnchors()</code> method in the example must return a <code>Set of TrustAnchors</code> that represent the CA root certificates you wish to use for validation. Here is one simple implementation that loads a single CA root certificate from a file:

```
public Set<TrustAnchor> getTrustAnchors()
    throws IOException, CertificateException {
    CertificateFactory cf = CertificateFactory.getInstance("X.509");
    X509Certificate c;
    try (InputStream in = new FileInputStream("x509_ca-certificate.cer")) {
        c = (X509Certificate)cf.generateCertificate(in);
    }
    TrustAnchor anchor = new TrustAnchor(c, null);
    return Collections.singleton(anchor);
}
```

Similarly, here is a simple implementation of getCertPath() that loads a certificate path from a file:

```
public CertPath getCertPath() throws IOException, CertificateException {
   CertificateFactory cf = CertificateFactory.getInstance("X.509");

   CertPath cp;
   try (InputStream in = new FileInputStream("certpath.pkcs7")) {
     cp = cf.generateCertPath(in, "PKCS7");
   }
   return cp;
}
```

Note that PKCS#7 does not require a specific order for the certificates in the file, so this code only works for certification path validation when the certificates are ordered starting from the entity to be validated and progressing back toward the CA root. If the certificates are not in the right order, you need to do some additional processing. CertificateFactory has a generateCertPath() method that accepts a Collection, which is useful for this type of processing.

Adding in a PKIXCertPathChecker

To customize certification path validation, add a PKIXCertPathChecker as follows. In this example, SimpleChecker is a PKIXCertPathChecker subclass. The new lines are shown in **bold**.

```
Set<TrustAnchor> trustAnchors = getTrustAnchors();
CertPath cp = getCertPath();
```



SimpleChecker is a rudimentary subclass of PKIXCertPathChecker. Its check() method is called for every certificate in the certification path that is being validated. SimpleChecker uses an AlgorithmConstraints implementation to examine the signature algorithm and public key of each certificate.

```
import java.security.AlgorithmConstraints;
import java.security.CryptoPrimitive;
import java.security.Key;
import java.security.cert.*;
import java.util.*;
public class SimpleChecker extends PKIXCertPathChecker {
  private final static Set<CryptoPrimitive> SIGNATURE PRIMITIVE SET =
      EnumSet.of(CryptoPrimitive.SIGNATURE);
  public void init(boolean forward) throws CertPathValidatorException {}
  public boolean isForwardCheckingSupported() { return true; }
  public Set<String> getSupportedExtensions() { return null; }
  public void check(Certificate cert,
      Collection<String> unresolvedCritExts)
      throws CertPathValidatorException {
    X509Certificate c = (X509Certificate)cert;
    String sa = c.getSigAlgName();
    Key key = c.getPublicKey();
    AlgorithmConstraints constraints = new SimpleConstraints();
    if (constraints.permits(SIGNATURE_PRIMITIVE_SET, sa, null) == false)
      throw new CertPathValidatorException("Forbidden algorithm: " + sa);
    if (constraints.permits(SIGNATURE_PRIMITIVE_SET, key) == false)
      throw new CertPathValidatorException("Forbidden key: " + key);
```

Finally, SimpleConstraints is an AlgorithmConstraints implementation that requires RSA 2048.

```
import java.security.AlgorithmConstraints;
import java.security.AlgorithmParameters;
import java.security.CryptoPrimitive;
import java.security.Key;
import java.security.interfaces.RSAKey;
import java.util.Set;
```



```
public class SimpleConstraints implements AlgorithmConstraints {
   public boolean permits(Set<CryptoPrimitive> primitives,
        String algorithm, AlgorithmParameters parameters) {
        return permits(primitives, algorithm, null, parameters);
   }

   public boolean permits(Set<CryptoPrimitive> primitives, Key key) {
        return permits(primitives, null, key, null);
   }

   public boolean permits(Set<CryptoPrimitive> primitives,
        String algorithm, Key key, AlgorithmParameters parameters) {
        if (algorithm == null) algorithm = key.getAlgorithm();

        if (algorithm.indexOf("RSA") == -1) return false;

        if (key != null) {
            RSAKey rsaKey = (RSAKey)key;
            int size = rsaKey.getModulus().bitLength();
            if (size < 2048) return false;
        }

        return true;
   }
}</pre>
```

Check Revocation Status of Certificates with PKIXRevocationChecker Class

An instance of PKIXRevocationChecker checks the revocation status of certificates with the Online Certificate Status Protocol (OCSP) or Certificate Revocation Lists (CRLs).

The PKIXRevocationChecker (introduced in JDK 8), which is a subclass of PKIXCertPathChecker, checks the revocation status of certificates with the PKIX algorithm.

An instance of PKIXRevocationChecker checks the revocation status of certificates with the Online Certificate Status Protocol (OCSP) or Certificate Revocation Lists (CRLs). OCSP is described in RFC 2560 and is a network protocol for determining the status of a certificate. A CRL is a time-stamped list identifying revoked certificates, and RFC 5280 describes an algorithm for determining the revocation status of certificates using CRLs.

Each PKIX CertPathValidator and CertPathBuilder instance provides a default revocation implementation that is enabled by default. If you want more control over the revocation settings used by that implementation, use the PKIXRevocationChecker class.

Follow these general steps to check the revocation status of a certificate path with the PKIXRevocationChecker Class:

- Obtain a PKIXRevocationChecker instance by calling the getRevocationChecker method of a PKIX CertPathValidator or CertPathBuilder instance.
- 2. Set additional parameters and options specific to certificate revocation with methods contained in the PKIXRevocationChecker class. These methods include setOCSPResponder(URI), which sets the URI that identifies the location of the OCSP responder (although normally the URI is included in the certificate and does not have to be set) and setOptions(Set<PKIXRevocationChecker.Option>), which sets revocation options. PKIXRevocationChecker.Option is an enumerated type used to specify the following options:



- ONLY_END_ENTITY: Only check the revocation status of end-entity certificates.
- PREFER_CRLS: By default, OCSP is the preferred mechanism for checking revocation status, with CRLs as the fallback mechanism. Switch this preference to CRLs with this option.
- SOFT_FAIL: Ignore network failures.
- 3. After obtaining an instance of PKIXRevocationChecker, add it to a PKIXParameters or PKIXBuilderParameters object with the addCertPathChecker or setCertPathCheckers method.
- **4.** Follow one of these steps depending on whether you are using a PKIX CertPathValidator **Of** CertPathBuilder **instance**:
 - If you are using a PKIX CertPathValidator instance, call the validate method using as arguments the certificate path you want to validate and the PKIXParameters object that contains a revocation checker.
 - If you are using a PKIX CertPathBuilder instance, call the build method using as arguments the PKIXBuilderParameters object that contains a revocation checker
- 5. Call the validate method of the PKIX CertPathValidator or CertPathBuilder instance using as arguments the certificate path you want to validate and the PKIXParameters or PKIXBuilderParameters object that contains a revocation checker.

The following excerpt checks the revocation status of certificates contained in a certificate path. The CertPath object path is the certificate path, and params is an object of type PKIXParameters:

```
CertPathValidator cpv = CertPathValidator.getInstance("PKIX");
PKIXRevocationChecker rc = (PKIXRevocationChecker)cpv.getRevocationChecker();
rc.setOptions(EnumSet.of(Option.SOFT_FAIL));
params.addCertPathChecker(rc);
params.setRevocationEnabled(false);
CertPathValidatorResult res = cpv.validate(path, params);
```

In this excerpt, the <code>SOFT_FAIL</code> option causes the revocation checker to ignore any network failures (such as failing to establish a connection to the OCSP server) when it checks the revocation status.

Implementing a Service Provider

Experienced programmers can create their own provider packages supplying certification path service implementations.

This section assumes that you have read Java Cryptography Architecture (JCA) Reference Guide.

The following engine classes are defined in the Java Certification Path API:

- CertPathValidator used to validate certification paths
- CertPathBuilder used to build certification paths
- Certstore used to retrieve certificates and CRLs from a repository

In addition, the pre-existing <code>CertificateFactory</code> engine class also supports the generation of certification paths.



The application interfaces supplied by an engine class are implemented in terms of a "Service Provider Interface" (SPI). The name of each SPI class is the same as that of the corresponding engine class, followed by "Spi". For example, the SPI class corresponding to the CertPathValidator engine class is the CertPathValidatorSpi class. Each SPI class is abstract. To supply the implementation of a particular type of service, for a specific algorithm or type, a provider must subclass the corresponding SPI class and provide implementations for all the abstract methods. For example, the CertStore class provides access to the functionality of retrieving certificates and CRLs from a repository. The actual implementation supplied in a CertStoreSpi subclass would be that for a specific type of certificate repository, such as LDAP.

Steps to Implement and Integrate a Provider

When implementing and integrating a provider for the certification path services, you must ensure that certain information is provided.

Developers should follow the Steps to Implement and Integrate a Provider. Here are some additional rules to follow for certain steps:

Step 3: Write your "Master Class", a subclass of Provider

In Step 3: Write Your Master Class, a Subclass of Provider these are the properties that must be defined for the certification path services, where the algorithm name is substituted for *algName*, and certstore type for *storeType*:

- CertPathValidator.algName
- CertPathBuilder.algName
- CertStore.storeType

See Java Security Standard Algorithm Names Specification for the standard names that are defined for *algName* and *storeType*. The value of each property must be the fully qualified name of the class implementing the specified algorithm, or certstore type. That is, it must be the package name followed by the class name, where the two are separated by a period. For example, a provider sets the <code>CertPathValidator.PKIX</code> property to have the value

"sun.security.provider.certpath.PKIXCertPathValidator" as follows:

```
put("CertPathValidator.PKIX", "sun.security.provider.certpath.PKIXCertPathValidator")
```

In addition, service attributes can be defined for the certification path services. These attributes can be used as filters for selecting service providers. See Appendix A for the definition of some standard service attributes. For example, a provider may set the ValidationAlgorithm service attribute to the name of an RFC or specification that defines the PKIX validation algorithm:

```
put("CertPathValidator.PKIX ValidationAlgorithm", "RFC5280");
```

Step 11: Document your Provider and its Supported Services

In Step 12: Document Your Provider and Its Supported Services, certification path service providers should document the following information for each SPI:

Certificate Factories



A provider should document what types of certification paths (and the version numbers of the certificates in the path, if relevant) can be created by the factory. A provider should describe the ordering of the certificates in the certification path, as well as the contents.

A provider should document the list of encoding formats supported. This is not technically necessary, since the client can request them by calling the <code>getCertPathEncodings</code> method. However, the documentation should describe each encoding format in more detail and reference any standards when applicable.

Certification Path Validators

A provider should document any relevant information regarding the <code>CertPathValidator</code> implementation, including the types of certification paths that it validates. In particular, a PKIX <code>CertPathValidator</code> implementation should document the following information:

- The RFC or specification it is compliant with.
- The mechanism it uses to check that certificates have not been revoked.
- Any optional certificate or CRL extensions that it recognizes and how it processes them.

Certification Path Builders

A provider should document any relevant information regarding the <code>CertPathBuilder</code> implementation, including the types of certification paths that it creates and whether or not they are validated. In particular a PKIX <code>CertPathBuilder</code> implementation should document the following information:

- The RFC or specification it is compliant with.
- The mechanism it uses to check that certificates have not been revoked.
- Any optional certificate or CRL extensions that it recognizes and how it processes them.
- Details on the algorithm it uses for finding certification paths. Ex: depth-first, breadth-first, forward (i.e., from target to trust anchor(s)), reverse (i.e., from trust anchor(s) to target).
- The algorithm it uses to select and sort potential certificates. For example, given
 two certificates that are potential candidates for the next certificate in the path,
 what criteria are used to select one before the other? What criteria are used to
 reject a certificate?
- If applicable, the algorithm it uses for backtracking or constructing another path (i.e., when potential paths do not meet constraints).
- The types of CertStore implementations that have been tested. The
 implementation should be designed to work with any CertStore type, but this
 information may still be useful.

All CertPathBuilder implementations should provide additional debugging support, in order to analyze and correct potential path building problems. Details on how to access this debugging information should be documented.

Certificate/CRL Stores

A provider should document what types of certificates and CRLs (and the version numbers, if relevant) are retrieved by the $\tt CertStore$.



A provider should also document any relevant information regarding the <code>certStore</code> implementation (such as protocols used or formats supported). For example, an LDAP <code>certStore</code> implementation should describe which versions of LDAP are supported and which standard attributes are used for finding certificates and CRLs. It should also document if the implementation caches results, and for how long (i.e., under what conditions are they refreshed).

If the implementation returns the certificates and CRLs in a particular order, it should describe the sorting algorithm. An implementation should also document any additional or default initialization parameters. Finally, an implementation should document if and how it uses information in the CertSelector or CRLSelector objects to find certificates and CRLs.

Service Interdependencies

Common types of algorithm interdependencies in certification path service implementations.

The following are some common types of algorithm interdependencies in certification path service implementations:

Certification Path Validation and Signature Algorithms

A CertPathValidator implementation often requires use of a signature algorithm to verify each certificate's digital signature. The setSigProvider method of the PKIXParameters class allows a user to specify a specific Signature provider.

Certification Path Builders and Certificate Factories

A CertPathBuilder implementation will often utilize a CertificateFactory to generate a certification path from a list of certificates.

· CertStores and Certificate Factories

A CertStore implementation will often utilize a CertificateFactory to generate certificates and CRLs from their encodings. For example, an LDAP CertStore implementation may use an X.509 CertificateFactory to generate X.509 certificates and CRLs from their ASN.1 encoded form.

Certification Path Parameter Specification Interfaces

The Certification Path API contains two interfaces representing *transparent* specifications of parameters, the CertPathParameters and CertStoreParameters interfaces.

Two implementations of the CertPathParameters interface are included, the PKIXParameters and PKIXBuilderParameters classes. If you are working with PKIX certification path validation and algorithm parameters, you can utilize these classes. If you need parameters for a different algorithm, you will need to supply your own CertPathParameters implementation for that algorithm.

Two implementations of the CertStoreParameters interface are included, the LDAPCertStoreParameters and the CollectionCertStoreParameters classes. These classes are to be used with LDAP and Collection CertStore implementations, respectively. If you need parameters for a different repository type, you will need to supply your own CertStoreParameters implementation for that type.

The CertPathParameters and CertStoreParameters interfaces each define a clone method that implementations should override. A typical implementation will perform a

"deep" copy of the object, such that subsequent changes to the copy will not affect the original (and vice versa). However, this is not an absolute requirement for implementations of CertStoreParameters. A shallow copy implementation of clone is more appropriate for applications that need to hold a reference to a parameter contained in the CertStoreParameters. For example, since CertStore.getInstance makes a clone of the specified CertStoreParameters, a shallow copy clone allows an application to hold a reference to and later release the resources of a particular CertStore initialization parameter, rather than waiting for the garbage collection mechanism. This should be done with the utmost care, since the CertStore may still be in use by other threads.

Certification Path Result Specification Interfaces

The Certification Path API contains two interfaces representing *transparent* specifications of results, the CertPathValidatorResult and CertPathBuilderResult interfaces.

One implementation for each of the interfaces is included: the PKIXCertPathValidatorResult and PKIXCertPathBuilderResult Classes. If you are implementing PKIX certification path service providers, you can utilize these classes. If you need certification path results for a different algorithm, you will need to supply your own CertPathValidatorResult or CertPathBuilderResult implementation for that algorithm.

A PKIX implementation of a CertPathValidator or a CertPathBuilder may find it useful to store additional information in the PKIXCertPathValidatorResult or PKIXCertPathBuilderResult, such as debugging traces. In these cases, the implementation should implement a subclass of the appropriate result class with methods to retrieve the relevant information. These classes must be shipped with the provider classes, for example, as part of the provider JAR file.

Certification Path Exception Classes

The Certification Path API contains a set of exception classes for handling errors. CertPathValidatorException, CertPathBuilderException, and CertStoreException are subclasses of GeneralSecurityException.

You may need to extend these classes in your service provider implementation.

For example, a CertPathBuilder implementation may provide additional information such as debugging traces when a CertPathBuilderException is thrown. The implementation may throw a subclass of CertPathBuilderException that holds this information. Likewise, a CertStore implementation can provide additional information when a failure occurs by throwing a subclass of CertStoreException . Also, you may want to implement a subclass of CertPathValidatorException to describe a particular failure mode of your CertPathValidator implementation.

In each case, the new exception classes must be shipped with the provider classes, for example, as part of the provider JAR file. Each provider should document the exception subclasses.

Appendix A: Standard Names

The Java Certification Path API requires and utilizes a set of standard names for certification path validation algorithms, encodings and certificate storage types.

The standard names previously found here in Appendix A and in the other security specifications (JCA/JSSE/etc.) have been combined in the Java Security Standard Algorithm Names Specification. Specific provider information can be found in the JDK Providers.

Please note that a service provider may choose to define a new name for a proprietary or non-standard algorithm that is not mentioned in the Standard Names document. However, to prevent name collisions, it is recommended that the name be prefixed with the reverse Internet domain name of the provider's organization (for example: com.sun.MyCertPathValidator).

Appendix B: CertPath Implementation in SUN Provider

The "SUN" provider supports the following standard algorithms, types and encodings:

- CertificateFactory: X.509 CertPath type with PKCS7 and PkiPath encodings
- CertPathValidator: PKIX algorithm
- CertPathBuilder: PKIX algorithm
- CertStore: Collection CertStore type

Each of these service provider interface implementations is discussed in more detail below.

CertificateFactory

The "SUN" provider for the <code>CertificateFactory</code> engine class supports generation of X. 509 <code>CertPath</code> objects. The PKCS7 and PkiPath encodings are supported. The PKCS#7 implementation supports a subset of RFC 2315 (only the SignedData ContentInfo type is supported). The certificates in the <code>CertPath</code> are ordered in the forward direction (from target to trust anchor). Each certificate in the <code>CertPath</code> is of type <code>java.security.cert.X509Certificate</code>, and versions 1, 2 and 3 are supported.

CertPathValidator

The "SUN" provider supplies a PKIX implementation of the CertPathValidator engine class. The implementation validates CertPaths of type X.509 and implements the certification path validation algorithm defined in RFC 5280: PKIX Certificate and CRL Profile. This implementation sets the ValidationAlgorithm service attribute to "RFC5280".

Weak cryptographic algorithms can be disabled in the "SUN" provider using the jdk.certpath.disabledAlgorithms Security Property. See Appendix E: Disabling Cryptographic Algorithms for a description and examples of this property.

The PKIX Certificate and CRL Profile has many optional features. The "SUN" provider implements support for the policy mapping, authority information access and CRL distribution point certificate extensions, the issuing distribution point CRL extension, and the reason code and certificate issuer CRL entry extensions. It does not implement support for the freshest CRL or subject information access certificate extensions. It also does not include support for the freshest CRL and delta CRL Indicator CRL extensions and the invalidity date and hold instruction code CRL entry extensions.

The implementation supports a CRL revocation checking mechanism that conforms to section 6.3 of the PKIX Certificate and CRL Profile. OCSP (RFC 2560) is also



currently supported as a built in revocation checking mechanism. See Appendix C: OCSP Support for more details on the implementation and configuration and how it works in conjunction with CRLs.

The implementation does not support the nameConstraints parameter of the TrustAnchor class and the validate method throws an InvalidAlgorithmParameterException if it is specified.

CertPathBuilder

The "SUN" provider supplies a PKIX implementation of the CertPathBuilder engine class. The implementation builds CertPaths of type X.509. Each CertPath is validated according to the PKIX algorithm defined in RFC 5280: PKIX Certificate and CRL Profile. This implementation sets the ValidationAlgorithm service attribute to "RFC5280".

The implementation requires that the targetConstraints parameter of a PKIXBuilderParameters object is an instance of X509CertSelector and the subject criterion is set to a non-null value. Otherwise the build method throws an InvalidAlgorithmParameterException.

The implementation builds <code>CertPath</code> objects in a forward direction using a depth-first algorithm. It backtracks to previous states and tries alternate paths when a potential path is determined to be invalid or exceeds the <code>PKIXBuilderParameters maxPathLength</code> parameter.

Validation of the path is performed in the same manner as the <code>certPathValidator</code> implementation. The implementation validates most of the path as it is being built, in order to eliminate invalid paths earlier in the process. Validation checks that cannot be executed on certificates ordered in a forward direction are delayed and executed on the path after it has been constructed (but before it is returned to the application).

As with CertPathValidator, the jdk.certpath.disabledAlgorithms Security Property can be used to exclude cryptographic algorithms that are not considered safe.

When two or more potential certificates are discovered that may lead to finding a path that meets the specified constraints, the implementation uses the following criteria to prioritize the certificates (in the examples below, assume a TrustAnchor distinguished name of "ou=D,ou=C,o=B,c=A" is specified):

- 1. The issuer DN of the certificate matches the DN of one of the specified TrustAnchors (ex: issuerDN = "ou=D,ou=C,o=B,c=A").
- 2. The issuer DN of the certificate is a descendant of the DN of one of the TrustAnchors, ordered by proximity to the anchor (ex: issuerDN = "ou=E,ou=D,ou=C,o=B,c=A").
- 3. The issuer DN of the certificate is an ancestor of the DN of one of the TrustAnchors, ordered by proximity to the anchor (ex: issuerDN = "ou=C,o=B,c=A".
- 4. The issuer DN of the certificate is in the same namespace of one of the TrustAnchors, ordered by proximity to the anchor (ex: issuerDN = "ou=G,ou=C,o=B,c=A").
- 5. The issuer DN of the certificate is an ancestor of the subject DN of the certificate, ordered by proximity to the subject.

These are followed by certificates which don't meet any of the above criteria.



This implementation has been tested with the LDAP and Collection <code>CertStore</code> implementations included in this release of the "SUN" provider.

Debugging support can be enabled by setting the <code>java.security.debug</code> property to <code>certpath</code>. For example:

java -Djava.security.debug=certpath BuildCertPath

This will print additional debugging information to standard error.

Collection CertStore

The SUN provider supports the Collection implementation of the certstore engine class.

The Collection CertStore implementation can hold any objects that are an instance of java.security.cert.Certificate Or java.security.cert.CRL.

The certificates and CRLs are not returned in any particular order and will not contain duplicates.

Support for the CRL Distribution Points Extension

Support for the CRL Distribution Points extension is available. It is disabled by default for compatibility and can be enabled by setting the system property com.sun.security.enableCRLDP to the value true.

If set to true, Sun's PKIX implementation uses the information in a certificate's CRL Distribution Points extension (in addition to CertStores that are specified) to find the CRL, provided the distribution point is an X.500 distinguished name or a URI of type Idap, http, or ftp.



Depending on your network and firewall setup, it may be necessary to also configure your networking proxy servers.

Support for the Authority Information Access (AIA) Extension

Support for the calssuers access method of the Authority Information Access extension is available. It is disabled by default for compatibility and can be enabled by setting the system property com.sun.security.enableAIAcaIssuers to the value true.

If set to true, Sun's PKIX implementation of <code>CertPathBuilder</code> uses the information in a certificate's AIA extension (in addition to CertStores that are specified) to find the issuing CA certificate, provided it is a URI of type Idap, http, or ftp.

Note:

Depending on your network and firewall setup, it may be necessary to also configure your networking proxy servers.



Appendix C: OCSP Support

Client-side support for the On-Line Certificate Status Protocol (OCSP) as defined in RFC 2560 is supported.

OCSP checking is controlled by the following five Security Properties:

Property Name	Description
ocsp.enable	This property's value is either true or false. If true, OCSP checking is enabled when doing certificate revocation checking; if false or not set, OCSP checking is disabled.
ocsp.responderURL	This property's value is a URL that identifies the location of the OCSP responder. Here is an example
	<pre>ocsp.responderURL=http:// ocsp.example.net:80</pre>
	By default, the location of the OCSP responder is determined implicitly from the certificate being validated. The property is used when the Authority Information Access extension (defined in RFC 5280) is absent from the certificate or when it requires overriding.
ocsp.responderCertSubjectName	This property's value is the subject name of the OCSP responder's certificate. Here is an example
	ocsp.responderCertSubjectName="CN=OCSP Responder, O=XYZ Corp"
	By default, the certificate of the OCSP responder is that of the issuer of the certificate being validated. This property identifies the certificate of the OCSP responder when the default does not apply. Its value is a string distinguished name (defined in RFC 2253) which identifies a certificate in the set of certificates supplied during cert path validation. In cases where the subject name alone is not sufficient to uniquely identify the certificate, then both the ocsp.responderCertIssuerName and ocsp.responderCertSerialNumber properties must be used instead. When this property is set, then those two properties are ignored.



Property Name	Description
ocsp.responderCertIssuerName	This property's value is the issuer name of the OCSP responder's certificate . Here is an example
	<pre>ocsp.responderCertIssuerName="CN=Enterpri se CA, O=XYZ Corp"</pre>
	By default, the certificate of the OCSP responder is that of the issuer of the certificate being validated. This property identifies the certificate of the OCSP responder when the default does not apply. Its value is a string distinguished name (defined in RFC 2253) which identifies a certificate in the set of certificates supplied during cert path validation. When this property is set then the ocsp.responderCertSerialNumber property must also be set. Note that this property is ignored when the ocsp.responderCertSubjectName property has been set.
ocsp.responderCertSerialNumber	This property's value is the serial number of the OCSP responder's certificate Here is an example
	ocsp.responderCertSerialNumber=2A:FF:00
	By default, the certificate of the OCSP responder is that of the issuer of the certificate being validated. This property identifies the certificate of the OCSP responder when the default does not apply. Its value is a string of hexadecimal digits (colon or space separators may be present) which identifies a certificate in the set of certificates supplied during cert path validation. When this property is set then the ocsp.responderCertIssuerName property must also be set. Note that this property is ignored when the ocsp.responderCertSubjectName property has been set.

These properties may be set either statically in the Java runtime's <java_home>/conf/security/java.security file, or dynamically using the java.security.Security.setProperty() method.

By default, OCSP checking is not enabled. It is enabled by setting the <code>ocsp.enable</code> property to "true". Use of the remaining properties is optional. Note that enabling OCSP checking only has an effect if revocation checking has also been enabled. Revocation checking is enabled via the

PKIXParameters.setRevocationEnabled() method.

OCSP checking works in conjunction with Certificate Revocation Lists (CRLs) during revocation checking. Below is a summary of the interaction of OCSP and CRLs. Failover to CRLs occurs only if an OCSP problem is encountered. Failover does not

occur if the OCSP responder confirms either that the certificate has been revoked or that it has not been revoked.

PKIXParameters RevocationEnabled (default=true)	ocsp.enable (default=false)	Behavior
true	true	Revocation checking using OCSP, failover to using CRLs
true	false	Revocation checking using CRLs only
false	true	No revocation checking
false	false	No revocation checking

Appendix D: CertPath Implementation in JdkLDAP Provider

The JdkLDAP provider supports the LDAP implementation of the <code>certStore</code> engine class.

LDAP CertStore

The LDAP certstore implementation retrieves certificates and CRLs from an LDAP directory using the LDAP schema defined in RFC 2587.

The LDAPSchema service attribute is set to "RFC2587".

The implementation fetches certificates from different locations, depending on the values of the subject, issuer, and basicConstraints selection criteria specified in the x509CertSelector. It performs as many of the following operations as possible:

- Subject non-null, basicConstraints <= -1
 <= -1
 <p>Looks for certificates in the subject DN's "userCertificate" attribute.
- Subject non-null, basicConstraints >= -1
 Looks for certificates in the forward element of the subject DN's
 "crossCertificatePair" attribute AND in the subject's "caCertificate" attribute.
- Issuer non-null, basicConstraints >= -1
 Looks for certificates in the reverse element of the issuer DN's
 "crossCertificatePair" attribute AND in the issuer DN's "caCertificate" attribute.

In each case, certificates are checked using X509CertSelector.match() before adding them to the resulting collection.

If none of the conditions specified above applies, then an exception is thrown to indicate that it was impossible to fetch certificates using the criteria supplied. Note that even if one or more of the conditions apply, the Collection returned may still be empty if there are no certificates in the directory.

The implementation fetches CRLs from the issuer DNs specified in the setCertificateChecking, addIssuerName or setIssuerNames methods of the x509CRLSelector class. If no issuer DNs have been specified using one of these methods, the implementation throws an exception indicating it was impossible to fetch CRLs using the criteria supplied. Otherwise, the CRLs are searched as follows:

 The implementation first creates a list of issuer names. If a certificate was specified in the setCertificateChecking method, it uses the issuer of that



- certificate. Otherwise, it uses the issuer names specified using the addIssuerName Or setIssuerNames methods.
- 2. Next, the implementation iterates through the list of issuer names. For each issuer name, it searches first in the issuer's "authorityRevocationList" attribute and then, if no matching CRL was found there, in the issuer's "certificateRevocationList" attribute. One exception to the above is that if the issuer name was obtained from the certificate specified in the setCertificateChecking method, it only checks the issuer's "authorityRevocationList" attribute if the specified certificate is a CA certificate.
- 3. All CRLs are checked using X509CRLSelector.match() before adding them to the resulting collection.
- 4. If no CRLs satisfying the selection criteria can be found, an empty Collection is returned.

Caching

By default each LDAP CertStore instance caches lookups for a maximum of 30 seconds. The cache lifetime can be changed by setting the system property sun.security.certpath.ldap.cache.lifetime to a value in seconds. A value of 0 disables the cache completely. A value of -1 means unlimited lifetime.

Appendix E: Disabling Cryptographic Algorithms

The jdk.certpath.disabledAlgorithms Security Property contains a list of cryptographic algorithms and key size constraints that are considered weak or broken. Certificates and other data (CRLs, OCSPResponses) containing any of these algorithms or key sizes will be blocked during certification path building and validation. This property is used by Oracle's PKIX implementation, other implementations might not examine and use it.

The exact syntax of the jdk.certpath.disabledAlgorithms property is described in the java.security file. In Java SE 9, the default value of the property is:

```
jdk.certpath.disabledAlgorithms=MD2, MD5, SHA1 jdkCA & usage TLSServer, \
    RSA keySize < 1024, DSA keySize < 1024, EC keySize < 224</pre>
```

In this syntax:

MD2

Any MD2-based algorithm will be blocked.

For example, a certificate, CRL, or OCSPResponse signed with an MD2withRSA signature algorithm.

MD5

Any MD5-based algorithm will be blocked.

For example, a certificate, CRL, or OCSPResponse signed with an MD5withRSA signature algorithm.

SHA1 jdkCA & usage TLSServer

All SHA1 certificates that chain to trust anchors pre-installed in the cacerts keystore and that are used for authentication of TLS Servers. See JEP 288.

RSA keySize < 1024

Any RSA key less than 1024 bits will be blocked.



For example, a certificate with a 768-bit RSA public key.

DSA keySize < 1024

Any DSA key less than 1024 bits will be blocked. For example, a certificate with a 512-bit DSA public key.

EC keySize < 224

Any EC key less than 224 bits will be blocked. For example, a certificate with a 160-bit EC public key.

Administrators or users can modify the value of the jdk.certpath.disabledAlgorithms property to address additional security requirements. However, removing any of the current algorithms or key sizes is not recommended.



10

Java SASL API Programming and Deployment Guide

Simple Authentication and Security Layer (SASL) protocol specifies the authentication and optional establishment of a security layer between client and server applications.

SASL, is an Internet standard (RFC 2222) that specifies a protocol for authentication and optional establishment of a security layer between client and server applications. SASL defines how authentication data is to be exchanged but does not itself specify the contents of that data. It is a framework into which specific *authentication mechanisms* that specify the contents and semantics of the authentication data can fit.

SASL is used by protocols, such as the Lightweight Directory Access Protocol, version 3 (LDAP v3), and the Internet Message Access Protocol, version 4 (IMAP v4) to enable pluggable authentication. Instead of hardwiring an authentication method into the protocol, LDAP v3 and IMAP v4 use SASL to perform authentication, thus enabling authentication via various SASL mechanisms.

There are a number of standard SASL mechanisms defined by the Internet community for various levels of security and deployment scenarios. These range from no security (e.g., anonymous authentication) to high security (e.g., Kerberos authentication) and levels in between.

The Java SASL API

The Java SASL API defines classes and interfaces for applications that use SASL mechanisms. It is defined to be mechanism-neutral: the application that uses the API need not be hardwired into using any particular SASL mechanism. The API supports both client and server applications. It allows applications to select the mechanism to use based on desired security features, such as whether they are susceptible to passive dictionary attacks or whether they accept anonymous authentication.

The Java SASL API also allows developers to use their own, custom SASL mechanisms. SASL mechanisms are installed by using the Java Cryptography Architecture (JCA) Reference Guide (JCA).

When to Use SASL

SASL provides pluggable authentication and security layer for network applications. There are other features in the Java SE that provide similar functionality, including the Java Secure Socket Extension (JSSE) Reference Guide and Java Generic Security Service.

Java GSS is the Java language bindings for the Generic Security Service Application Programming Interface GSS-API.

The only mechanism currently supported underneath this API on Java SE is Kerberos v5.

When compared with JSSE and Java GSS, SASL is relatively lightweight and is popular among more recent protocols. It also has the advantage that several popular,

lightweight (in terms infrastructure support) SASL mechanisms have been defined. Primary JSSE and Java GSS mechanisms, on the other hand, have relatively heavyweight mechanisms that require more elaborate infrastructures (Public Key Infrastructure and Kerberos, respectively).

SASL, JSSE, and Java GSS are often used together. For example, a common pattern is for an application to use JSSE for establishing a secure channel, and to use SASL for client, username/password-based authentication. There are also SASL mechanisms layered on top of GSS-API mechanisms; one popular example is a SASL GSS-API/Kerberos v5 mechanism that is used with LDAP.

Except when defining and building protocols from scratch, often the biggest factor determining which API to use is the protocol definition. For example, LDAP and IMAP are defined to use SASL, so software related to these protocols should use the Java SASL API. When building Kerberos applications and services, the API to use is Java GSS. When building applications and services that use SSL/TLS as their protocol, the API to use is JSSE.

Java SASL API Overview

How SASL Mechanisms are Installed and Selected

The SunSASL Provider

Implementing a SASL Security Provider

Java SASL API Overview

The Java SASL API has interfaces SaslClient and SaslServer that represent the client-side and server-side APIs.

SASL is a challenge-response protocol. The server issues a challenge to the client, and the client sends a response based on the challenge. This exchange continues until the server is satisfied and issues no further challenge. These challenges and responses are binary tokens of arbitrary length. The encapsulating protocol (such as LDAP or IMAP) specifies how these tokens are encoded and exchanged. For example, LDAP specifies how SASL tokens are encapsulated within LDAP bind requests and responses.

The Java SASL API is modeled according to this style of interaction and usage. It has interfaces, <code>SaslClient</code> and <code>SaslServer</code>, that represent client-side and server-side mechanisms, respectively. The application interacts with the mechanisms via byte arrays that represent the challenges and responses. The server-side mechanism iterates, issuing challenges and processing responses, until it is satisfied, while the client-side mechanism iterates, evaluating challenges and issuing responses, until the server is satisfied. The application that is using the mechanism drives each iteration. That is, it extracts the challenge or response from a protocol packet and supplies it to the mechanism, and then puts the response or challenge returned by the mechanism into a protocol packet and sends it to the peer.

Creating the Mechanisms

Passing Input to the Mechanisms

Using the Mechanisms

Using the Negotiated Security Layer



Creating the Mechanisms

The client and server code that uses the SASL mechanisms are not hardwired to use specific mechanism(s). In many protocols that use SASL, the server advertises (either statically or dynamically) a list of SASL mechanisms that it supports. The client then selects one of these based on its security requirements.

The Sas1 class is used for creating instances of Sas1Client and Sas1Server. Here is an example of how an application creates a SASL client mechanism using a list of possible SASL mechanisms.

Based on the availability of the mechanisms supported by the platform and other configuration information provided via the parameters, the Java SASL framework selects one of the listed mechanisms and return an instance of SaslClient.

The name of the selected mechanism is usually transmitted to the server via the application protocol. Upon receiving the mechanism name, the server creates a corresponding SaslServer object to process client-sent responses. Here is an example of how the server would create an instance of SaslServer.

Passing Input to the Mechanisms

The API provides three ways by which an application gives input to a mechanism.

Because the Java SASL API is a general framework, it must be able to accommodate many different types of mechanisms. Each mechanism needs to be initialized with input and may need input to make progress. The API provides three means by which an application gives input to a mechanism:

- 1. Common input parameters The application uses predefined parameters to supply information that are defined by the SASL specification and commonly required by mechanisms. For Saslclient mechanisms, the input parameters are authorization id, protocol id, and server name. For SaslServer mechanisms, the common input parameters are prototol id and (its own fully qualified) server name.
- 2. Properties parameter The application uses the properties parameter, a mapping of property names to (possibly non-string) property values, to supply configuration information. The Java SASL API defines some standard properties, such as Sasl.QOP, Sasl.STRENGTH, and Sasl.MAX_BUFFER. The parameter can also be used to pass in non-standard properties that are specific to particular mechanisms.
- 3. Callbacks The application uses the Interface CallbackHandler parameter to supply input that cannot be predetermined or might not be common across mechanisms. When a mechanism requires input data, it uses the callback handler supplied by the application to collect the data, possibly from the end-user of the application. For example, a mechanism might require the end-user of the application to supply a name and password.



Mechanisms can use the callbacks defined in the <code>javax.security.auth.callback</code> package; these are generic callbacks useful for building applications that perform authentication. Mechanisms might also need SASL-specific callbacks, such as those for collecting realm and authorization information, or even (nonstandardized) mechanism-specific callbacks. The application should be able to accommodate a variety of mechanisms. Consequently, its callback handler must be able to service all of the callbacks that the mechanisms might request. This is not possible in general for arbitrary mechanisms, but is usually feasible due to the limited number of mechanisms that are typically deployed and used.

Using the Mechanisms

After the application has created a mechanism, it uses the mechanism to obtain SASL tokens to exchange with the peer.

Some protocols allows the client to accompany the request with an optional initial response for mechanisms that have an initial response. This feature can be used to lower the number of message exchanges required for authentication. Example 10-1 illustrates how a client might use SaslClient for authentication.

The client application iterates through each step of the authentication by using the mechanism (sc) to evaluate the challenge gotten from the server and to get a response to send back to the server. It continues this cycle until either the mechanism or application-level protocol indicates that the authentication has completed, or if the mechanism cannot evaluate a challenge. If the mechanism cannot evaluate the challenge, it throws an exception to indicate the error and terminates the authentication. Disagreement between the mechanism and protocol about the completion state must be treated as an error because it might indicate a compromise of the authentication exchange.

Example 10-2 illustrates how a server might use SaslServer.

The server application iterates through each step of the authentication by giving the client's response to the mechanism (ss) to process. If the response is incorrect, the mechanism indicates the error by throwing a Saslexception so that the server can report the error and terminate the authentication. If the response is correct, the mechanism returns challenge data to be sent to the client and indicates whether the authentication is complete. Note that challenge data can accompany a "success" indication. This might be used, for example, to tell the client to finalize some negotiated state.

Example 10-1 Sample Code for Using SASL Client for Authentication

```
// Get optional initial response
byte[] response =
        (sc.hasInitialResponse() ? sc.evaluateChallenge(new byte[]) : null);

String mechanism = sc.getMechanismName();

// Send selected mechanism name and optional initial response to server send(mechanism, response);

// Read response
msg = receive();
while (!sc.isComplete() && (msg.status == CONTINUE || msg.status == SUCCESS)) {
            // Evaluate server challenge
            response = sc.evaluateChallenge(msg.contents);
```



Example 10-2 Sample Code for Using SASL Server for Authentication

```
// Read request that contains mechanism name and optional initial response
msg.receive();
// Obtain a SaslServer to perform authentication
SaslServer ss = Sasl.createSaslServer(msg.mechanism,
   protocol, myName, props, callbackHandler);
// Perform authentication steps until done
while (!ss.isComplete()) {
    try {
        // Process response
       byte[] challenge = sc.evaluateResponse(msg.contents);
        if (ss.isComplete()) {
            send(mechanism, challenge, SUCCESS);
            send(mechanism, challenge, CONTINUE);
            msq.receive();
    } catch (SaslException e) {
        send(ERROR);
        sc.dispose();
       break;
```

Using the Negotiated Security Layer

Some SASL mechanisms support only authentication while others support use of a negotiated security layer after authentication. The security layer feature is often not used when the application uses some other means, such as SSL/TLS, to communicate securely with the peer.

When a security layer has been negotiated, all subsequent communication with the peer must take place using the security layer. To determine whether a security layer has been negotiated, get the negotiated Sasl.QOP from the mechanism. Here is an example of how to determine whether a security layer has been negotiated.

A security layer has been negotiated if the Sasl.QOP property indicates that either integrity and/or confidentiality has been negotiated.

To communicate with the peer using the negotiated layer, the application first uses the wrap method to encode the data to be sent to the peer to produce a "wrapped" buffer. It then transfers a length field representing the number of octets in the wrapped buffer followed by the contents of the wrapped buffer to the peer. The peer receiving the stream of octets passes the buffer (without the length field) to unwrap to obtain the decoded bytes sent by the peer. Details of this protocol are described in RFC 2222. Example 10-3 illustrates how a client application sends and receives application data using a security layer.

Example 10-3 Sample Code for SASL Client Send and Receive Data

```
// Send outgoing application data to peer
byte[] outgoing = ...;
byte[] netOut = sc.wrap(outgoing, 0, outgoing.length);
send(netOut.length, netOut); // send to peer

// Receive incoming application data from peer
byte[] netIn = receive(); // read length and ensuing bytes from peer
byte[] incoming = sc.unwrap(netIn, 0, netIn.length);
```

How SASL Mechanisms are Installed and Selected

SASL mechanism implementations are provided by SASL security providers. Each provider may support one or more SASL mechanisms and is registered with the JCA.

By default, in J2SE 5, the SunSASL provider is automatically registered as a JCA provider. To remove it or reorder its priority as a JCA provider, change the line

```
security.provider.7=com.sun.security.sasl.Provider
```

in the Java security properties file (java-home/conf/security/java.security).

To add or remove a SASL provider, you add or remove the corresponding line in the security properties file. For example, if you want to add a SASL provider and have its mechanisms be chosen over the same ones implemented by the SunSASL provider, then you would add a line to the security properties file with a lower number.

```
security.provider.7=com.example.MyProvider
security.provider.8=com.sun.security.sasl.Provider
```

Alternatively, you can programmatically add your own provider using the <code>java.security.Security</code> class. For example, the following sample code registers the <code>com.example.MyProvider</code> to the list of available SASL security providers.

```
Security.addProvider(new com.example.MyProvider());
```

When an application requests a SASL mechanism by supplying one or more mechanism names, the SASL framework looks for registered SASL providers that support that mechanism by going through, in order, the list of registered providers. The providers must then determine whether the requested mechanism matches the selection policy properties in the Sas1 and if so, return an implementation for the mechanism.



The selection policy properties specify the security aspects of a mechanism, such as its susceptibility to certain attacks. These are characteristics of the mechanism (definition), rather than its implementation so all providers should come to the same conclusion about a particular mechanism. For example, the PLAIN mechanism is susceptible to plaintext attacks regardless of how it is implemented. If no selection policy properties are supplied, there are no restrictions on the selected mechanism. Using these properties, an application can ensure that it does not use unsuitable mechanisms that might be deployed in the execution environment. For example, an application might use the following sample code if it does not want to allow the use of mechanisms susceptible to plaintext attacks.

```
Map props = new HashMap();
props.add(Sasl.POLICY_NOPLAINTEXT, "true");
SaslClient sc = Sasl.createSaslClient(mechanisms,
    authzid, protocol, serverName, props, callbackHandler);
```

The SunSASL Provider

Information about the SunSASL provider client and server mechanisms.

The SunSASL provider supports the following client and server mechanisms.

- Client Mechanisms
 - PLAIN (RFC 2595). This mechanism supports cleartext username/password authentication.
 - CRAM-MD5 (RFC 2195). This mechanism supports a hashed username/ password authentication scheme.
 - DIGEST-MD5 (RFC 2831). This mechanism defines how HTTP Digest Authentication can be used as a SASL mechanism.
 - EXTERNAL (RFC 2222). This mechanism obtains authentication information from an external channel (such as TLS or IPsec).
- Server Mechanisms
 - CRAM-MD5
 - DIGEST-MD5

The SunSASL Provider Client Mechanisms

The SunSASL provider supports several SASL client mechanisms used in popular protocols such as LDAP, IMAP, and SMTP.

The following table summarizes the client mechanisms and their required input.

Table 10-1 SunSASL Provider Client Mechanisms

Client Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy
CRAM-MD5 authorization id (as default username)	•	PasswordCallback	None	Sasl.POLICY_NOANON
	NameCallback		YMOUS	
				Sasl.POLICY_NOPLAI



Table 10-1 (Cont.) SunSASL Provider Client Mechanisms

Client Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy
DIGEST-MD5	authorization id protocol id	NameCallback	Sasl.QOP	Sasl.POLICY_NOANON
		PasswordCallback	Sasl.STRENGTH	YMOUS
	server name	RealmCallback	Sasl.MAX_BUFFER	Sasl.POLICY_NOPLAI
		RealmChoiceCallbac	Sasl.SERVER_AUTH	NTEXT
		k	"javax.security.sasl.s	
		RealmChoiceCallbac	endmaxbuffer"	
		k	"com.sun.security.sas I.digest.cipher"	
EXTERNAL	authorization id	None	None	Sasl.POLICY_NOPLAI
	external channel			NTEXT
				Sasl.POLICY_NOACTI VE
				Sasl.POLICY_NODICT IONARY
PLAIN	authorization id	NameCallback	None	Sasl.POLICY_NOANON
		PasswordCallback		YMOUS

An application that uses these mechanisms from the SunSASL provider must supply the required parameters, callbacks and properties. The properties have reasonable defaults and only need to be set if the application wants to override the defaults. Most of the parameters, callbacks, and properties are described in the API documentation. The following sections describe mechanism-specific behaviors and parameters not already covered by the API documentation.

Cram-MD5

The Cram-MD5 client mechanism uses the authorization id parameter, if supplied, as the default username in the NameCallback to solicit the application/end-user for the authentication id. The authorization id is otherwise not used by the Cram-MD5 mechanism; only the authentication id is exchanged with the server.

Digest-MD5

The Digest-MD5 mechanism is used for digest authentication and optional establishment of a security layer. It specifies the following ciphers for use with the security layer: Triple DES, DES and RC4 (128, 56, and 40 bits). The Digest-MD5 mechanism can support only ciphers that are available on the platform. For example, if the platform does not support the RC4 ciphers, then the Digest-MD5 mechanism will not use those ciphers.

The Sasl.Strength property supports "high", "medium", and "low" settings; its default is "high,medium,low". The ciphers are mapped to the strength settings as follows:



Table 10-2 Cipher Strength

Strength	Cipher	Cipher Id	
high	Triple DES RC4 128 bits	3des rc4	
medium	DES RC4 56 bits	des rc4-56	
low	RC4 40 bits	rc4-40	

When there is more than one choice for a particular strength, the cipher selected depends on the availability of the ciphers in the underlying platform. To explicitly name the cipher to use, set the "com.sun.security.sasl.digest.cipher" property to the corresponding cipher id. Note that this property setting must be compatible with Sasl.STRENGTH and the ciphers available in the underlying platform. For example, Sasl.STRENGTH being set to "low" and "com.sun.security.sasl.digest.cipher" being set to "3des" are incompatible. The "com.sun.security.sasl.digest.cipher" property has no default.

The "javax.security.sasl.sendmaxbuffer" property specifies (the string representation of) the maximum send buffer size in bytes. The default is 65536. The actual maximum number of bytes will be the minimum of this number and the peer's maximum receive buffer size.

The SunSASL Provider Server Mechanisms

The SunSASL provider supports several SASL server mechanisms used in popular protocols such as LDAP, IMAP, and SMTP.

The following table summarizes the server mechanisms and the required input:

Table 10-3 Server Mechanisms

Server Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy
CRAM-MD5	server name	AuthorizeCallback NameCallback	None	Sasl.POLICY_NOANON YMOUS
		PasswordCallback		Sasl.POLICY_NOPLAI NTEXT
DIGEST-MD5	protocol id server name	AuthorizeCallback	Sasl.QOP	Sasl.POLICY_NOANON
		NameCallback	Sasl.STRENGTH	YMOUS
		PasswordCallback	Sasl.MAX_BUFFER	Sasl.POLICY_NOPLAI
	RealmCallback	"javax.security.sa sl.sendmaxbuffer"	NTEXT	
			"com.sun.security. sasl.digest.realm"	
			"com.sun.security. sasl.digest.utf8"	

An application that uses these mechanisms from the SunSASL provider must supply the required parameters, callbacks and properties. The properties have reasonable defaults and only need to be set if the application wants to override the defaults.

All users of server mechanisms must have a callback handler that deals with the AuthorizeCallback. This is used by the mechanisms to determine whether the authenticated user is allowed to act on behalf of the requested authorization id, and also to obtain the canonicalized name of the authorized user (if canonicalization is applicable).

Most of the parameters, callbacks, and properties are described in the API documentation. The following sections describe mechanism-specific behaviors and parameters not already covered by the API documentation.

Cram-MD5

The Cram-MD5 server mechanism uses the NameCallback and PasswordCallback to obtain the password required to verify the SASL client's response. The callback handler should use the NameCallback.getDefaultName() as the key to fetch the password.

Digest-MD5

The Digest-MD5 server mechanism uses the RealmCallback, NameCallback, and PasswordCallback to obtain the password required to verify the SASL client's response. The callback handler should use RealmCallback.getDefaultText() and NameCallback.getDefaultName() as keys to fetch the password.

The "javax.security.sasl.sendmaxbuffer" property specifies (the string representation of) the maximum send buffer size in bytes. The default is 65536. The actual maximum number of bytes will be the minimum of this number and the peer's maximum receive buffer size.

The "com.sun.security.sasl.digest.realm" property is used to specify a list of space-separated realm names that the server supports. The list is sent to the client as part of the challenge. If this property has not been set, the default realm is the server's name (supplied as a parameter).

The "com.sun.security.sasl.digest.utf8" property is used to specify the character encoding to use. "true" means to use UTF-8 encoding; "false" means to use ISO Latin 1 (ISO-8859-1). The default is "true".

Debugging and Monitoring

The SunSASL and JdkSASL providers uses the Logging APIs to provide implementation logging output. This output can be controlled by using the logging configuration file and programmatic API (java.util.logging)

The logger name used by the SunSASL provider is "javax.security.sasl"

Here is a sample logging configuration file that enables the ${\tt FINEST}$ logging level for the SunSASL provider:

javax.security.sasl.level=FINEST
handlers=java.util.logging.ConsoleHandler
java.util.logging.ConsoleHandler.level=FINEST

The table below shows the mechanisms and the logging output that they generate:



Table 10-4 Logging Output

Mechanism	Logging Level	Information Logged	
CRAM-MD5	FINE	Configuration properties; challenge/ response messages	
DIGEST-MD5	INFO	Message discarded due to encoding problem (e.g., unmatched MACs, incorrect padding)	
DIGEST-MD5	FINE	Configuration properties; challenge/ response messages	
DIGEST-MD5	FINER	More detailed information about challenge/response messages	
DIGEST-MD5	FINEST	Buffers exchanged at the security layer	
GSSAPI	FINE	Configuration properties; challenge, response messages	
GSSAPI	FINER	More detailed information about challenge/response messages	
GSSAPI	FINEST	Buffers exchanged at the security layer	

The JdkSASL Provider

Information about the JdkSASL provider client and server mechanisms.

The JdkSASL provider supports the following client and server mechanisms.

- Client Mechanisms
 - GSSAPI (RFC 2222). This mechanism uses the GSSAPI for obtaining authentication information. It supports Kerberos v5 authentication.
- Server Mechanisms
 - GSSAPI (Kerberos v5)

The JdkSASL Provider Client Mechanism

The JdkSASL provider supports the GSSAPI client mechanism used in popular protocols such as LDAP, IMAP, and SMTP.

The following table summarizes the GSSAPI client mechanism and its required input.



Table 10-5 JdkSASL Provider Client Mechanism

Client Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy
GSSAPI JAAS Subject authorization id protocol id server name	None	Sasl.QOP Sasl.MAX_BUFFER Sasl.SERVER_AUTH	Sasl.POLICY_NOACTI VE Sasl.POLICY_NOANON	
		"javax.security.sasl.s endmaxbuffer"	YMOUS Sasl.POLICY_NOPLAI NTEXT	

An application that uses the GSSAPI mechanism from the JdkSASL provider must supply the required parameters, callbacks and properties. The properties have reasonable defaults and only need to be set if the application wants to override the defaults. Most of the parameters, callbacks, and properties are described in the API documentation. The following section describes further GSSAPI behaviors and parameters not already covered by the API documentation.

GSSAPI

The GSSAPI mechanism is used for Kerberos v5 authentication and optional establishment of a security layer. The mechanism expects the calling thread's <code>subject</code> to contain the client's Kerberos credentials or that the credentials could be obtained by implicitly logging in to Kerberos. To obtain the client's Kerberos credentials, use the <code>Java Authentication</code> and <code>Authorization Service (JAAS)</code> to log in using the Kerberos login module. See the <code>JAAS</code> and <code>Java GSS-API Tutorial</code> for details and examples. After using <code>JAAS</code> authentication to obtain the Kerberos credentials, you put the code that uses the <code>SASL GSSAPI</code> mechanism within <code>doAs</code> or <code>doAsPrivileged</code>.

To obtain Kerberos credentials without doing explicit JAAS programming, see the JAAS and Java GSS-API Tutorial . When using this approach, there is no need to wrap the code within doAs or doAsPrivileged

The "javax.security.sasl.sendmaxbuffer" property specifies (the string representation of) the maximum send buffer size in bytes. The default is 65536. The actual maximum number of bytes will be the minimum of this number and the peer's maximum receive buffer size.



The JdkSASL Provider Server Mechanism

The JdkSASL provider supports the GSSAPI mechanism used in popular protocols such as LDAP, IMAP, and SMTP.

The following table summarizes the GSSAPI server mechanism and the required input:

Table 10-6 Server mechanism

Server Mechanism Name	Parameters/Input	Callbacks	Configuration Properties	Selection Policy
GSSAPI	Subject protocol id server name	AuthorizeCallback	Sasl.QOP Sasl.MAX_BUFFER "javax.security.sa sl.sendmaxbuffer"	Sasl.POLICY_NOACTI VE Sasl.POLICY_NOANON YMOUS
				Sasl.POLICY_NOPLAI NTEXT

An application that uses the GSSAPI mechanism from the JdkSASL provider must supply the required parameters, callbacks and properties. The properties have reasonable defaults and only need to be set if the application wants to override the defaults.

All users of server mechanism must have a callback handler that deals with the <code>AuthorizeCallback</code>. This is used by the mechanism to determine whether the authenticated user is allowed to act on behalf of the requested authorization id, and also to obtain the canonicalized name of the authorized user (if canonicalization is applicable).

Most of the parameters, callbacks, and properties are described in the API documentation. The following section describes GSSAPI mechanism-specific behaviors and parameters not already covered by the API documentation.

GSSAPI

The GSSAPI mechanism is used for Kerberos v5 authentication and optional establishment of a security layer. The mechanism expects the calling thread's <code>Subject</code> to contain the client's Kerberos credentials or that the credentials could be obtained by implicitly logging in to Kerberos. To obtain the client's Kerberos credentials, use the <code>Java Authentication</code> and <code>Authorization Service</code> (<code>JAAS</code>) to log in using the Kerberos login module. See the <code>JAAS</code> and <code>Java GSS-API Tutorial</code> for details and examples. After using <code>JAAS</code> authentication to obtain the Kerberos credentials, you put the code that uses the <code>SASL GSSAPI</code> mechanism within <code>doAs</code> or <code>doAsPrivileged</code>.

```
LoginContext lc = new LoginContext("JaasSample", new TextCallbackHandler());
lc.login();
lc.getSubject().doAs(new SaslAction());

class SaslAction implements java.security.PrivilegedAction {
   public class run() {
        ...
        String[] mechanisms = new String[]{"GSSAPI"};
        SaslClient sc = Sasl.createSaslClient(mechanisms,
```



```
authzid, protocol, serverName, props, callbackHandler);
...
}
```

To obtain Kerberos credentials without doing explicit JAAS programming, see the JAAS and Java GSS-API Tutorial . When using this approach, there is no need to wrap the code within $doAs\ Or\ doAsPrivileged$

The "javax.security.sasl.sendmaxbuffer" property specifies (the string representation of) the maximum send buffer size in bytes. The default is 65536. The actual maximum number of bytes will be the minimum of this number and the peer's maximum receive buffer size.

Implementing a SASL Security Provider

Procedure to implement a SASL Security Provider.

There are three basic steps in implementing a SASL security provider:

- 1. Write a class that implements the SaslClient or SaslServer interface.
 - This involves providing an implementation for the SASL mechanism. To implement a client mechanism, you need to implement the methods declared in the SaslClient interface. Similarly, for a server mechanism, you need to implement the methods declared in the SaslServer interface. For the purposes of this discussion, suppose you are developing an implementation for the client mechanism "SAMPLE-MECH", implemented by the class, com.example.SampleMechClient. You must decide what input are needed by the mechanism and how the implementation is going to collect them. For example, if the mechanism is username/password-based, then the implementation would likely need to collect that information via the callback handler parameter.
- 2. Write a factory class (that implements SaslClientFactory Or SaslServerFactory) that creates instances of the class.
 - This involves providing a factory class that will create instances of <code>com.example.SampleMechClient</code>. The factory needs to determine the characteristics of the mechanism that it supports (as described by the <code>Sasl.POLICY_*</code> properties) so that it can return an instance of the mechanism when the API user requests it using compatible policy properties. The factory may also check for validity of the parameters before creating the mechanism. For the purposes of this discussion, suppose the factory class is named <code>com.example.MySampleClientFactory</code>. Although our sample factory is responsible for only one mechanism, a single factory can be responsible for any number of mechanisms.
- 3. Write a JCA provider that registers the factory.

This involves creating a JCA provider. The steps for creating a JCA provider is described in detail in the document, Steps to Implement and Integrate a Provider. SASL client factories are registered using property names of the form SaslClientFactory.mechName while SASL server factories are registered using property names of the form SaslServerFactory.mechName

mechName is the SASL mechanism's name. This is what's returned by SaslClient.getMechanismName() and SaslServer.getMechanismName(). Continuing with our example, here is how the provider would register the "SAMPLE-MECH" mechanism.



 $\verb"put("SaslClientFactory.SAMPLE-MECH", "com.example.MySampleClientFactory");\\$

A single SASL provider might be responsible for many mechanisms. Therefore, it might have many invocations of put to register the relevant factories. The completed SASL provider can then be made available to applications using the instructions given above.



11

XML Digital Signature

Java XML Digital Signature API Specification describes a standard API for generating and validating XML signatures. It also describes how to create a concrete implementation of the XML Digital Signature API and register it as a cryptographic service of a Java Cryptography Architecture (JCA) provider.

XML Digital Signature API Overview and Tutorial demonstrates how to validate and generate an XML signature with the API.

Java XML Digital Signature API Specification

This document describes the Java XML Digital Signature API Specification (JSR 105). The purpose of this JSR is to define a standard Java API for generating and validating XML signatures.

When this specification is final, there will be a Reference Implementation which will demonstrate the capabilities of this API and will provide an operational definition of this specification. A Technology Compatibility Kit (TCK) will also be available that will verify whether an implementation of the specification is compliant. These are required as per the Java Community Process 2.1.

The JSR 105 API is intended to target the following two types of users:

- Java programmers who want to use the JSR 105 API to generate and validate XML signatures.
- Java programmers who want to create a concrete implementation of the JSR 105 API and register it as a cryptographic service of a JCA provider (see The Provider Class).

Acknowledgements

The JSR 105 Expert Group:

- Nicolas Catania, Hewlett-Packard
- Donald E. Eastlake 3rd, Motorola
- Christian Geuer-Pollmann, Apache Software Foundation
- Hans Granqvist, VeriSign
- Kazuyuki Harada, Fujitsu
- Anthony Ho, DSTC
- Merlin Hughes, Baltimore Technologies
- Joyce Leung, IBM
- Gregor Karlinger, IAIK



- Serge Mister, Entrust Technologies
- Takuya Mori, NEC Corporation
- Sean Mullan, Sun Microsystems (co-specification lead)
- Anthony Nadalin, IBM (co-specification lead)
- Erwin van der Koogh, Apache Software Foundation
- Chris Yeung, XML Asia

Also, special thanks to: Valerie Peng, Vincent Ryan, Sharon Liu, Chok Poh, K. Venugopal Rao., Paul Rank, Alexey Gavrilov, Bill Situ, Eric Jendrock, Andrew Fan, Manveen Kaur, Tom Amiro, Michael Mi, Dmitri Silaev, Roman Makarchuk, Vanitha Venkatraman, Arkadiy Sutchilin, and Scott Fordin from Sun Microsystems, Vishal Mahajan from Apache, and Martin Centner from IAIK.

Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

- 1. W3C Recommendation, XML-Signature Syntax and Processing.
 - The API MUST allow a programmer to generate and validate XML Signatures such that all of the SHOULD and MUST requirements specified by the W3C recommendation can be satisfied.
 - The API MUST allow an implementation of the API to be created such that all
 of the SHOULD and MUST requirements specified by the W3C
 recommendation can be satisfied.
- 2. An implementation SHOULD support the W3C Recommendation, XML-Signature XPath Filter Transform 2.0.
- 3. An implementation SHOULD support the W3C Recommendation, Exclusive XML Canonicalization Version 1.0.
- 4. DOM-independent API. The API MUST NOT have dependencies on a specific XML representation, such as DOM. It MUST be possible to create implementations of the API for different XML processing and mechanism representations, such as DOM, JDOM or dom4j.
- 5. Extensible, provider-based API. It MUST be possible for a third-party to create and plug in an implementation responsible for managing and creating cryptographic and transform algorithms, dereferencing URIs, and marshalling objects to/from XML.
- 6. Support for a default XML mechanism type: DOM. An implementation MUST minimally support the default mechanism type: DOM. This ensures that all implementations of JSR 105 are guaranteed a minimal level of functionality. Implementations MAY support other mechanism types.
- Interoperability for the default XML mechanism type: DOM. The API SHOULD ensure that applications using a DOM implementation are portable and interoperable.
- **8. J2SE requirements.** Implementations of this technology MAY support J2SE 1.2 or later but MUST at a minimum support version 1.4 or later of J2SE.



API Dependencies

- J2SE (JDK) 1.2 or higher
- W3C DOM Level 2 API. This dependency is required by classes of the javax.xml.crypto.dom and javax.xml.crypto.dsig.dom packages.

Non-Goals

- Support for non-DOM implementations. While the API SHOULD allow non-DOM implementations to be created, it is beyond the scope of the first version to ensure interoperability between implementations other than DOM. Additional standard service provider types MAY be added in the future and necessary API enhancements MAY be considered for a maintenance revision of JSR 105.
- 2. Support for a higher-level API. We expect that programmers MAY design high-level APIs which will be built on the JSR 105 API to hide low-level details, address common use-cases or apply profiling constraints. However, it is beyond the scope of the first version to support these requirements. A high-level API MAY be considered for a maintenance release of JSR 105.
- 3. Support for user-pluggable algorithms (other than transform and canonicalization algorithms which is supported by the javax.xml.crypto.dsig.TransformService class): Allowing developers to plug in their own implementations of XML Signature algorithms without requiring them to create a complete JSR 105 implementation seems like a worthy goal but SHALL NOT be REQUIRED for this release of JSR 105. A solution we are investigating for a subsequent release of Java SE is to enhance the underlying JCA/JCE to add better support for registering, parsing and processing XML security algorithms, parameters, and key information.

Package Overview

The JSR 105 API consists of 6 packages:

- javax.xml.crypto
- javax.xml.crypto.dom
- javax.xml.crypto.dsig
- javax.xml.crypto.dsig.dom
- javax.xml.crypto.dsig.keyinfo
- javax.xml.crypto.dsig.spec

The <code>javax.xml.crypto</code> package contains common classes that are used to perform XML cryptographic operations, such as generating an XML signature or encrypting XML data. Two notable classes in this package are the <code>KeySelector</code> class, the purpose of which is to allow developers to supply implementations which locate and optionally validate keys using the information contained in a <code>KeyInfo</code> object, and the <code>URIDereferencer</code> class which allows developers to create and specify their own URI dereferencing implementations.



The <code>javax.xml.crypto.dsig</code> package includes interfaces that represent the core elements defined in the W3C XML digital signature specification. Of primary significance is the <code>XMLSignature</code> class, which allows you to sign and validate an XML digital signature. Most of the XML signature structures or elements are represented by a corresponding interface (except for the <code>KeyInfo</code> structures, which are included in their own package, and discussed in the next paragraph). These interfaces include: <code>SignedInfo</code>, <code>CanonicalizationMethod</code>, <code>SignatureMethod</code>, <code>Reference</code>, <code>Transform</code>, <code>DigestMethod</code>, <code>XMLObject</code>, <code>Manifest</code>, <code>SignatureProperty</code>, and <code>SignatureProperties</code>. The <code>XMLSignatureFactory</code> class is an abstract factory that is used to create objects that implement these interfaces.

The <code>javax.xml.crypto.dsig.keyinfo</code> package contains interfaces that represent most of the <code>KeyInfo</code> structures defined in the W3C XML digital signature recommendation, including <code>KeyInfo</code>, <code>KeyName</code>, <code>KeyValue</code>, <code>X509Data</code>, <code>X509IssuerSerial</code>, <code>RetrievalMethod</code>, and <code>PGPData</code>. The <code>KeyInfoFactory</code> class is an abstract factory that is used to create objects that implement these interfaces.

The <code>javax.xml.crypto.dsig.spec</code> package contains interfaces and classes representing input parameters for the digest, signature, transform, or canonicalization algorithms used in the processing of XML signatures.

Finally, the <code>javax.xml.crypto.dom</code> and <code>javax.xml.crypto.dsig.dom</code> packages contains DOM-specific classes for the <code>javax.xml.crypto</code> and <code>javax.xml.crypto.dsig</code> packages, respectively. Only developers and users who are creating or using a DOM-based <code>XMLSignatureFactory</code> or <code>KeyInfoFactory</code> implementation should need to make direct use of these packages.

Service Providers

A JSR 105 cryptographic service is a concrete implementation of the abstract XMLSignatureFactory and KeyInfoFactory classes and is responsible for creating objects and algorithms that parse, generate and validate XML Signatures and KeyInfo structures. A concrete implementation of XMLSignatureFactory MUST provide support for each of the REQUIRED algorithms as specified by the W3C recommendation for XML Signatures. It MAY support other algorithms as defined by the W3C recommendation or other specifications.

JSR 105 leverages the JCA provider model (see The Provider Class) for registering and loading XMLSignatureFactory and KeyInfoFactory implementations.

Each concrete XMLSignatureFactory or KeyInfoFactory implementation supports a specific XML mechanism type that identifies the XML processing mechanism that an implementation uses internally to parse and generate XML signature and KeyInfo structures. This JSR supports one standard type: **DOM**. Support for new standard types (such as JDOM) MAY be added in the future.

A JSR 105 implementation SHOULD use underlying JCA engine classes, such as java.security.Signature and java.security.MessageDigest to perform cryptographic operations.

In addition to the XMLSignatureFactory and KeyInfoFactory classes, JSR 105 supports a service provider interface for transform and canonicalization algorithms. The TransformService class allows you to develop and plug in an implementation of a specific transform or canonicalization algorithm for a particular XML mechanism



type. The TransformService class uses the standard JCA provider model for registering and loading implementations. Each JSR 105 implementation SHOULD use the TransformService class to find a provider that supports transform and canonicalization algorithms in XML Signatures that it is generating or validating.

DOM Mechanism Requirements

The following requirements MUST be abided by when implementing a DOM-based XMLSignatureFactory, KeyInfoFactory Or TransformService in order to minimize interoperability problems:

- The unmarshalXMLSignature method of XMLSignatureFactory MUST support DOMValidateContext types. If the type is DOMValidateContext, it SHOULD contain an Element of type Signature. Additionally, the unmarshalXMLSignature method MAY populate the Id/Element mappings of the passed-in DOMValidateContext.
- 2. The sign method of XMLSignatures produced by XMLSignatureFactory MUST support DOMSignContext types and the validate method MUST support DOMValidateContext types. This requirement also applies to the validate method of SignatureValue and the validate method of Reference.
- 3. The implementation MUST support DOMStructures as the mechanism for the application to specify extensible content (any elements or mixed content).
- 4. If the dereference method of user-specified URIDereferencers returns NodeSetData objects, the iterator method MUST return an iteration over objects of type org.w3c.dom.Node.
- 5. URIReference objects passed to the dereference method of user-specified URIDereferencers MUST be of type DOMURIReference and XMLCryptoContext objects MUST implement DOMCryptoContext.
- 6. The previous 2 requirements also apply to URIDereferencers returned by the getURIDereferencer method of XMLSignatureFactory and KeyInfoFactory.
- 7. The unmarshalKeyInfo method of KeyInfoFactory MUST support DOMStructure types. If the type is DOMStructure, it SHOULD contain an Element of type KeyInfo.
- **8.** The transform method of Transform MUST support DOMCryptoContext context parameter types.
- **9.** The newtransform and newCanonicalizationMethod methods of XMLSignatureFactory MUST support DOMStructure parameter types.
- 10. The init, and marshalParams methods of TransformService MUST support DOMStructure and DOMCryptoContext types.
- 11. The unmarshalXMLSignature method of XMLSignatureFactory MUST support DOMStructure types. If the type is DOMStructure, it SHOULD contain an Element of type Signature.
- **12.** The marshal method of KeyInfo MUST support DOMStructure and DOMCryptoContext parameter types.



Note that a DOM implementation MAY internally use other XML parsing APIs other than DOM as long as it doesn't affect interoperability. For example, a DOM implementation of XMLSignatureFactory might use a SAX parser internally to canonicalize data.

Open API Issues

The following is a list of open API issues.

 ID attribute registration of external XML document references is not supported. Consider the following reference:

Dereferencing the external document results in an octet stream which is subsequently converted to a NodeSet by the JSR 105 implementation. But the API does not provide a mechanism for registering ID attributes of external documents and therefore the XPath Transform implementation may be unable to identify the "foo" ID.

Programming Examples

Examples 1-3 below demonstrate how to generate different types of simple XML Digital Signature using the JSR 105 API. Example 1 describes how to generate a detached signature using the DSA signature algorithm. Example 2 describes how to generate an enveloped signature. Example 3 describes how to generate an enveloping signature. Example 4 describes how to validate an XML Signature.

Example 11-1 1. Generating a detached XML Digital Signature

```
import javax.xml.crypto.*;
import javax.xml.crypto.dsig.*;
import javax.xml.crypto.dom.*;
import javax.xml.crypto.dsig.dom.DOMSignContext;
import javax.xml.crypto.dsig.keyinfo.*;
import javax.xml.crypto.dsig.spec.C14NMethodParameterSpec;
import java.io.FileOutputStream;
import java.io.OutputStream;
import java.security.*;
import java.util.Collections;
import javax.xml.parsers.DocumentBuilderFactory;
import javax.xml.transform.*;
import javax.xml.transform.dom.DOMSource;
import javax.xml.transform.stream.StreamResult;
import org.w3c.dom.Document;
 ^{\star} This is a simple example of generating a Detached XML
 ^{\star} Signature using the JSR 105 API. The resulting signature will look
 * like (key and signature values will be different):
 * <code>
```



```
* <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
       <CanonicalizationMethod Algorithm="http://www.w3.org/TR/2001/REC-xml-</pre>
c14n-20010315"/>
       <SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#dsa-sha256"/>
       <Reference URI="http://www.w3.org/TR/xml-stylesheet">
         <DigestMethod Algorithm="http://www.w3.org/2001/04/xmlenc#sha256"/>
         <DigestValue>60NvZvtdTB+7UnlLp/H24p7h4bs=</DigestValue>
       </Reference>
     </SignedInfo>
     <SignatureValue>
       DpEylhQoiUKBoKWmYfajXO7LZxiDYqVtUtCNyTqwZqoChzorA2nhkQ==
     </SignatureValue>
     <KeyInfo>
       <KevValue>
         <DSAKeyValue>
           <P>
             rFto8uPQM6y34FLPmDh40BLJ1rVrC8VeRquuhPZ6jYNFkQuwxnu/wCvIAMhukPBL
             FET8bJf/b2ef+oqxZajEb+88zlZoyG8g/wMfDBHTxz+CnowLahnCCTYBp5kt7G8q
             UobJuvjylwj1st7V9Lsu03iXMXtbiriUjFa5gURasN8=
           </P>
           <0>
             kEjAFpCe4lcUOdwphpzf+tBaUds=
           </Q>
           <G>
             oe14R2OtyKx+s+6005BRNMOYpIg2TU/f15N3bsDErKOWtKXeNK9FS7dWStreDxo2
             SSgOonqAd4FuJ/4uva7GgNL4ULIqY7E+mW5iwJ7n/WTELh98mEocsLXkNh24HcH4
             BZfSCTruuzmCyjdV1KSqX/Eux04HfCWYmdxN3SQ/qqw=
           </G>
           < V >
             pA5NnZvcd574WRXuOA7ZfC/7Lqt4cB0MRLWtHubtJoVOao9ib5ry4rTk0r6ddnOv
             AIGKktutzK3ymvKleS3DOrwZQgJ+/BDWDW8kO9R66o6rdjiSobBi/0c2V1+dkqOg
             jFmKz395mvCOZGhC7fqAVhHat2EjGPMfgSZyABa7+1k=
           </Y>
         </DSAKeyValue>
       </KeyValue>
     </KeyInfo>
 * </Signature>
 * </code>
 * /
public class GenDetached {
    // Synopsis: java GenDetached [output]
    // where output is the name of the file that will contain the detached
    // signature. If not specified, standard output is used.
    public static void main(String[] args) throws Exception {
        // First, create a DOM XMLSignatureFactory that will be used to
        // generate the XMLSignature and marshal it to DOM.
        XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");
        // Create a Reference to an external URI that will be digested
        // using the SHA256 digest algorithm
        Reference ref = fac.newReference("http://www.w3.org/TR/xml-stylesheet",
            fac.newDigestMethod(DigestMethod.SHA256, null));
        // Create the SignedInfo
        SignedInfo si = fac.newSignedInfo(
```

```
(CanonicalizationMethod.INCLUSIVE_WITH_COMMENTS,
                 (C14NMethodParameterSpec) null),
            fac.newSignatureMethod("http://www.w3.org/2000/09/xmldsig#dsa-sha256",
null),
            Collections.singletonList(ref));
        // Create a DSA KeyPair
        KeyPairGenerator kpg = KeyPairGenerator.getInstance("DSA");
        kpg.initialize(2048);
        KeyPair kp = kpg.generateKeyPair();
        // Create a KeyValue containing the DSA PublicKey that was generated
        KeyInfoFactory kif = fac.getKeyInfoFactory();
        KeyValue kv = kif.newKeyValue(kp.getPublic());
        // Create a KeyInfo and add the KeyValue to it
        KeyInfo ki = kif.newKeyInfo(Collections.singletonList(kv));
        // Create the XMLSignature (but don't sign it yet)
        XMLSignature signature = fac.newXMLSignature(si, ki);
        // Create the Document that will hold the resulting XMLSignature
        DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
        dbf.setNamespaceAware(true); // must be set
        Document doc = dbf.newDocumentBuilder().newDocument();
        // Create a DOMSignContext and set the signing Key to the DSA
        // PrivateKey and specify where the XMLSignature should be inserted
        // in the target document (in this case, the document root)
        DOMSignContext signContext = new DOMSignContext(kp.getPrivate(), doc);
        // Marshal, generate (and sign) the detached XMLSignature. The DOM
        // Document will contain the XML Signature if this method returns
        // successfully.
        signature.sign(signContext);
        // output the resulting document
        OutputStream os;
        if (args.length > 0) {
           os = new FileOutputStream(args[0]);
        } else {
           os = System.out;
        TransformerFactory tf = TransformerFactory.newInstance();
        Transformer trans = tf.newTransformer();
        trans.transform(new DOMSource(doc), new StreamResult(os));
}
```

fac.newCanonicalizationMethod

Example 11-2 2. Generating an enveloped XML Digital Signature

```
import javax.xml.crypto.*;
import javax.xml.crypto.dsig.*;
import javax.xml.crypto.dom.*;
import javax.xml.crypto.dsig.dom.DOMSignContext;
import javax.xml.crypto.dsig.keyinfo.*;
import javax.xml.crypto.dsig.spec.*;
import java.io.FileInputStream;
import java.io.FileOutputStream;
```



```
import java.io.OutputStream;
import java.security.*;
import java.util.Collections;
import java.util.Iterator;
import javax.xml.parsers.DocumentBuilderFactory;
import javax.xml.transform.*;
import javax.xml.transform.dom.DOMSource;
import javax.xml.transform.stream.StreamResult;
import org.w3c.dom.Document;
 * This is a simple example of generating an Enveloped XML
 * Signature using the JSR 105 API. The resulting signature will look
 * like (key and signature values will be different):
 * <code>
 *<Envelope xmlns="urn:envelope">
 * <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
     <SignedInfo>
       <CanonicalizationMethod Algorithm="http://www.w3.org/TR/2001/REC-xml-</pre>
c14n-20010315"/>
       <SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#dsa-sha256"/>
       <Reference URI="">
         <Transforms>
           <Transform Algorithm="http://www.w3.org/2000/09/xmldsig#enveloped-</pre>
signature"/>
         <DigestMethod Algorithm="http://www.w3.org/2001/04/xmlenc#sha256"/>
         <DigestValue>K8M/lPbKnuMDsO0Uzuj75lQtzQI=<DigestValue>
       </Reference>
     </SignedInfo>
     <SignatureValue>
       DpEylhQoiUKBoKWmYfajXO7LZxiDYgVtUtCNyTgwZgoChzorA2nhkQ==
     </SignatureValue>
     <KeyInfo>
       <KeyValue>
         <DSAKeyValue>
           <P>
             rFto8uPQM6y34FLPmDh40BLJ1rVrC8VeRquuhPZ6jYNFkQuwxnu/wCvIAMhukPBL
             FET8bJf/b2ef+oqxZajEb+88zlZoyG8g/wMfDBHTxz+CnowLahnCCTYBp5kt7G8q
             UobJuvjylwj1st7V9Lsu03iXMXtbiriUjFa5gURasN8=
           </P>
           <Q>
             kEjAFpCe4lcUOdwphpzf+tBaUds=
           </0>
           <G>
             oe14R2OtyKx+s+6005BRNMOYpIq2TU/f15N3bsDErKOWtKXeNK9FS7dWStreDxo2
             SSqOongAd4FuJ/4uva7GqNL4ULIqY7E+mW5iwJ7n/WTELh98mEocsLXkNh24HcH4
             BZfSCTruuzmCyjdV1KSqX/Eux04HfCWYmdxN3SQ/qqw=
           </G>
           < Y >
             pA5NnZvcd574WRXu0A7ZfC/7Lqt4cB0MRLWtHubtJoVOao9ib5ry4rTk0r6ddnOv
             AIGKktutzK3ymvKleS3DOrwZQgJ+/BDWDW8kO9R66o6rdjiSobBi/0c2V1+dkqOg
             jFmKz395mvCOZGhC7fqAVhHat2EjGPMfgSZyABa7+1k=
           </Y>
         </DSAKeyValue>
       </KeyValue>
     </KeyInfo>
 * </Signature>
 *</Envelope>
 * </code>
```

```
public class GenEnveloped {
    // Synopsis: java GenEnveloped [document] [output]
    //
    //
          where "document" is the name of a file containing the XML document
    //
          to be signed, and "output" is the name of the file to store the
    //
          signed document. The 2nd argument is optional - if not specified,
    //
          standard output will be used.
    public static void main(String[] args) throws Exception {
        // Create a DOM XMLSignatureFactory that will be used to generate the
        // enveloped signature
        XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");
        // Create a Reference to the enveloped document (in this case we are
        // signing the whole document, so a URI of "" signifies that) and
        // also specify the SHA256 digest algorithm and the ENVELOPED Transform.
        Reference ref = fac.newReference
            ("", fac.newDigestMethod(DigestMethod.SHA256, null),
             Collections.singletonList
              (fac.newTransform
                (Transform.ENVELOPED, (TransformParameterSpec) null)),
             null, null);
        // Create the SignedInfo
        SignedInfo si = fac.newSignedInfo
            (fac.newCanonicalizationMethod
             (CanonicalizationMethod.INCLUSIVE_WITH_COMMENTS,
              (C14NMethodParameterSpec) null),
             fac.newSignatureMethod("http://www.w3.org/2000/09/xmldsig#dsa-sha256",
null),
             Collections.singletonList(ref));
        // Create a DSA KeyPair
        KeyPairGenerator kpg = KeyPairGenerator.getInstance("DSA");
        kpg.initialize(2048);
        KeyPair kp = kpg.generateKeyPair();
        // Create a KeyValue containing the DSA PublicKey that was generated
        KeyInfoFactory kif = fac.getKeyInfoFactory();
        KeyValue kv = kif.newKeyValue(kp.getPublic());
        // Create a KeyInfo and add the KeyValue to it
        KeyInfo ki = kif.newKeyInfo(Collections.singletonList(kv));
        // Instantiate the document to be signed
        DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
        dbf.setNamespaceAware(true);
        Document doc =
            dbf.newDocumentBuilder().parse(new FileInputStream(args[0]));
        // Create a DOMSignContext and specify the DSA PrivateKey and
        // location of the resulting XMLSignature's parent element
        DOMSignContext dsc = new DOMSignContext
            (kp.getPrivate(), doc.getDocumentElement());
        // Create the XMLSignature (but don't sign it yet)
        XMLSignature signature = fac.newXMLSignature(si, ki);
```



```
// Marshal, generate (and sign) the enveloped signature
signature.sign(dsc);

// output the resulting document
OutputStream os;
if (args.length > 1) {
    os = new FileOutputStream(args[1]);
} else {
    os = System.out;
}

TransformerFactory tf = TransformerFactory.newInstance();
Transformer trans = tf.newTransformer();
trans.transform(new DOMSource(doc), new StreamResult(os));
}
```

Example 11-3 3. Generating an enveloping XML Digital Signature

```
import javax.xml.crypto.*;
import javax.xml.crypto.dsig.*;
import javax.xml.crypto.dom.*;
import javax.xml.crypto.dsig.dom.DOMSignContext;
import javax.xml.crypto.dsig.keyinfo.*;
import javax.xml.crypto.dsig.spec.C14NMethodParameterSpec;
import java.io.FileInputStream;
import java.io.FileOutputStream;
import java.io.OutputStream;
import java.security.*;
import java.util.Arrays;
import java.util.Collections;
import javax.xml.parsers.DocumentBuilderFactory;
import javax.xml.transform.*;
import javax.xml.transform.dom.DOMSource;
import javax.xml.transform.stream.StreamResult;
import org.w3c.dom.Document;
import org.w3c.dom.Node;
* This is a simple example of generating an Enveloping XML
 * Signature using the JSR 105 API. The signature in this case references a
 * local URI that points to an Object element.
 * The resulting signature will look like (certificate and
 * signature values will be different):
 * <code>
 * <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
     <SignedInfo>
       <CanonicalizationMethod Algorithm="http://www.w3.org/TR/2001/REC-xml-</pre>
c14n-20010315#WithComments"/>
       <SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#dsa-sha256"/>
       <Reference URI="#object">
         <DigestMethod Algorithm="http://www.w3.org/2001/04/xmlenc#sha256"/>
         <DigestValue>7/XTsHaBSOnJ/jXD5v0zL6VKYsk=</DigestValue>
      </Reference>
     </SignedInfo>
     <SignatureValue>
      RpMRbtMHLa0siSS+BwUpLIEmTfh/0fsld2JYQWZzCzfa5kBTz25+XA==
     </SignatureValue>
```



```
<KeyInfo>
       <KeyValue>
         <DSAKeyValue>
             /KaCzo4Syrom78z3EQ5SbbB4sF7ey80etKII864WF64B81uRpH5t9jQTxeEu0Imbz
             RMqzVDZkVG9xD7nN1kuFw==
           </P>
           <0>
             li7dzDacuo67Jg7mtqEm2TRuOMU=
           </Q>
           <G>
             Z4Rxsnqc9E7pGknFFH2xqaryRPBaQ01khpMdLRQnG541Awtx/XPaF5Bpsy4pNWMOH
             CBiNU0NogpsQW5QvnlMpA==
           </G>
             wbEUaCgHZXqK4qLvbdYrAc6+Do0XVcsziCJqxzn4cJJRxwc3E1xnEXHscVgr1Cq19
             i5fanOKQbFXzmb+bChqiq==
           </Y>
         </DSAKeyValue>
       </KeyValue>
     </KeyInfo>
     <Object Id="object">some text</Object>
 * </Signature>
 * </code>
public class GenEnveloping {
    // Synopis: java GenEnveloping [output]
         where "output" is the name of a file that will contain the
         generated signature. If not specified, standard ouput will be used.
    public static void main(String[] args) throws Exception {
        // First, create the DOM XMLSignatureFactory that will be used to
        // generate the XMLSignature
        XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");
        // Next, create a Reference to a same-document URI that is an Object
        // element and specify the SHA256 digest algorithm
        Reference ref = fac.newReference("#object",
            fac.newDigestMethod(DigestMethod.SHA256, null));
        // Next, create the referenced Object
        DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
        dbf.setNamespaceAware(true);
        Document doc = dbf.newDocumentBuilder().newDocument();
        Node text = doc.createTextNode("some text");
        XMLStructure content = new DOMStructure(text);
        XMLObject obj = fac.newXMLObject
            (Collections.singletonList(content), "object", null, null);
        // Create the SignedInfo
        SignedInfo si = fac.newSignedInfo(
            fac.newCanonicalizationMethod
                (CanonicalizationMethod.INCLUSIVE_WITH_COMMENTS,
                 (C14NMethodParameterSpec) null),
            fac.newSignatureMethod("http://www.w3.org/2000/09/xmldsig#dsa-sha256",
null),
```

```
Collections.singletonList(ref));
// Create a DSA KeyPair
KeyPairGenerator kpg = KeyPairGenerator.getInstance("DSA");
kpg.initialize(2048);
KeyPair kp = kpg.generateKeyPair();
// Create a KeyValue containing the DSA PublicKey that was generated
KeyInfoFactory kif = fac.getKeyInfoFactory();
KeyValue kv = kif.newKeyValue(kp.getPublic());
// Create a KeyInfo and add the KeyValue to it
KeyInfo ki = kif.newKeyInfo(Collections.singletonList(kv));
// Create the XMLSignature (but don't sign it yet)
XMLSignature signature = fac.newXMLSignature(si, ki,
    Collections.singletonList(obj), null, null);
// Create a DOMSignContext and specify the DSA PrivateKey for signing
// and the document location of the XMLSignature
DOMSignContext dsc = new DOMSignContext(kp.getPrivate(), doc);
// Lastly, generate the enveloping signature using the PrivateKey
signature.sign(dsc);
// output the resulting document
OutputStream os;
if (args.length > 0) {
   os = new FileOutputStream(args[0]);
} else {
   os = System.out;
TransformerFactory tf = TransformerFactory.newInstance();
Transformer trans = tf.newTransformer();
trans.transform(new DOMSource(doc), new StreamResult(os));
```

Example 11-4 4. Validating an XML Digital Signature

```
import javax.xml.crypto.*;
import javax.xml.crypto.dsig.*;
import javax.xml.crypto.dom.*;
import javax.xml.crypto.dsig.dom.DOMValidateContext;
import javax.xml.crypto.dsig.keyinfo.*;
import java.io.FileInputStream;
import java.security.*;
import java.util.Collections;
import java.util.Iterator;
import java.util.List;
import javax.xml.parsers.DocumentBuilderFactory;
import org.w3c.dom.Document;
import org.w3c.dom.NodeList;
 \mbox{\scriptsize *} This is a simple example of validating an XML
 * Signature using the JSR 105 API. It assumes the key needed to
 * validate the signature is contained in a KeyValue KeyInfo.
public class Validate {
```



```
// Synopsis: java Validate [document]
//
//
      where "document" is the name of a file containing the XML document
//
      to be validated.
//
public static void main(String[] args) throws Exception {
    // Instantiate the document to be validated
    DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
    dbf.setNamespaceAware(true);
    Document doc =
        dbf.newDocumentBuilder().parse(new FileInputStream(args[0]));
    // Find Signature element
    NodeList nl =
        doc.getElementsByTagNameNS(XMLSignature.XMLNS, "Signature");
    if (nl.getLength() == 0) {
        throw new Exception("Cannot find Signature element");
    // Create a DOM XMLSignatureFactory that will be used to unmarshal the
    // document containing the XMLSignature
    XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");
    // Create a DOMValidateContext and specify a KeyValue KeySelector
    // and document context
    DOMValidateContext valContext = new DOMValidateContext
        (new KeyValueKeySelector(), nl.item(0));
    // unmarshal the XMLSignature
    XMLSignature signature = fac.unmarshalXMLSignature(valContext);
    // Validate the XMLSignature (generated above)
    boolean coreValidity = signature.validate(valContext);
    // Check core validation status
    if (coreValidity == false) {
        System.err.println("Signature failed core validation");
       boolean sv = signature.getSignatureValue().validate(valContext);
       System.out.println("signature validation status: " + sv);
        // check the validation status of each Reference
        Iterator i = signature.getSignedInfo().getReferences().iterator();
        for (int j=0; i.hasNext(); j++) {
            boolean refValid =
                ((Reference) i.next()).validate(valContext);
            System.out.println("ref["+j+"] validity status: " + refValid);
        }
    } else {
        System.out.println("Signature passed core validation");
}
 * KeySelector which retrieves the public key out of the
 * KeyValue element and returns it.
 * NOTE: If the key algorithm doesn't match signature algorithm,
 * then the public key will be ignored.
private static class KeyValueKeySelector extends KeySelector {
```

```
public KeySelectorResult select(KeyInfo keyInfo,
                                        KeySelector.Purpose purpose,
                                        AlgorithmMethod method,
                                        XMLCryptoContext context)
            throws KeySelectorException {
            if (keyInfo == null) {
                throw new KeySelectorException("Null KeyInfo object!");
            SignatureMethod sm = (SignatureMethod) method;
            List list = keyInfo.getContent();
            for (int i = 0; i < list.size(); i++) {
                XMLStructure xmlStructure = (XMLStructure) list.get(i);
                if (xmlStructure instanceof KeyValue) {
                    PublicKey pk = null;
                        pk = ((KeyValue)xmlStructure).getPublicKey();
                    } catch (KeyException ke) {
                        throw new KeySelectorException(ke);
                    // make sure algorithm is compatible with method
                    if (algEquals(sm.getAlgorithm(), pk.getAlgorithm())) {
                        return new SimpleKeySelectorResult(pk);
            throw new KeySelectorException("No KeyValue element found!");
        }
        //@@@FIXME: this should also work for key types other than DSA/RSA
        static boolean algEquals(String algURI, String algName) {
            if (algName.equalsIgnoreCase("DSA") &&
                algURI.equalsIgnoreCase("http://www.w3.org/2000/09/xmldsig#dsa-
sha256")) {
                return true;
            } else if (algName.equalsIgnoreCase("RSA") &&
                       algURI.equalsIgnoreCase("http://www.w3.org/2000/09/
xmldsig#dsa-sha256")) {
                return true;
            } else {
                return false;
    private static class SimpleKeySelectorResult implements KeySelectorResult {
       private PublicKey pk;
        SimpleKeySelectorResult(PublicKey pk) {
            this.pk = pk;
       public Key getKey() { return pk; }
```

XML Digital Signature API Overview and Tutorial

The Java XML Digital Signature API is a standard Java API for generating and validating XML Signatures. This API was defined under the Java Community Process

as JSR 105. This JSR is final and this release of Java SE contains an FCS access implementation of the Final version of the APIs.

XML Signatures can be applied to data of any type, XML or binary (see XML Signature Syntax and Processing). The resulting signature is represented in XML. An XML Signature can be used to secure your data and provide data integrity, message authentication, and signer authentication.

After providing a brief overview of XML Signatures and the XML Digital Signature API, this document presents two examples that demonstrate how to use the API to validate and generate an XML Signature. This document assumes that you have a basic knowledge of cryptography and digital signatures.

The API is designed to support all of the required or recommended features of the W3C Recommendation for XML-Signature Syntax and Processing. The API is extensible and pluggable and is based on the Java Cryptography Service Provider Architecture. The API is designed for two types of developers:

- Developers who want to use the XML Digital Signature API to generate and validate XML signatures
- Developers who want to create a concrete implementation of the XML Digital Signature API and register it as a cryptographic service of a JCA provider (see The Provider Class).

Package Hierarchy

The six packages listed below comprise the XML Digital Signature API:

- javax.xml.crypto
- javax.xml.crypto.dsig
- javax.xml.crypto.dsig.keyinfo
- javax.xml.crypto.dsig.spec
- javax.xml.crypto.dom
- javax.xml.crypto.dsig.dom

The <code>javax.xml.crypto</code> package contains common classes that are used to perform XML cryptographic operations, such as generating an XML signature or encrypting XML data. Two notable classes in this package are the <code>KeySelector</code> class, which allows developers to supply implementations that locate and optionally validate keys using the information contained in a <code>KeyInfo</code> object, and the <code>URIDereferencer</code> class, which allows developers to create and specify their own URI dereferencing implementations.

The <code>javax.xml.crypto.dsig</code> package includes interfaces that represent the core elements defined in the W3C XML digital signature specification. Of primary significance is the <code>XMLSignature</code> class, which allows you to sign and validate an XML digital signature. Most of the XML signature structures or elements are represented by a corresponding interface (except for the KeyInfo structures, which are included in their own package and are discussed in the next paragraph). These interfaces include: <code>SignedInfo</code>, <code>CanonicalizationMethod</code>, <code>SignatureMethod</code>,



Reference, Transform, DigestMethod, XMLObject, Manifest, SignatureProperty, and SignatureProperties. The XMLSignatureFactory class is an abstract factory that is used to create objects that implement these interfaces.

The <code>javax.xml.crypto.dsig.keyinfo</code> package contains interfaces that represent most of the <code>KeyInfo</code> structures defined in the W3C XML digital signature recommendation, including <code>KeyInfo</code>, <code>KeyName</code>, <code>KeyValue</code>, <code>X509Data</code>, <code>X509IssuerSerial</code>, <code>RetrievalMethod</code>, and <code>PGPData</code>. The <code>KeyInfoFactory</code> class is an abstract factory that is used to create objects that implement these interfaces.

The <code>javax.xml.crypto.dsig.spec</code> package contains interfaces and classes representing input parameters for the digest, signature, transform, or canonicalization algorithms used in the processing of XML signatures.

Finally, the <code>javax.xml.crypto.dom</code> and <code>javax.xml.crypto.dsig.dom</code> packages contains DOM-specific classes for the <code>javax.xml.crypto</code> and <code>javax.xml.crypto.dsig</code> packages, respectively. Only developers and users who are creating or using a DOM-based <code>XMLSignatureFactory</code> or <code>KeyInfoFactory</code> implementation will need to make direct use of these packages.

Service Providers

A JSR 105 cryptographic service is a concrete implementation of the abstract XMLSignatureFactory and KeyInfoFactory classes and is responsible for creating objects and algorithms that parse, generate and validate XML Signatures and KeyInfo structures. A concrete implementation of XMLSignatureFactory must provide support for each of the required algorithms as specified by the W3C recommendation for XML Signatures. It can optionally support other algorithms as defined by the W3C recommendation or other specifications.

JSR 105 leverages the JCA provider model for registering and loading XMLSignatureFactory and KeyInfoFactory implementations.

Each concrete XMLSignatureFactory or KeyInfoFactory implementation supports a specific XML mechanism type that identifies the XML processing mechanism that an implementation uses internally to parse and generate XML signature and KeyInfo structures. This JSR supports one standard type, DOM. The XML Digital Signature provider implementation that is bundled with Java SE supports the DOM mechanism. Support for new standard types, such as JDOM, may be added in the future.

An XML Digital Signature API implementation should use underlying JCA engine classes, such as java.security.Signature and java.security.MessageDigest, to perform cryptographic operations.

In addition to the XMLSignatureFactory and KeyInfoFactory classes, JSR 105 supports a service provider interface for transform and canonicalization algorithms. The TransformService class allows you to develop and plug in an implementation of a specific transform or canonicalization algorithm for a particular XML mechanism type. The TransformService class uses the standard JCA provider model for registering and loading implementations. Each JSR 105 implementation should use the TransformService class to find a provider that supports transform and canonicalization algorithms in XML Signatures that it is generating or validating.



Introduction to XML Signatures

You can use an XML Signature to sign any arbitrary data, whether it is XML or binary. The data is identified via URIs in one or more Reference elements. XML Signatures are described in one or more of three forms: detached, enveloping, or enveloped. A detached signature is over data that is external, or outside of the signature element itself. Enveloping signatures are signatures over data that is inside the signature element, and an enveloped signature is a signature that is contained inside the data that it is signing.

Example of an XML Signature

The easiest way to describe the contents of an XML Signature is to show an actual sample and describe each component in more detail. The following is an example of an enveloped XML Signature generated over the contents of an XML document. The contents of the document before it is signed are:

```
<Envelope xmlns="urn:envelope"> </Envelope>
```

The resulting enveloped XML Signature, indented and formatted for readability, is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<Envelope xmlns="urn:envelope">
  <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
    <SignedInfo>
      <CanonicalizationMethod
        Algorithm="http://www.w3.org/TR/2001/REC-xml-c14n-20010315#WithComments"/>
      <SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#dsa-sha1"/>
      <Reference URT="">
        <Transforms>
          <Transform Algorithm="http://www.w3.org/2000/09/xmldsig#enveloped-</pre>
signature"/>
        <DigestMethod Algorithm="http://www.w3.org/2000/09/xmldsig#sha1"/>
        <DigestValue>uooqbWYa5VCqcJCbuymBKqm17vY=</DigestValue>
      </Reference>
    </SignedInfo>
    <SignatureValue>
      KedJuTob5gtvYx9qM3k3gm7kbLBwVbEQRl26S2tmXjqNND7MRGtoew==
    </SignatureValue>
    <KeyInfo>
      <KeyValue>
        <DSAKeyValue>
          <P>
            /KaCzo4Syrom78z3EQ5SbbB4sF7ey80etKII864WF64B81uRpH5t9jQTxe
            Eu0ImbzRMqzVDZkVG9xD7nN1kuFw==
          </P>
          <Q>li7dzDacuo67Jg7mtqEm2TRuOMU=</Q>
            Z4Rxsnqc9E7pGknFFH2xqaryRPBaQ01khpMdLRQnG541Awtx/
            XPaF5Bpsy4pNWMOHCBiNU0NogpsQW5QvnlMpA==
          </G>
          <Y>
```



The Signature element has been inserted inside the content that it is signing, thereby making it an enveloped signature. The required SignedInfo element contains the information that is actually signed:

The required <code>canonicalizationMethod</code> element defines the algorithm used to canonicalize the <code>signedInfo</code> element before it is signed or validated. Canonicalization is the process of converting XML content to a canonical form, to take into account changes that can invalidate a signature over that data. Canonicalization is necessary due to the nature of XML and the way it is parsed by different processors and intermediaries, which can change the data such that the signature is no longer valid but the signed data is still logically equivalent.

The required SignatureMethod element defines the digital signature algorithm used to generate the signature, in this case DSA with SHA-1.

One or more Reference elements identify the data that is digested. Each Reference element identifies the data via a URI. In this example, the value of the URI is the empty String (""), which indicates the root of the document. The optional Transforms element contains a list of one or more Transform elements, each of which describes a transformation algorithm used to transform the data before it is digested. In this example, there is one Transform element for the enveloped transform algorithm. The enveloped transform is required for enveloped signatures so that the signature element itself is removed before calculating the signature value. The required DigestMethod element defines the algorithm used to digest the data, in this case SHA1. Finally the required DigestValue element contains the actual base64-encoded digested value.

The required SignatureValue element contains the base64-encoded signature value of the signature over the SignedInfo element.

The optional KeyInfo element contains information about the key that is needed to validate the signature:

```
<KeyInfo>
<KeyValue>
<DSAKeyValue>
```



This <code>KeyInfo</code> element contains a <code>KeyValue</code> element, which in turn contains a <code>DSAKeyValue</code> element consisting of the public key needed to validate the signature. <code>KeyInfo</code> can contain various content such as X.509 certificates and PGP key identifiers. See The <code>KeyInfo</code> Element in XML Signature Syntax and Processing for more information on the different <code>KeyInfo</code> types.

XML Digital Signature API Examples

The following sections describe two examples that show how to use the XML Digital Signature API:

Validate Example

To compile and run the example, execute the following commands:

```
$ javac Validate.java
$ java Validate signature.xml
```

The sample program will validate the signature in the file signature.xml in the current working directory.

Example 11-5 Validate.java

```
import javax.xml.crypto.*;
import javax.xml.crypto.dsig.*;
import javax.xml.crypto.dom.*;
import javax.xml.crypto.dsig.dom.DOMValidateContext;
import javax.xml.crypto.dsig.keyinfo.*;
import java.io.FileInputStream;
import java.security.*;
import java.util.Collections;
import java.util.Iterator;
import java.util.List;
import javax.xml.parsers.DocumentBuilderFactory;
import org.w3c.dom.Document;
import org.w3c.dom.NodeList;

/**
    * This is a simple example of validating an XML
    * Signature using the JSR 105 API. It assumes the key needed to
```



```
* validate the signature is contained in a KeyValue KeyInfo.
public class Validate {
    //
    // Synopsis: java Validate [document]
    //
    //
          where "document" is the name of a file containing the XML document
    //
          to be validated.
    //
    public static void main(String[] args) throws Exception {
        // Instantiate the document to be validated
        DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
        dbf.setNamespaceAware(true);
        Document doc =
            dbf.newDocumentBuilder().parse(new FileInputStream(args[0]));
        // Find Signature element
       NodeList nl =
            doc.getElementsByTagNameNS(XMLSignature.XMLNS, "Signature");
        if (nl.getLength() == 0) {
            throw new Exception("Cannot find Signature element");
        }
        \//\ Create a DOM XMLSignatureFactory that will be used to unmarshal the
        // document containing the XMLSignature
       XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");
        // Create a DOMValidateContext and specify a KeyValue KeySelector
        // and document context
        DOMValidateContext valContext = new DOMValidateContext
            (new KeyValueKeySelector(), nl.item(0));
        // unmarshal the XMLSignature
        XMLSignature signature = fac.unmarshalXMLSignature(valContext);
        // Validate the XMLSignature (generated above)
       boolean coreValidity = signature.validate(valContext);
        // Check core validation status
        if (coreValidity == false) {
            System.err.println("Signature failed core validation");
            boolean sv = signature.getSignatureValue().validate(valContext);
            System.out.println("signature validation status: " + sv);
            // check the validation status of each Reference
            Iterator i = signature.getSignedInfo().getReferences().iterator();
            for (int j=0; i.hasNext(); j++) {
                boolean refValid =
                    ((Reference) i.next()).validate(valContext);
                System.out.println("ref["+j+"] validity status: " + refValid);
            }
        } else {
            System.out.println("Signature passed core validation");
     * KeySelector which retrieves the public key out of the
     * KeyValue element and returns it.
     * NOTE: If the key algorithm doesn't match signature algorithm,
```

```
* then the public key will be ignored.
    private static class KeyValueKeySelector extends KeySelector {
        public KeySelectorResult select(KeyInfo keyInfo,
                                        KeySelector.Purpose purpose,
                                        AlgorithmMethod method,
                                        XMLCryptoContext context)
            throws KeySelectorException {
            if (keyInfo == null) {
                throw new KeySelectorException("Null KeyInfo object!");
            SignatureMethod sm = (SignatureMethod) method;
            List list = keyInfo.getContent();
            for (int i = 0; i < list.size(); i++) {</pre>
                XMLStructure xmlStructure = (XMLStructure) list.get(i);
                if (xmlStructure instanceof KeyValue) {
                    PublicKey pk = null;
                    try {
                        pk = ((KeyValue)xmlStructure).getPublicKey();
                    } catch (KeyException ke) {
                        throw new KeySelectorException(ke);
                    // make sure algorithm is compatible with method
                    if (algEquals(sm.getAlgorithm(), pk.getAlgorithm())) {
                        return new SimpleKeySelectorResult(pk);
            throw new KeySelectorException("No KeyValue element found!");
        //@@@FIXME: this should also work for key types other than DSA/RSA
        static boolean algEquals(String algURI, String algName) {
            if (algName.equalsIgnoreCase("DSA") &&
                algURI.equalsIgnoreCase("http://www.w3.org/2009/xmldsig11#dsa-
sha256")) {
                return true;
            } else if (algName.equalsIgnoreCase("RSA") &&
                       algURI.equalsIgnoreCase("http://www.w3.org/2001/04/xmldsig-
more#rsa-sha256")) {
                return true;
            } else {
                return false;
    private static class SimpleKeySelectorResult implements KeySelectorResult {
        private PublicKey pk;
        SimpleKeySelectorResult(PublicKey pk) {
            this.pk = pk;
        public Key getKey() { return pk; }
```



Example 11-6 envelope.xml

```
<Envelope xmlns="urn:envelope">
</Envelope>
```

Example 11-7 signature.xml

```
<?xml version="1.0" encoding="UTF-8"?>
<Envelope xmlns="urn:envelope">
    <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
        <SignedInfo xmlns="http://www.w3.org/2000/09/xmldsig#">
             <CanonicalizationMethod
                 Algorithm="http://www.w3.org/TR/2001/REC-xml-c14n-20010315#WithComments"
                xmlns="http://www.w3.org/2000/09/xmldsig#"/>
             <SignatureMethod
                 Algorithm="http://www.w3.org/2009/xmldsig11#dsa-sha256"
                 xmlns="http://www.w3.org/2000/09/xmldsig#"/>
             <Reference URI="" xmlns="http://www.w3.org/2000/09/xmldsig#">
                 <Transforms xmlns="http://www.w3.org/2000/09/xmldsig#">
                     <Transform
                         Algorithm="http://www.w3.org/2000/09/xmldsig#enveloped-signature"
                         xmlns="http://www.w3.org/2000/09/xmldsig#"/>
                 </Transforms>
                 <DigestMethod
                     Algorithm="http://www.w3.org/2001/04/xmlenc#sha256"
                     xmlns="http://www.w3.org/2000/09/xmldsig#"/>
                 <DigestValue xmlns="http://www.w3.org/2000/09/</pre>
xmldsig#">uoogbWYa5VCgcJCbuymBKqm17vY=</DigestValue>
            </Reference>
        </SignedInfo>
        <SignatureValue xmlns="http://www.w3.org/2000/09/</pre>
xmldsig#">e07K1BdC0kzNvr1HpMf4hKoWsv1+oI04nMw55G0+Z5hyI6By30ihow==</SignatureValue>
        <KeyInfo xmlns="http://www.w3.org/2000/09/xmldsig#">
             <KeyValue xmlns="http://www.w3.org/2000/09/xmldsig#">
                 <DSAKeyValue xmlns="http://www.w3.org/2000/09/xmldsig#">
                     <P xmlns="http://www.w3.org/2000/09/xmldsig#">/
KaCzo4Syrom78z3EQ5SbbB4sF7ey80etKII864WF64B81uRpH5t9jQTxeEu0ImbzRMqzVDZkVG9xD7nN1kuFw
==</P>
                     <Q xmlns="http://www.w3.org/2000/09/
xmldsig#">li7dzDacuo67Jg7mtqEm2TRuOMU=</Q>
                     <G xmlns="http://www.w3.org/2000/09/
\verb|xmldsig#"> \verb|Z4Rxsnqc9E7pGknFFH2xqaryRPBaQ01khpMdLRQnG541Awtx/|
XPaF5Bpsy4pNWMOHCBiNU0NogpsQW5QvnlMpA==</G>
                     <Y xmlns="http://www.w3.org/2000/09/
\verb|xm|| dsig#">OqFi0sGpvroi6Ut3m154QNWc6gavH3j2ZoRPDW7qVBbgk7XompuKvZe1owz0yvxq+1K+mWbL7ST| dsigmux d
+t5nr6UFBCq==</Y>
                 </DSAKeyValue>
             </KeyValue>
        </KeyInfo>
    </Signature>
</Envelope>
```

Validating an XML Signature

This example shows you how to validate an XML Signature using the JSR 105 API. The example uses DOM (the Document Object Model) to parse an XML document containing a Signature element and a JSR 105 DOM implementation to validate the signature.

Instantiating the Document that Contains the Signature

First we use a JAXP DocumentBuilderFactory to parse the XML document containing the Signature. An application obtains the default implementation for DocumentBuilderFactory by calling the following line of code:

```
DocumentBuilderFactory dbf =
  DocumentBuilderFactory.newInstance();
```

We must also make the factory namespace-aware:

```
dbf.setNamespaceAware(true);
```

Next, we use the factory to get an instance of a DocumentBuilder, which is used to parse the document:

```
DocumentBuilder builder = dbf.newDocumentBuilder();
Document doc = builder.parse(new FileInputStream(argv[0]));
```

Specifying the Signature Element to be Validated

We need to specify the Signature element that we want to validate, since there could be more than one in the document. We use the DOM method Document.getElementsByTagNameNS, passing it the XML Signature namespace URI and the tag name of the Signature element, as shown:

```
NodeList nl = doc.getElementsByTagNameNS
  (XMLSignature.XMLNS, "Signature");
if (nl.getLength() == 0) {
  throw new Exception("Cannot find Signature element");
}
```

This returns a list of all Signature elements in the document. In this example, there is only one Signature element.

Creating a Validation Context

We create an XMLValidateContext instance containing input parameters for validating the signature. Since we are using DOM, we instantiate a DOMValidateContext instance (a subclass of XMLValidateContext), and pass it two parameters, a KeyValueKeySelector object and a reference to the Signature element to be validated (which is the first entry of the NodeList we generated earlier):

```
DOMValidateContext valContext = new DOMValidateContext
  (new KeyValueKeySelector(), nl.item(0));
```

The KeyValueKeySelector is explained in greater detail in Using KeySelectors.

Unmarshalling the XML Signature

We extract the contents of the Signature element into an XMLSignature object. This process is called unmarshalling. The Signature element is unmarshalled using an

XMLSignatureFactory object. An application can obtain a DOM implementation of XMLSignatureFactory by calling the following line of code:

```
XMLSignatureFactory factory =
   XMLSignatureFactory.getInstance("DOM");
```

We then invoke the unmarshalXMLSignature method of the factory to unmarshal an XMLSignature object, and pass it the validation context we created earlier:

```
XMLSignature signature =
  factory.unmarshalXMLSignature(valContext);
```

Validating the XML Signature

Now we are ready to validate the signature. We do this by invoking the validate method on the XMLSignature object, and pass it the validation context as follows:

```
boolean coreValidity = signature.validate(valContext);
```

The validate method returns "true" if the signature validates successfully according to the core validation rules in the W3C XML Signature Recommendation, and false otherwise.

Using KeySelectors

KeySelectors are used to find and select keys that are needed to validate an XMLSignature. Earlier, when we created a DOMValidateContext object, we passed a KeySelector object as the first argument:

```
DOMValidateContext valContext = new DOMValidateContext
  (new KeyValueKeySelector(), nl.item(0));
```

Alternatively, we could have passed a PublicKey as the first argument if we already knew what key is needed to validate the signature. However, we often don't know.

The KeyValueKeySelector is a concrete implementation of the abstract KeySelector class. The KeyValueKeySelector implementation tries to find an appropriate validation key using the data contained in KeyValue elements of the KeyInfo element of an XMLSignature. It does not determine if the key is trusted. This is a very simple KeySelector implementation, designed for illustration rather than real-world usage. A more practical example of a KeySelector is one that searches a KeyStore for trusted keys that match X509Data information (for example, X509SubjectName, X509IssuerSerial, X509SKI, or X509Certificate elements) contained in a KeyInfo.

The implementation of the KeyValueKeySelector is as follows:

```
private static class KeyValueKeySelector extends KeySelector {
  public KeySelectorResult select(KeyInfo keyInfo,
     KeySelector.Purpose purpose,
     AlgorithmMethod method,
     XMLCryptoContext context)
  throws KeySelectorException {
  if (keyInfo == null) {
    throw new KeySelectorException("Null KeyInfo object!");
}
```



```
SignatureMethod sm = (SignatureMethod) method;
   List list = keyInfo.getContent();
   for (int i = 0; i < list.size(); i++) {</pre>
     XMLStructure xmlStructure = (XMLStructure) list.get(i);
     if (xmlStructure instanceof KeyValue) {
       PublicKey pk = null;
        try {
         pk = ((KeyValue)xmlStructure).getPublicKey();
        } catch (KeyException ke) {
         throw new KeySelectorException(ke);
        // make sure algorithm is compatible with method
       if (algEquals(sm.getAlgorithm(),
           pk.getAlgorithm())) {
         return new SimpleKeySelectorResult(pk);
     }
   throw new KeySelectorException("No KeyValue element
found!");
  static boolean algEquals(String algURI, String algName) {
   if (algName.equalsIgnoreCase("DSA") &&
       algURI.equalsIgnoreCase(SignatureMethod.DSA_SHA1)) {
     return true;
   } else if (algName.equalsIgnoreCase("RSA") &&
        algURI.equalsIgnoreCase(SignatureMethod.RSA_SHA1)) {
     return true;
    } else {
     return false;
```

GenEnveloped Example

To compile and run this sample, execute the following command:

```
$ javac GenEnveloped.java
$ java GenEnveloped envelope.xml envelopedSignature.xml
```

The sample program will generate an enveloped signature of the document in the file envelope.xml and store it in the file envelopedSignature.xml in the current working directory.

Example 11-8 GenEnveloped.java

```
import javax.xml.crypto.*;
import javax.xml.crypto.dsig.*;
import javax.xml.crypto.dom.*;
import javax.xml.crypto.dsig.dom.DOMSignContext;
import javax.xml.crypto.dsig.keyinfo.*;
import javax.xml.crypto.dsig.spec.*;
import java.io.FileInputStream;
import java.io.FileOutputStream;
import java.io.OutputStream;
import java.security.*;
```



```
import java.util.Collections;
import java.util.Iterator;
import javax.xml.parsers.DocumentBuilderFactory;
import javax.xml.transform.*;
import javax.xml.transform.dom.DOMSource;
import javax.xml.transform.stream.StreamResult;
import org.w3c.dom.Document;
 \mbox{\ensuremath{^{\star}}} This is a simple example of generating an Enveloped XML
 * Signature using the JSR 105 API. The resulting signature will look
 * like (key and signature values will be different):
 * <code>
 *<Envelope xmlns="urn:envelope">
 * <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
     <SignedInfo>
       <CanonicalizationMethod Algorithm="http://www.w3.org/TR/2001/REC-xml-
c14n-20010315"/>
       <SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#dsa-sha256"/>
       <Reference URI="">
         <Transforms>
           <Transform Algorithm="http://www.w3.org/2000/09/xmldsig#enveloped-</pre>
signature"/>
         <DigestMethod Algorithm="http://www.w3.org/2001/04/xmlenc#sha256"/>
         <DigestValue>K8M/lPbKnuMDsO0Uzuj75lQtzQI=<DigestValue>
       </Reference>
     </SignedInfo>
     <SignatureValue>
       DpEylhQoiUKBoKWmYfajXO7LZxiDYgVtUtCNyTgwZgoChzorA2nhkQ==
     </SignatureValue>
     <KeyInfo>
       <KeyValue>
         <DSAKeyValue>
           <P>
             rFto8uPQM6y34FLPmDh40BLJ1rVrC8VeRquuhPZ6jYNFkQuwxnu/wCvIAMhukPBL
             FET8bJf/b2ef+ogxZajEb+88zlZoyG8g/wMfDBHTxz+CnowLahnCCTYBp5kt7G8g
             UobJuvjylwj1st7V9Lsu03iXMXtbiriUjFa5gURasN8=
           </P>
           <Q>
             kEjAFpCe4lcUOdwphpzf+tBaUds=
           </0>
           <G>
             oe14R2OtyKx+s+6005BRNMOYpIq2TU/f15N3bsDErKOWtKXeNK9FS7dWStreDxo2
             SSgOonqAd4FuJ/4uva7GgNL4ULIqY7E+mW5iwJ7n/WTELh98mEocsLXkNh24HcH4
             BZfSCTruuzmCyjdV1KSqX/Eux04HfCWYmdxN3SQ/qqw=
           </G>
           <Y>
             \verb"pA5NnZvcd574WRXuOA7ZfC/7Lqt4cB0MRLWtHubtJoVOao9ib5ry4rTk0r6ddnOv"
             AIGKktutzK3ymvKleS3DOrwZQgJ+/BDWDW8kO9R66o6rdjiSobBi/0c2V1+dkqOg
             jFmKz395mvCOZGhC7fqAVhHat2EjGPMfgSZyABa7+1k=
           </Y>
         </DSAKeyValue>
       </KeyValue>
     </KeyInfo>
 * </Signature>
 *</Envelope>
 * </code>
public class GenEnveloped {
```

```
// Synopsis: java GenEnveloped [document] [output]
    //
    //
         where "document" is the name of a file containing the XML document
    //
         to be signed, and "output" is the name of the file to store the
    //
         signed document. The 2nd argument is optional - if not specified,
    //
          standard output will be used.
    //
    public static void main(String[] args) throws Exception {
        // Create a DOM XMLSignatureFactory that will be used to generate the
        // enveloped signature
        XMLSignatureFactory fac = XMLSignatureFactory.getInstance("DOM");
        // Create a Reference to the enveloped document (in this case we are
        // signing the whole document, so a URI of "" signifies that) and
        // also specify the SHA256 digest algorithm and the ENVELOPED Transform.
        Reference ref = fac.newReference
            ("", fac.newDigestMethod(DigestMethod.SHA256, null),
             Collections.singletonList
              (fac.newTransform
                (Transform.ENVELOPED, (TransformParameterSpec) null)),
             null, null);
        // Create the SignedInfo
        SignedInfo si = fac.newSignedInfo
            (fac.newCanonicalizationMethod
             (CanonicalizationMethod.INCLUSIVE_WITH_COMMENTS,
              (C14NMethodParameterSpec) null),
             fac.newSignatureMethod("http://www.w3.org/2000/09/xmldsig#dsa-sha256",
null),
             Collections.singletonList(ref));
        // Create a DSA KeyPair
        KeyPairGenerator kpg = KeyPairGenerator.getInstance("DSA");
        kpg.initialize(2048);
        KeyPair kp = kpg.generateKeyPair();
        // Create a KeyValue containing the DSA PublicKey that was generated
        KeyInfoFactory kif = fac.getKeyInfoFactory();
        KeyValue kv = kif.newKeyValue(kp.getPublic());
        // Create a KeyInfo and add the KeyValue to it
        KeyInfo ki = kif.newKeyInfo(Collections.singletonList(kv));
        // Instantiate the document to be signed
        DocumentBuilderFactory dbf = DocumentBuilderFactory.newInstance();
        dbf.setNamespaceAware(true);
        Document doc =
            dbf.newDocumentBuilder().parse(new FileInputStream(args[0]));
        // Create a DOMSignContext and specify the DSA PrivateKey and
        // location of the resulting XMLSignature's parent element
        DOMSignContext dsc = new DOMSignContext
            (kp.getPrivate(), doc.getDocumentElement());
        // Create the XMLSignature (but don't sign it yet)
        XMLSignature signature = fac.newXMLSignature(si, ki);
        // Marshal, generate (and sign) the enveloped signature
```



```
signature.sign(dsc);

// output the resulting document
OutputStream os;
if (args.length > 1) {
    os = new FileOutputStream(args[1]);
} else {
    os = System.out;
}

TransformerFactory tf = TransformerFactory.newInstance();
Transformer trans = tf.newTransformer();
trans.transform(new DOMSource(doc), new StreamResult(os));
}
```

Example 11-9 envelope.xml

```
<Envelope xmlns="urn:envelope">
</Envelope>
```

Generating an XML Signature

This example shows you how to generate an XML Signature using the XML Digital Signature API. More specifically, the example generates an enveloped XML Signature of an XML document. An enveloped signature is a signature that is contained inside the content that it is signing. The example uses DOM (the Document Object Model) to parse the XML document to be signed and a JSR 105 DOM implementation to generate the resulting signature.

A basic knowledge of XML Signatures and their different components is helpful for understanding this section. See http://www.w3.org/TR/xmldsig-core/ for more information.

Instantiating the Document to be Signed

First, we use a JAXP DocumentBuilderFactory to parse the XML document that we want to sign. An application obtains the default implementation for DocumentBuilderFactory by calling the following line of code:

```
DocumentBuilderFactory dbf =
  DocumentBuilderFactory.newInstance();
```

We must also make the factory namespace-aware:

```
dbf.setNamespaceAware(true);
```

Next, we use the factory to get an instance of a DocumentBuilder, which is used to parse the document:

```
DocumentBuilder builder = dbf.newDocumentBuilder();
Document doc = builder.parse(new FileInputStream(argv[0]));
```

Creating a Public Key Pair



We generate a public key pair. Later in the example, we will use the private key to generate the signature. We create the key pair with a KeyPairGenerator. In this example, we will create a DSA KeyPair with a length of 2048 bytes:

```
KeyPairGenerator kpg = KeyPairGenerator.getInstance("DSA");
kpg.initialize(2048);
KeyPair kp = kpg.generateKeyPair();
```

In practice, the private key is usually previously generated and stored in a KeyStore file with an associated public key certificate.

Creating a Signing Context

We create an XML Digital Signature XMLSignContext containing input parameters for generating the signature. Since we are using DOM, we instantiate a DOMSignContext (a subclass of XMLSignContext), and pass it two parameters, the private key that will be used to sign the document and the root of the document to be signed:

```
DOMSignContext dsc = new DOMSignContext
  (kp.getPrivate(), doc.getDocumentElement());
```

Assembling the XML Signature

We assemble the different parts of the Signature element into an XMLSignature object. These objects are all created and assembled using an XMLSignatureFactory object. An application obtains a DOM implementation of XMLSignatureFactory by calling the following line of code:

```
XMLSignatureFactory fac =
   XMLSignatureFactory.getInstance("DOM");
```

We then invoke various factory methods to create the different parts of the XMLSignature object as shown below. We create a Reference object, passing to it the following:

- The URI of the object to be signed (We specify a URI of "", which implies the root of the document.)
- The DigestMethod (we use SHA1)
- A single Transform, the enveloped Transform, which is required for enveloped signatures so that the signature itself is removed before calculating the signature value

Next, we create the SignedInfo object, which is the object that is actually signed, as shown below. When creating the SignedInfo, we pass as parameters:

- The CanonicalizationMethod (we use inclusive and preserve comments)
- The SignatureMethod (we use DSA)
- A list of References (in this case, only one)



Next, we create the optional KeyInfo object, which contains information that enables the recipient to find the key needed to validate the signature. In this example, we add a KeyValue object containing the public key. To create KeyInfo and its various subtypes, we use a KeyInfoFactory object, which can be obtained by invoking the getKeyInfoFactory method of the XMLSignatureFactory, as follows:

```
KeyInfoFactory kif = fac.getKeyInfoFactory();
```

We then use the KeyInfoFactory to create the KeyValue object and add it to a KeyInfo object:

```
KeyValue kv = kif.newKeyValue(kp.getPublic());
KeyInfo ki = kif.newKeyInfo(Collections.singletonList(kv));
```

Finally, we create the XMLSignature object, passing as parameters the SignedInfo and KeyInfo objects that we created earlier:

```
XMLSignature signature = fac.newXMLSignature(si, ki);
```

Notice that we haven't actually generated the signature yet; we'll do that in the next step.

Generating the XML Signature

Now we are ready to generate the signature, which we do by invoking the sign method on the XMLSignature object, and pass it the signing context as follows:

```
signature.sign(dsc);
```

The resulting document now contains a signature, which has been inserted as the last child element of the root element.

Printing or Displaying the Resulting Document

You can use the following code to print the resulting signed document to a file or standard output:

```
OutputStream os;
if (args.length > 1) {
  os = new FileOutputStream(args[1]);
} else {
  os = System.out;
}
TransformerFactory tf = TransformerFactory.newInstance();
Transformer trans = tf.newTransformer();
trans.transform(new DOMSource(doc), new StreamResult(os));
```



Security API Specification

General Security

- java.security
- javax.crypto
- javax.security.cert
- javax.crypto.spec
- java.security.spec
- java.security.interfaces
- javax.crypto.interfaces
- javax.rmi.ssl

Certification Path

java.security.cert

JAAS

- javax.security.auth
- javax.security.auth.callback
- javax.security.auth.kerberos
- javax.security.auth.login
- javax.security.auth.spi
- javax.security.auth.x500

Java GSS-API

- org.ietf.jgss
- com.sun.security.jgss

JSSE

- javax.net
- javax.net.ssl
- java.security.cert

Java SASL

javax.security.sasl

SSL/TLS-Based RMI Socket Factories

javax.rmi.ssl

XML Digital Signature

- javax.xml.crypto
- javax.xml.crypto.dom
- javax.xml.crypto.dsig
- javax.xml.crypto.dsig.dom
- javax.xml.crypto.dsig.keyinfo
- javax.xml.crypto.dsig.spec

Smart Card I/O

javax.smartcardio



13

Deprecated Security APIs Marked for Removal

The following APIs are deprecated and eligible to be removed in a future release.

You can check the API dependencies using the jdeprscan tool. See jdeprscan in Java Platform, Standard Edition Tools Reference.

The following classes are deprecated and marked for removal:

- com.sun.security.auth.PolicyFile
- com.sun.security.auth.SolarisNumericGroupPrincipal
- com.sun.security.auth.SolarisNumericUserPrincipal
- com.sun.security.auth.SolarisPrincipal
- com.sun.security.auth.X500Principal
- com.sun.security.auth.module.SolarisLoginModule
- com.sun.security.auth.module.SolarisSystem

The following methods are deprecated and marked for removal:

- java.lang.SecurityManager.getInCheck
- java.lang.SecurityManager.checkMemberAccess
- java.lang.SecurityManager.classDepth
- java.lang.SecurityManager.currentClassLoader
- java.lang.SecurityManager.currentLoadedClass
- java.lang.SecurityManager.inClass
- java.lang.SecurityManager.inClassLoader
- java.lang.SecurityManager.checkAwtEventQueueAccess
- java.lang.SecurityManager.checkTopLevelWindow
- java.lang.SecurityManager.checkSystemClipboardAccess

The following field is deprecated and marked for removal:

• java.lang.SecurityManager.incheck

14

Security Tools

See Security Tools and Commands in *Java Platform, Standard Edition Tools Reference*.



15

Security Tutorials

- The Security Features in Java SE trail of the Java Tutorial
- JAAS Tutorials
- Java GSS-API and JAAS Tutorials for Use with Kerberos

