



Spring Security Reference

5.2.1.RELEASE

Ben Alex , Luke Taylor , Rob Winch , Gunnar Hillert , Joe Grandja , Jay Bryant , Eddú Meléndez , Josh Cummings

Copyright © 2004-2019

Copies of this document may be made for your own use and for distribution to others, provided that you do not charge any fee for such copies and further provided that each copy contains this Copyright Notice, whether distributed in print or electronically.

Table of Contents

.....	xviii
I. Preface	1
1. Prerequisites	2
2. Spring Security Community	3
2.1. Getting Help	3
2.2. Becoming Involved	3
2.3. Source Code	3
2.4. Apache 2 License	3
2.5. Social Media	3
3. What's New in Spring Security 5.2	4
3.1. Servlet	4
3.2. WebFlux	4
3.3. Core	5
4. Getting Spring Security	6
4.1. Release Numbering	6
4.2. Usage with Maven	6
Spring Boot with Maven	6
Maven Without Spring Boot	7
Maven Repositories	8
4.3. Gradle	8
Spring Boot with Gradle	9
Gradle Without Spring Boot	9
Gradle Repositories	10
5. Features	11
5.1. Protection Against Exploits	11
Cross Site Request Forgery (CSRF)	11
What is a CSRF Attack?	11
Protecting Against CSRF Attacks	12
When to use CSRF protection	15
CSRF Considerations	15
6. Project Modules	18
6.1. Core — <code>spring-security-core.jar</code>	18
6.2. Remoting — <code>spring-security-remoting.jar</code>	18
6.3. Web — <code>spring-security-web.jar</code>	18
6.4. Config — <code>spring-security-config.jar</code>	18
6.5. LDAP — <code>spring-security-ldap.jar</code>	18
6.6. OAuth 2.0 Core — <code>spring-security-oauth2-core.jar</code>	18
6.7. OAuth 2.0 Client — <code>spring-security-oauth2-client.jar</code>	19
6.8. OAuth 2.0 JOSE — <code>spring-security-oauth2-jose.jar</code>	19
6.9. OAuth 2.0 Resource Server — <code>spring-security-oauth2-resource-</code> <code>server.jar</code>	19
6.10. ACL — <code>spring-security-acl.jar</code>	19
6.11. CAS — <code>spring-security-cas.jar</code>	19
6.12. OpenID — <code>spring-security-openid.jar</code>	20
6.13. Test — <code>spring-security-test.jar</code>	20
7. Samples	21
II. Servlet Applications	22

8. Hello Spring Security	23
8.1. Hello Spring Security (Boot)	23
Updating Dependencies	23
Starting Hello Spring Security Boot	23
Spring Boot Auto Configuration	23
8.2. Hello Spring Security (Java Configuration)	24
Updating Dependencies	24
Minimal <code>@EnableWebSecurity</code> Configuration	24
Using <code>AbstractSecurityWebApplicationInitializer</code>	26
8.3. Hello Spring Security (XML)	27
Updating Dependencies	27
Minimal <code><http></code> Configuration	27
<code>web.xml</code> Configuration	28
9. Architecture and Implementation	30
9.1. Technical Overview	30
Runtime Environment	30
Core Components	30
SecurityContextHolder, SecurityContext and Authentication Objects	30
The UserDetailsService	31
GrantedAuthority	32
Summary	32
Authentication	32
What is authentication in Spring Security?	33
Setting the SecurityContextHolder Contents Directly	35
Authentication in a Web Application	35
ExceptionHandler	36
AuthenticationEntryPoint	36
Authentication Mechanism	36
Storing the SecurityContext between requests	36
Access-Control (Authorization) in Spring Security	37
Security and AOP Advice	37
Secure Objects and the AbstractSecurityInterceptor	38
9.2. Core Services	40
The AuthenticationManager, ProviderManager and AuthenticationProvider	40
Erasing Credentials on Successful Authentication	41
DaoAuthenticationProvider	41
UserDetailsService Implementations	41
In-Memory Authentication	42
JdbcDaoImpl	42
10. Authentication	44
10.1. In-Memory Authentication	44
10.2. JDBC Authentication	44
10.3. LDAP Authentication	44
Overview	44
Using LDAP with Spring Security	45
10.4. Configuring an LDAP Server	45
Using an Embedded Test Server	45
Using Bind Authentication	46
Loading Authorities	46
10.5. Implementation Classes	46

LdapAuthenticator Implementations	47
Common Functionality	47
BindAuthenticator	47
PasswordComparisonAuthenticator	47
Connecting to the LDAP Server	47
LDAP Search Objects	48
FilterBasedLdapUserSearch	48
LdapAuthoritiesPopulator	48
Spring Bean Configuration	48
LDAP Attributes and Customized UserDetails	49
10.6. Active Directory Authentication	49
ActiveDirectoryLdapAuthenticationProvider	50
Active Directory Error Codes	50
10.7. LDAP Java Configuration	50
10.8. AuthenticationProvider	51
AuthenticationProvider Java Configuration	51
AuthenticationProvider XML Configuration	52
10.9. UserDetailsService	52
10.10. Password Encoding	53
Password History	53
DelegatingPasswordEncoder	54
Password Storage Format	54
Password Encoding	55
Password Matching	55
Getting Started Experience	56
Troubleshooting	56
BCryptPasswordEncoder	57
Argon2PasswordEncoder	57
Pbkdf2PasswordEncoder	57
SCryptPasswordEncoder	58
Other PasswordEncoders	58
Password Encoder XML Configuration	58
10.11. The Authentication Manager and the Namespace	58
10.12. Session Management	59
Detecting Timeouts	59
Concurrent Session Control	60
Session Fixation Attack Protection	60
SessionManagementFilter	61
SessionAuthenticationStrategy	61
Concurrency Control	62
Querying the SessionRegistry for currently authenticated users and their sessions	63
10.13. Remember-Me Authentication	64
Overview	64
Simple Hash-Based Token Approach	64
Persistent Token Approach	65
Remember-Me Interfaces and Implementations	65
TokenBasedRememberMeServices	65
PersistentTokenBasedRememberMeServices	66
10.14. OpenID Support	66

Attribute Exchange	67
10.15. Anonymous Authentication	67
Overview	67
Configuration	68
AuthenticationTrustResolver	69
10.16. Pre-Authentication Scenarios	69
Pre-Authentication Framework Classes	70
AbstractPreAuthenticatedProcessingFilter	70
PreAuthenticatedAuthenticationProvider	70
Http403ForbiddenEntryPoint	71
Concrete Implementations	71
Request-Header Authentication (Siteminder)	71
Java EE Container Authentication	72
10.17. Java Authentication and Authorization Service (JAAS) Provider	72
Overview	72
AbstractJaasAuthenticationProvider	72
JAAS CallbackHandler	72
JAAS AuthorityGranter	73
DefaultJaasAuthenticationProvider	73
InMemoryConfiguration	73
DefaultJaasAuthenticationProvider Example Configuration	74
JaasAuthenticationProvider	74
Running as a Subject	75
10.18. CAS Authentication	75
Overview	75
How CAS Works	75
Spring Security and CAS Interaction Sequence	76
Configuration of CAS Client	78
Service Ticket Authentication	78
Single Logout	79
Authenticating to a Stateless Service with CAS	81
Proxy Ticket Authentication	82
10.19. X.509 Authentication	84
Overview	84
Adding X.509 Authentication to Your Web Application	84
Setting up SSL in Tomcat	85
10.20. Run-As Authentication Replacement	85
Overview	85
Configuration	86
10.21. Form Login	87
Form Login Java Configuration	87
Form Login XML Configuration	88
Form and Basic Login Options	88
10.22. Basic and Digest Authentication	89
BasicAuthenticationFilter	89
Configuration	90
10.23. DigestAuthenticationFilter	90
Configuration	91
10.24. Handling Logouts	92
Logout Java Configuration	92

Logout XML Configuration	93
LogoutHandler	93
LogoutSuccessHandler	93
Further Logout-Related References	94
10.25. Setting a Custom AuthenticationEntryPoint	94
11. Authorization	95
11.1. Authorization Architecture	95
Authorities	95
Pre-Invocation Handling	95
The AccessDecisionManager	96
Voting-Based AccessDecisionManager Implementations	96
After Invocation Handling	98
Hierarchical Roles	99
11.2. Secure Object Implementations	100
AOP Alliance (MethodInvocation) Security Interceptor	100
Explicit MethodSecurityInterceptor Configuration	100
AspectJ (JoinPoint) Security Interceptor	101
11.3. Expression-Based Access Control	102
Overview	103
Common Built-In Expressions	103
Web Security Expressions	104
Referring to Beans in Web Security Expressions	104
Path Variables in Web Security Expressions	105
Method Security Expressions	105
@Pre and @Post Annotations	105
Built-In Expressions	107
11.4. Authorize Requests	108
11.5. Method Security	109
EnableGlobalMethodSecurity	109
GlobalMethodSecurityConfiguration	110
The <global-method-security> Element	111
Adding Security Pointcuts using protect-pointcut	112
11.6. Domain Object Security (ACLs)	112
Overview	112
Key Concepts	113
Getting Started	115
12. OAuth2	117
12.1. OAuth 2.0 Login	117
Spring Boot 2.x Sample	117
Initial setup	117
Setting the redirect URI	117
Configure application.yml	118
Boot up the application	118
Spring Boot 2.x Property Mappings	118
CommonOAuth2Provider	119
Configuring Custom Provider Properties	120
Overriding Spring Boot 2.x Auto-configuration	121
Register a ClientRegistrationRepository @Bean	121
Provide a WebSecurityConfigurerAdapter	122
Completely Override the Auto-configuration	122

Java Configuration without Spring Boot 2.x	123
Advanced Configuration	124
OAuth 2.0 Login Page	126
Redirection Endpoint	127
UserInfo Endpoint	128
ID Token Signature Verification	134
OpenID Connect 1.0 Logout	135
12.2. OAuth 2.0 Client	136
Core Interfaces / Classes	138
ClientRegistration	138
ClientRegistrationRepository	139
OAuth2AuthorizedClient	140
OAuth2AuthorizedClientRepository / OAuth2AuthorizedClientService	140
OAuth2AuthorizedClientManager / OAuth2AuthorizedClientProvider	141
Authorization Grant Support	143
Authorization Code	143
Refresh Token	149
Client Credentials	151
Resource Owner Password Credentials	154
Additional Features	158
Resolving an Authorized Client	158
WebClient integration for Servlet Environments	159
Providing the Authorized Client	159
Defaulting the Authorized Client	160
12.3. OAuth 2.0 Resource Server	161
Dependencies	161
Minimal Configuration for JWTs	161
Specifying the Authorization Server	161
Startup Expectations	162
Runtime Expectations	162
Specifying the Authorization Server JWK Set Uri Directly	163
Overriding or Replacing Boot Auto Configuration	163
Using <code>jwtSetUri()</code>	164
Using <code>decoder()</code>	165
Exposing a <code>JwtDecoder @Bean</code>	165
Configuring Trusted Algorithms	165
Via Spring Boot	165
Using a Builder	165
From JWK Set response	166
Trusting a Single Asymmetric Key	166
Via Spring Boot	166
Using a Builder	167
Trusting a Single Symmetric Key	167
Configuring Authorization	167
Extracting Authorities Manually	168
Configuring Validation	168
Customizing Timestamp Validation	169
Configuring a Custom Validator	169
Configuring Claim Set Mapping	170
Customizing the Conversion of a Single Claim	170

Adding a Claim	170
Removing a Claim	171
Renaming a Claim	171
Configuring Timeouts	171
Minimal Configuration for Introspection	171
Specifying the Authorization Server	172
Startup Expectations	172
Runtime Expectations	172
Looking Up Attributes Post-Authentication	173
Looking Up Attributes Via SpEL	173
Overriding or Replacing Boot Auto Configuration	173
Using <code>introspectionUri()</code>	174
Using <code>introspector()</code>	174
Exposing a <code>OpaqueTokenIntrospector @Bean</code>	175
Configuring Authorization	175
Extracting Authorities Manually	175
Configuring Timeouts	176
Using Introspection with JWTs	176
Calling a <code>/userinfo</code> Endpoint	177
Supporting both JWT and Opaque Token	178
Multi-tenancy	179
Resolving the Tenant By Request Material	179
Resolving the Tenant By Claim	180
Parsing the Claim Only Once	181
Bearer Token Resolution	184
Reading the Bearer Token from a Custom Header	184
Reading the Bearer Token from a Form Parameter	184
Bearer Token Propagation	184
<code>RestTemplate</code> support	185
13. SAML2	186
13.1. SAML 2.0 Login	186
SAML 2 Support in Spring Security	186
Saml 2 Login - High Level Concepts	186
Saml 2 Login - Current Feature Set	186
Saml 2 Login - Not Yet Supported	187
Saml 2 Login - Introduction to Java Configuration	187
<code>RelyingPartyRegistration</code>	188
Service Provider Metadata	189
Authentication Requests - SP Initiated Flow	190
Spring Boot 2.x Sample	190
Multiple Identity Provider Sample	190
14. Protection Against Exploits	192
14.1. Cross Site Request Forgery (CSRF) for Servlet Environments	192
Using Spring Security CSRF Protection	192
Use proper HTTP verbs	192
Configure CSRF Protection	192
Include the CSRF Token	193
CSRF Considerations	195
Logging In	196
Logging Out	196

CSRF and Session Timeouts	196
Multipart (file upload)	196
HiddenHttpMethodFilter	198
14.2. Security HTTP Response Headers	198
Default Security Headers	198
Cache Control	200
Content Type Options	201
HTTP Strict Transport Security (HSTS)	202
HTTP Public Key Pinning (HPKP)	203
X-Frame-Options	205
X-XSS-Protection	206
Content Security Policy (CSP)	206
Configuring Content Security Policy	207
Additional Resources	209
Referrer Policy	209
Configuring Referrer Policy	209
Feature Policy	209
Configuring Feature Policy	210
Clear Site Data	210
Configuring Clear Site Data	210
Custom Headers	211
Static Headers	211
Headers Writer	212
DelegatingRequestMatcherHeaderWriter	212
14.3. HTTPS	213
Adding HTTP/HTTPS Channel Security	213
15. Integrations	215
15.1. Servlet API integration	215
Servlet 2.5+ Integration	215
HttpServletRequest.getRemoteUser()	215
HttpServletRequest.getUserPrincipal()	215
HttpServletRequest.isUserInRole(String)	215
Servlet 3+ Integration	216
HttpServletRequest.authenticate(HttpServletRequest, HttpServletResponse)	216
HttpServletRequest.login(String, String)	216
HttpServletRequest.logout()	216
AsyncContext.start(Runnable)	216
Async Servlet Support	217
Servlet 3.1+ Integration	218
HttpServletRequest#changeSessionId()	218
15.2. Spring Data Integration	218
Spring Data & Spring Security Configuration	218
Security Expressions within @Query	218
15.3. Concurrency Support	218
DelegatingSecurityContextRunnable	219
DelegatingSecurityContextExecutor	220
Spring Security Concurrency Classes	221
15.4. Jackson Support	221
15.5. Localization	222

15.6. Spring MVC Integration	223
@EnableWebMvcSecurity	223
MvcRequestMatcher	223
@AuthenticationPrincipal	225
Spring MVC Async Integration	227
Spring MVC and CSRF Integration	228
Automatic Token Inclusion	228
Resolving the CsrfToken	228
15.7. WebSocket Security	229
WebSocket Configuration	229
WebSocket Authentication	230
WebSocket Authorization	230
WebSocket Authorization Notes	231
Outbound Messages	232
Enforcing Same Origin Policy	232
Why Same Origin?	232
Spring WebSocket Allowed Origin	233
Adding CSRF to Stomp Headers	233
Disable CSRF within WebSockets	233
Working with SockJS	234
SockJS & frame-options	234
SockJS & Relaxing CSRF	234
15.8. CORS	235
15.9. JSP Tag Libraries	236
Declaring the Taglib	236
The authorize Tag	237
Disabling Tag Authorization for Testing	237
The authentication Tag	238
The accesscontrollist Tag	238
The csrfInput Tag	238
The csrfMetaTags Tag	239
16. Java Configuration	241
16.1. Hello Web Security Java Configuration	241
AbstractSecurityWebApplicationInitializer	242
AbstractSecurityWebApplicationInitializer without Existing Spring	242
AbstractSecurityWebApplicationInitializer with Spring MVC	243
16.2. HttpSecurity	243
16.3. Multiple HttpSecurity	244
16.4. Custom DSLs	245
16.5. Post Processing Configured Objects	246
17. Security Namespace Configuration	247
17.1. Introduction	247
Design of the Namespace	248
17.2. Getting Started with Security Namespace Configuration	248
web.xml Configuration	248
A Minimal <http> Configuration	249
Setting a Default Post-Login Destination	250
17.3. Advanced Web Features	251
Adding in Your Own Filters	251
17.4. Method Security	253

17.5. The Default AccessDecisionManager	253
Customizing the AccessDecisionManager	253
18. Testing	254
18.1. Testing Method Security	254
Security Test Setup	254
@WithMockUser	255
@WithAnonymousUser	256
@WithUserDetails	257
@WithSecurityContext	258
Test Meta Annotations	259
18.2. Spring MVC Test Integration	259
Setting Up MockMvc and Spring Security	259
SecurityMockMvcRequestPostProcessors	260
Testing with CSRF Protection	260
Running a Test as a User in Spring MVC Test	260
Running as a User in Spring MVC Test with RequestPostProcessor	261
Testing HTTP Basic Authentication	262
SecurityMockMvcRequestBuilders	263
Testing Form Based Authentication	263
Testing Bearer Authentication	263
Testing Logout	265
SecurityMockMvcResultMatchers	265
Unauthenticated Assertion	265
Authenticated Assertion	265
19. Spring Security Crypto Module	267
19.1. Introduction	267
19.2. Encryptors	267
BytesEncryptor	267
TextEncryptor	267
19.3. Key Generators	268
BytesKeyGenerator	268
StringKeyGenerator	268
19.4. Password Encoding	268
20. Appendix	270
20.1. Security Database Schema	270
User Schema	270
For Oracle database	270
Group Authorities	270
Persistent Login (Remember-Me) Schema	271
ACL Schema	271
HyperSQL	272
PostgreSQL	273
MySQL and MariaDB	274
Microsoft SQL Server	275
Oracle Database	276
20.2. The Security Namespace	277
Web Application Security	277
<debug>	277
<http>	277
<access-denied-handler>	279

<cors>	280
<headers>	280
<cache-control>	281
<hsts>	282
<hpkp>	282
<pins>	282
<pin>	282
<content-security-policy>	283
<referrer-policy>	283
<feature-policy>	283
<frame-options>	283
<xss-protection>	284
<content-type-options>	285
<header>	285
<anonymous>	285
<csrf>	286
<custom-filter>	286
<expression-handler>	286
<form-login>	287
<http-basic>	288
<http-firewall> Element	288
<intercept-url>	288
<jee>	289
<logout>	290
<openid-login>	290
<attribute-exchange>	291
<openid-attribute>	292
<port-mappings>	292
<port-mapping>	292
<remember-me>	293
<request-cache> Element	294
<session-management>	294
<concurrency-control>	295
<x509>	295
<filter-chain-map>	296
<filter-chain>	296
<filter-security-metadata-source>	296
WebSocket Security	297
<websocket-message-broker>	297
<intercept-message>	298
Authentication Services	298
<authentication-manager>	298
<authentication-provider>	299
<jdbc-user-service>	299
<password-encoder>	300
<user-service>	300
<user>	301
Method Security	301
<global-method-security>	301
<after-invocation-provider>	302

<pre-post-annotation-handling>	302
<invocation-attribute-factory>	303
<post-invocation-advice>	303
<pre-invocation-advice>	303
Securing Methods using	303
<intercept-methods>	304
<method-security-metadata-source>	304
<protect>	304
LDAP Namespace Options	304
Defining the LDAP Server using the	305
<ldap-authentication-provider>	305
<password-compare>	306
<ldap-user-service>	307
20.3. Spring Security Dependencies	308
spring-security-core	308
spring-security-remoting	308
spring-security-web	309
spring-security-ldap	309
spring-security-config	310
spring-security-acl	310
spring-security-cas	311
spring-security-openid	311
spring-security-taglibs	311
20.4. Proxy Server Configuration	312
20.5. Spring Security FAQ	312
General Questions	312
Will Spring Security take care of all my application security requirements?	313
Why not just use web.xml security?	313
What Java and Spring Framework versions are required?	314
I'm new to Spring Security and I need to build an application that supports CAS single sign-on over HTTPS, while allowing Basic authentication locally for certain URLs, authenticating against multiple back end user information sources (LDAP and JDBC). I've copied some configuration files I found but it doesn't work.	314
Common Problems	314
When I try to log in, I get an error message that says "Bad Credentials". What's wrong?	315
My application goes into an "endless loop" when I try to login, what's going on?	316
I get an exception with the message "Access is denied (user is anonymous);". What's wrong?	316
Why can I still see a secured page even after I've logged out of my application?	316
I get an exception with the message "An Authentication object was not found in the SecurityContext". What's wrong?	316
I can't get LDAP authentication to work.	317
Session Management	317
I'm using Spring Security's concurrent session control to prevent users from logging in more than once at a time.	317

Why does the session Id change when I authenticate through Spring Security?	317
I'm using Tomcat (or some other servlet container) and have enabled HTTPS for my login page, switching back to HTTP afterwards.	318
I'm not switching between HTTP and HTTPS but my session is still getting lost	318
I'm trying to use the concurrent session-control support but it won't let me log back in, even if I'm sure I've logged out and haven't exceeded the allowed sessions.	318
Spring Security is creating a session somewhere, even though I've configured it not to, by setting the create-session attribute to never.	318
I get a 403 Forbidden when performing a POST	319
I'm forwarding a request to another URL using the RequestDispatcher, but my security constraints aren't being applied.	319
I have added Spring Security's <global-method-security> element to my application context but if I add security annotations to my Spring MVC controller beans (Struts actions etc.) then they don't seem to have an effect.	319
I have a user who has definitely been authenticated, but when I try to access the SecurityContextHolder during some requests, the Authentication is null.	319
The authorize JSP Tag doesn't respect my method security annotations when using the URL attribute.	319
Spring Security Architecture Questions	319
How do I know which package class X is in?	320
How do the namespace elements map to conventional bean configurations?	320
What does "ROLE_" mean and why do I need it on my role names?	320
How do I know which dependencies to add to my application to work with Spring Security?	320
What dependencies are needed to run an embedded ApacheDS LDAP server?	321
What is a UserDetailsService and do I need one?	321
Common "Howto" Requests	321
I need to login in with more information than just the username.	322
How do I apply different intercept-url constraints where only the fragment value of the requested URLs differs (e.g./foo#bar and /foo#blah?	322
How do I access the user's IP Address (or other web-request data) in a UserDetailsService?	322
How do I access the HttpSession from a UserDetailsService?	322
How do I access the user's password in a UserDetailsService?	323
How do I define the secured URLs within an application dynamically?	323
How do I authenticate against LDAP but load user roles from a database?	324
I want to modify the property of a bean that is created by the namespace, but there is nothing in the schema to support it.	324
III. Reactive Applications	326
21. WebFlux Security	327
21.1. Minimal WebFlux Security Configuration	327
21.2. Explicit WebFlux Security Configuration	327

22. Protection Against Exploits	329
22.1. Cross Site Request Forgery (CSRF) for WebFlux Environments	329
Using Spring Security CSRF Protection	329
Use proper HTTP verbs	329
Configure CSRF Protection	329
Include the CSRF Token	330
CSRF Considerations	332
Logging In	332
Logging Out	332
CSRF and Session Timeouts	333
Multipart (file upload)	333
HiddenHttpMethodFilter	333
22.2. Security HTTP Response Headers	333
Default Security Headers	334
Cache Control	335
Content Type Options	336
HTTP Strict Transport Security (HSTS)	336
X-Frame-Options	337
X-XSS-Protection	338
Content Security Policy (CSP)	338
Configuring Content Security Policy	339
Additional Resources	340
Referrer Policy	340
Configuring Referrer Policy	340
Feature Policy	341
Configuring Feature Policy	341
Clear Site Data	341
Configuring Clear Site Data	341
22.3. Redirect to HTTPS	342
23. OAuth2 WebFlux	343
23.1. OAuth 2.0 Login	343
Spring Boot 2.0 Sample	343
Initial setup	343
Setting the redirect URI	343
Configure <code>application.yml</code>	344
Boot up the application	344
Using OpenID Provider Configuration	344
Explicit OAuth2 Login Configuration	345
23.2. OAuth2 Client	346
23.3. OAuth 2.0 Resource Server	346
Dependencies	346
Minimal Configuration for JWTs	346
Specifying the Authorization Server	347
Startup Expectations	347
Runtime Expectations	347
Specifying the Authorization Server JWK Set Uri Directly	348
Overriding or Replacing Boot Auto Configuration	348
Configuring Trusted Algorithms	350
Via Spring Boot	350
Using a Builder	351

Trusting a Single Asymmetric Key	351
Trusting a Single Symmetric Key	352
Configuring Authorization	352
Configuring Validation	353
Minimal Configuration for Introspection	354
Looking Up Attributes Post-Authentication	356
Overriding or Replacing Boot Auto Configuration	356
Configuring Authorization	358
Using Introspection with JWTs	359
Calling a <code>/userinfo</code> Endpoint	360
Bearer Token Propagation	361
24. <code>@RegisteredOAuth2AuthorizedClient</code>	363
25. Reactive X.509 Authentication	364
26. <code>WebClient</code>	365
26.1. <code>WebClient OAuth2 Setup</code>	365
26.2. <code>Implicit OAuth2AuthorizedClient</code>	365
26.3. <code>Explicit OAuth2AuthorizedClient</code>	366
26.4. <code>clientRegistrationId</code>	366
27. <code>EnableReactiveMethodSecurity</code>	367
28. Reactive Test Support	369
28.1. Testing Reactive Method Security	369
28.2. <code>WebTestClientSupport</code>	369
Authentication	370
CSRF Support	370
Testing Bearer Authentication	371
<code>mockJwt()</code> <code>WebTestClientConfigurer</code>	371
<code>authentication()</code> <code>WebTestClientConfigurer</code>	372
29. <code>RSocket Security</code>	373
29.1. Minimal <code>RSocket Security Configuration</code>	373
29.2. Adding <code>SecuritySocketAcceptorInterceptor</code>	373
29.3. <code>RSocket Authentication</code>	373
Authentication at Setup vs Request Time	373
Basic Authentication	374
JWT	375
29.4. <code>RSocket Authorization</code>	375

Spring Security is a framework that provides authentication, authorization, and protection against common attacks. With first class support for both imperative and reactive applications, it is the de-facto standard for securing Spring-based applications.

Part I. Preface

This section discusses the logistics of Spring Security.

1. Prerequisites

Spring Security requires a Java 8 or higher Runtime Environment.

As Spring Security aims to operate in a self-contained manner, you do not need to place any special configuration files in your Java Runtime Environment. In particular, you need not configure a special Java Authentication and Authorization Service (JAAS) policy file or place Spring Security into common classpath locations.

Similarly, if you use an EJB Container or Servlet Container, you need not put any special configuration files anywhere nor include Spring Security in a server classloader. All the required files are contained within your application.

This design offers maximum deployment time flexibility, as you can copy your target artifact (be it a JAR, WAR, or EAR) from one system to another and it immediately works.

2. Spring Security Community

Welcome to the Spring Security Community! This section discusses how you can make the most of our vast community.

2.1 Getting Help

If you need help with Spring Security, we are here to help. The following are some of the best ways to get help:

- Read through this documentation.
- Try one of our many [sample applications](#).
- Ask a question on <https://stackoverflow.com> with the `spring-security` tag.
- Report bugs and enhancement requests at <https://github.com/spring-projects/spring-security/issues>

2.2 Becoming Involved

We welcome your involvement in the Spring Security project. There are many ways to contribute, including answering questions on StackOverflow, writing new code, improving existing code, assisting with documentation, developing samples or tutorials, reporting bugs, or simply making suggestions. For more information, see our [Contributing](#) documentation.

2.3 Source Code

You can find Spring Security's source code on GitHub at <https://github.com/spring-projects/spring-security/>

2.4 Apache 2 License

Spring Security is Open Source software released under the [Apache 2.0 license](#).

2.5 Social Media

You can follow [@SpringSecurity](#) and the [Spring Security team](#) on Twitter to stay up to date with the latest news. You can also follow [@SpringCentral](#) to keep up to date with the entire Spring portfolio.

3. What's New in Spring Security 5.2

Spring Security 5.2 provides a number of new features. Below are the highlights of the release.

3.1 Servlet

- Added [nested builder](#) support in HTTP Security DSL
- OAuth 2.0 Client
 - Introducing [OAuth2AuthorizedClientManager / OAuth2AuthorizedClientProvider](#)
 - Added [AuthorizedClientServiceOAuth2AuthorizedClientManager](#) which is capable of operating outside of a `HttpServletRequest` context
 - Public Client support with [PKCE](#)
 - Support for [Resource Owner Password Credentials](#) grant
 - Support for ID Token verification using a [Symmetric Key](#) via `NimbusJwtDecoder`
 - Added [nonce](#) to OpenID Connect Authentication Request
 - OpenID Connect [RP-Initiated Logout](#)
 - Updated [documentation](#)
- OAuth 2.0 Resource Server
 - Introducing [Token Introspection](#) (Opaque Tokens)
 - [Multi-tenancy](#) support
 - Added `ExchangeFilterFunction` that performs [Bearer Token propagation](#) (Token Relay)
 - Support for multiple [JWS algorithms](#) via `NimbusJwtDecoder`
 - Test support for [mock JWT](#)
 - Added [JWE](#) sample
 - Updated [documentation](#)

3.2 WebFlux

- Added [nested builder](#) support in HTTP Security DSL
- OAuth 2.0 Client
 - Introducing [ReactiveOAuth2AuthorizedClientManager / ReactiveOAuth2AuthorizedClientProvider](#)
 - Public Client support with [PKCE](#)
 - Support for [Resource Owner Password Credentials](#) grant
 - Support for ID Token verification using a [Symmetric Key](#) via `NimbusReactiveJwtDecoder`

- Added [nonce](#) to OpenID Connect Authentication Request
- OpenID Connect [RP-Initiated Logout](#)
- OAuth 2.0 Resource Server
 - Introducing [Token Introspection](#) (Opaque Tokens)
 - [Multi-tenancy](#) support
 - Added ExchangeFilterFunction that performs [Bearer Token propagation](#) (Token Relay)
 - Support for multiple [JWS algorithms](#) via NimbusReactiveJwtDecoder
- Support for [X509](#)

3.3 Core

- Introducing [RSocket](#) support
- Introducing [SAML Service Provider](#) support
- Introducing [AuthenticationManagerResolver](#)
- Introducing [AuthenticationFilter](#)
- Introducing [@CurrentSecurityContext](#) for method arguments
- Converting [key material](#) to Key instances
- Support for [Clear-Site-Data](#) header
- Introducing [CompositeHeaderWriter](#)
- Added [nohttp](#) to build
- [JDK 12](#) support
- Support for [path variables](#) in message expressions
- Configuration classes are proxy-less and support [proxyBeanMethods=false](#)
- Added [Argon2PasswordEncoder](#)
- Support upgrading between different [BCrypt encodings](#)
- Support upgrading between different [SCrypt encodings](#)

4. Getting Spring Security

This section discusses all you need to know about getting the Spring Security binaries. See Section 2.3, “Source Code” for how to obtain the source code.

4.1 Release Numbering

Spring Security versions are formatted as MAJOR.MINOR.PATCH such that:

- MAJOR versions may contain breaking changes. Typically, these are done to provide improved security to match modern security practices.
- MINOR versions contain enhancements but are considered passive updates
- PATCH level should be perfectly compatible, forwards and backwards, with the possible exception of changes that fix bugs.

4.2 Usage with Maven

As most open source projects, Spring Security deploys its dependencies as Maven artifacts. The topics in this section provide detail on how to consume Spring Security when using Maven.

Spring Boot with Maven

Spring Boot provides a `spring-boot-starter-security` starter that aggregates Spring Security-related dependencies together. The simplest and preferred way to use the starter is to use [Spring Initializr](#) by using an IDE integration ([Eclipse](#), [IntelliJ](#), [NetBeans](#)) or through <https://start.spring.io>.

Alternatively, you can manually add the starter, as the following example shows:

```
<dependencies>
  <!-- ... other dependency elements ... -->
  <dependency>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-starter-security</artifactId>
  </dependency>
</dependencies>
```

Example 4.1 pom.xml

Since Spring Boot provides a Maven BOM to manage dependency versions, you do not need to specify a version. If you wish to override the Spring Security version, you may do so by providing a Maven property, as the following example shows:

```
<properties>
  <!-- ... -->
  <spring-security.version>5.2.1.RELEASE</spring-security.version>
</dependencies>
```

Example 4.2 pom.xml

Since Spring Security makes breaking changes only in major releases, it is safe to use a newer version of Spring Security with Spring Boot. However, at times, you may need to update the version of Spring Framework as well. You can do so by adding a Maven property, as the following example shows:


```

<properties>
  <!-- ... -->
  <spring.version>5.2.1.RELEASE</spring.version>
</dependencies>

```

Example 4.3 pom.xml

If you use additional features (such as LDAP, OpenID, and others), you need to also include the appropriate Chapter 6, *Project Modules*.

Maven Without Spring Boot

When you use Spring Security without Spring Boot, the preferred way is to use Spring Security's BOM to ensure a consistent version of Spring Security is used throughout the entire project. The following example shows how to do so:

```

<dependencyManagement>
  <dependencies>
    <!-- ... other dependency elements ... -->
    <dependency>
      <groupId>org.springframework.security</groupId>
      <artifactId>spring-security-bom</artifactId>
      <version>5.2.1.RELEASE</version>
      <type>pom</type>
      <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>

```

Example 4.4 pom.xml

A minimal Spring Security Maven set of dependencies typically looks like the following:

```

<dependencies>
  <!-- ... other dependency elements ... -->
  <dependency>
    <groupId>org.springframework.security</groupId>
    <artifactId>spring-security-web</artifactId>
  </dependency>
  <dependency>
    <groupId>org.springframework.security</groupId>
    <artifactId>spring-security-config</artifactId>
  </dependency>
</dependencies>

```

Example 4.5 pom.xml

If you use additional features (such as LDAP, OpenID, and others), you need to also include the appropriate Chapter 6, *Project Modules*.

Spring Security builds against Spring Framework 5.2.1.RELEASE but should generally work with any newer version of Spring Framework 5.x. Many users are likely to run afoul of the fact that Spring Security's transitive dependencies resolve Spring Framework 5.2.1.RELEASE, which can cause strange classpath problems. The easiest way to resolve this is to use the `spring-framework-bom` within the `<dependencyManagement>` section of your `pom.xml` as the following example shows:

```

<dependencyManagement>
  <dependencies>
    <!-- ... other dependency elements ... -->
    <dependency>
      <groupId>org.springframework</groupId>
      <artifactId>spring-framework-bom</artifactId>
      <version>5.2.1.RELEASE</version>
      <type>pom</type>
      <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>

```

Example 4.6 pom.xml

The preceding example ensures that all the transitive dependencies of Spring Security use the Spring 5.2.1.RELEASE modules.

Note

This approach uses Maven's "bill of materials" (BOM) concept and is only available in Maven 2.0.9+. For additional details about how dependencies are resolved, see [Maven's Introduction to the Dependency Mechanism documentation](#).

Maven Repositories

All GA releases (that is, versions ending in .RELEASE) are deployed to Maven Central, so no additional Maven repositories need to be declared in your pom.

If you use a SNAPSHOT version, you need to ensure that you have the Spring Snapshot repository defined, as the following example shows:

```

<repositories>
  <!-- ... possibly other repository elements ... -->
  <repository>
    <id>spring-snapshot</id>
    <name>Spring Snapshot Repository</name>
    <url>https://repo.spring.io/snapshot</url>
  </repository>
</repositories>

```

Example 4.7 pom.xml

If you use a milestone or release candidate version, you need to ensure that you have the Spring Milestone repository defined, as the following example shows:

```

<repositories>
  <!-- ... possibly other repository elements ... -->
  <repository>
    <id>spring-milestone</id>
    <name>Spring Milestone Repository</name>
    <url>https://repo.spring.io/milestone</url>
  </repository>
</repositories>

```

Example 4.8 pom.xml

4.3 Gradle

As most open source projects, Spring Security deploys its dependencies as Maven artifacts, which allows for first-class Gradle support. The following topics provide detail on how to consume Spring Security when using Gradle.

Spring Boot with Gradle

Spring Boot provides a `spring-boot-starter-security` starter that aggregates Spring Security related dependencies together. The simplest and preferred method to use the starter is to use [Spring Initializr](#) by using an IDE integration ([Eclipse](#), [IntelliJ](#), [NetBeans](#)) or through <https://start.spring.io>.

Alternatively, you can manually add the starter, as the following example shows:

```
dependencies {
    compile "org.springframework.boot:spring-boot-starter-security"
}
```

Example 4.9 build.gradle

Since Spring Boot provides a Maven BOM to manage dependency versions, you need not specify a version. If you wish to override the Spring Security version, you may do so by providing a Gradle property, as the following example shows:

```
ext['spring-security.version']='5.2.1.RELEASE'
```

Example 4.10 build.gradle

Since Spring Security makes breaking changes only in major releases, it is safe to use a newer version of Spring Security with Spring Boot. However, at times, you may need to update the version of Spring Framework as well. You can do so by adding a Gradle property, as the following example shows:

```
ext['spring.version']='5.2.1.RELEASE'
```

Example 4.11 build.gradle

If you use additional features (such as LDAP, OpenID, and others), you need to also include the appropriate Chapter 6, *Project Modules*.

Gradle Without Spring Boot

When you use Spring Security without Spring Boot, the preferred way is to use Spring Security's BOM to ensure a consistent version of Spring Security is used throughout the entire project. You can do so by using the [Dependency Management Plugin](#), as the following example shows:

```
plugins {
    id "io.spring.dependency-management" version "1.0.6.RELEASE"
}

dependencyManagement {
    imports {
        mavenBom 'org.springframework.security:spring-security-bom:5.2.1.RELEASE'
    }
}
```

Example 4.12 build.gradle

A minimal Spring Security Maven set of dependencies typically looks like the following:

```
dependencies {
    compile "org.springframework.security:spring-security-web"
    compile "org.springframework.security:spring-security-config"
}
```

Example 4.13 build.gradle

If you use additional features (such as LDAP, OpenID, and others), you need to also include the appropriate Chapter 6, *Project Modules*.

Spring Security builds against Spring Framework 5.2.1.RELEASE but should generally work with any newer version of Spring Framework 5.x. {JB} Many users are likely to run afoul of the fact that Spring Security's transitive dependencies resolve Spring Framework 5.2.1.RELEASE, which can cause strange classpath problems. The easiest way to resolve this is to use the `spring-framework-bom` within your `<dependencyManagement>` section of your `pom.xml`. You can do so by using the [Dependency Management Plugin](#), as the following example shows:

```
plugins {
    id "io.spring.dependency-management" version "1.0.6.RELEASE"
}

dependencyManagement {
    imports {
        mavenBom 'org.springframework:spring-framework-bom:5.2.1.RELEASE'
    }
}
```

Example 4.14 build.gradle

The preceding example ensures that all the transitive dependencies of Spring Security use the Spring 5.2.1.RELEASE modules.

Gradle Repositories

All GA releases (that is, versions ending in `.RELEASE`) are deployed to Maven Central, so using the `mavenCentral()` repository is sufficient for GA releases. The following example shows how to do so:

```
repositories {
    mavenCentral()
}
```

Example 4.15 build.gradle

If you use a SNAPSHOT version, you need to ensure you have the Spring Snapshot repository defined, as the following example shows:

```
repositories {
    maven { url 'https://repo.spring.io/snapshot' }
}
```

Example 4.16 build.gradle

If you use a milestone or release candidate version, you need to ensure that you have the Spring Milestone repository defined, as the following example shows:

```
repositories {
    maven { url 'https://repo.spring.io/milestone' }
}
```

Example 4.17 build.gradle

5. Features

Spring Security provides comprehensive support for authentication, authorization, and protection against [common exploits](#). It also provides integration with other libraries to simplify its usage.

5.1 Protection Against Exploits

Spring Security provides protection against common exploits. Whenever possible, the protection is enabled by default. Below you will find high level description of the various exploits that Spring Security protects against.

Cross Site Request Forgery (CSRF)

Spring provides comprehensive support for protecting against [Cross Site Request Forgery \(CSRF\)](#) attacks. In the following sections we will explore:

- the section called “What is a CSRF Attack?”
- the section called “Protecting Against CSRF Attacks”
- the section called “CSRF Considerations”

Note

This portion of the documentation discusses the general topic of CSRF protection. Refer to the relevant sections for specific information on CSRF protection for [servlet](#) and [WebFlux](#) based applications.

What is a CSRF Attack?

The best way to understand a CSRF attack is by taking a look at a concrete example.

Assume that your bank’s website provides a form that allows transferring money from the currently logged in user to another bank account. For example, the transfer form might look like:

```
<form method="post"
  action="/transfer">
<input type="text"
  name="amount"/>
<input type="text"
  name="routingNumber"/>
<input type="text"
  name="account"/>
<input type="submit"
  value="Transfer"/>
</form>
```

Example 5.1 Transfer form

The corresponding HTTP request might look like:

```
POST /transfer HTTP/1.1
Host: bank.example.com
Cookie: JSESSIONID=randomid
Content-Type: application/x-www-form-urlencoded

amount=100.00&routingNumber=1234&account=9876
```

Example 5.2 Transfer HTTP request

Now pretend you authenticate to your bank's website and then, without logging out, visit an evil website. The evil website contains an HTML page with the following form:

```
<form method="post"
      action="https://bank.example.com/transfer">
  <input type="hidden"
        name="amount"
        value="100.00"/>
  <input type="hidden"
        name="routingNumber"
        value="evilsRoutingNumber"/>
  <input type="hidden"
        name="account"
        value="evilsAccountNumber"/>
  <input type="submit"
        value="Win Money!"/>
</form>
```

Example 5.3 Evil transfer form

You like to win money, so you click on the submit button. In the process, you have unintentionally transferred \$100 to a malicious user. This happens because, while the evil website cannot see your cookies, the cookies associated with your bank are still sent along with the request.

Worst yet, this whole process could have been automated using JavaScript. This means you didn't even need to click on the button. Furthermore, it could just as easily happen when visiting an honest site that is a victim of a [XSS attack](#). So how do we protect our users from such attacks?

Protecting Against CSRF Attacks

The reason that a CSRF attack is possible is that the HTTP request from the victim's website and the request from the attacker's website are exactly the same. This means there is no way to reject requests coming from the evil website and allow requests coming from the bank's website. To protect against CSRF attacks we need to ensure there is something in the request that the evil site is unable to provide so we can differentiate the two requests.

Spring provides two mechanisms to protect against CSRF attacks:

- The the section called "Synchronizer Token Pattern"
- Specifying the the section called "SameSite Attribute" on your session cookie

Note

Both protections require require that the section called "Safe Methods Must be Idempotent"

Safe Methods Must be Idempotent

In order for [either protection](#) against CSRF to work, the application must ensure that ["safe" HTTP methods are idempotent](#). This means that requests with the HTTP method GET, HEAD, OPTIONS, and TRACE should not change the state of the application.

Synchronizer Token Pattern

The predominant and most comprehensive way to protect against CSRF attacks is to use the [Synchronizer Token Pattern](#). This solution is to ensure that each HTTP request requires, in addition to our session cookie, a secure random generated value called a CSRF token must be present in the HTTP request.

When an HTTP request is submitted, the server must look up the expected CSRF token and compare it against the actual CSRF token in the HTTP request. If the values do not match, the HTTP request should be rejected.

The key to this working is that the actual CSRF token should be in a part of the HTTP request that is not automatically included by the browser. For example, requiring the actual CSRF token in an HTTP parameter or an HTTP header will protect against CSRF attacks. Requiring the actual CSRF token in a cookie does not work because cookies are automatically included in the HTTP request by the browser.

We can relax the expectations to only require the actual CSRF token for each HTTP request that updates state of the application. For that to work, our application must ensure that [safe HTTP methods are idempotent](#). This improves usability since we want to allow linking to our website using links from external sites. Additionally, we do not want to include the random token in HTTP GET as this can cause the tokens to be leaked.

Let's take a look at how [our example](#) would change when using the Synchronizer Token Pattern. Assume the actual CSRF token is required to be in an HTTP parameter named `_csrf`. Our application's transfer form would look like:

```
<form method="post"
  action="/transfer">
<input type="hidden"
  name="_csrf"
  value="4bfd1575-3ad1-4d21-96c7-4ef2d9f86721"/>
<input type="text"
  name="amount"/>
<input type="text"
  name="routingNumber"/>
<input type="hidden"
  name="account"/>
<input type="submit"
  value="Transfer"/>
</form>
```

Example 5.4 Synchronizer Token Form

The form now contains a hidden input with the value of the CSRF token. External sites cannot read the CSRF token since the same origin policy ensures the evil site cannot read the response.

The corresponding HTTP request to transfer money would look like this:

```
POST /transfer HTTP/1.1
Host: bank.example.com
Cookie: JSESSIONID=randomid
Content-Type: application/x-www-form-urlencoded

amount=100.00&routingNumber=1234&account=9876&_csrf=4bfd1575-3ad1-4d21-96c7-4ef2d9f86721
```

Example 5.5 Synchronizer Token request

You will notice that the HTTP request now contains the `_csrf` parameter with a secure random value. The evil website will not be able to provide the correct value for the `_csrf` parameter (which must be explicitly provided on the evil website) and the transfer will fail when the server compares the actual CSRF token to the expected CSRF token.

SameSite Attribute

An emerging way to protect against [CSRF Attacks](#) is to specify the [SameSite Attribute](#) on cookies. A server can specify the `SameSite` attribute when setting a cookie to indicate that the cookie should not be sent when coming from external sites.

Note

Spring Security does not directly control the creation of the session cookie, so it does not provide support for the `SameSite` attribute. [Spring Session](#) provides support for the `SameSite` attribute in servlet based applications. Spring Framework's [CookieWebSessionIdResolver](#) provides out of the box support for the `SameSite` attribute in WebFlux based applications.

An example, HTTP response header with the `SameSite` attribute might look like:

```
Set-Cookie: JSESSIONID=randomid; Domain=bank.example.com; Secure; HttpOnly; SameSite=Lax
```

Example 5.6 SameSite HTTP response

Valid values for the `SameSite` attribute are:

- `Strict` - when specified any request coming from the [same-site](#) will include the cookie. Otherwise, the cookie will not be included in the HTTP request.
- `Lax` - when specified cookies will be sent when coming from the [same-site](#) or when the request comes from top-level navigations and the [method is idempotent](#). Otherwise, the cookie will not be included in the HTTP request.

Let's take a look at how [our example](#) could be protected using the `SameSite` attribute. The bank application can protect against CSRF by specifying the `SameSite` attribute on the session cookie.

With the `SameSite` attribute set on our session cookie, the browser will continue to send the `JSESSIONID` cookie with requests coming from the banking website. However, the browser will no longer send the `JSESSIONID` cookie with a transfer request coming from the evil website. Since the session is no longer present in the transfer request coming from the evil website, the application is protected from the CSRF attack.

There are some important [considerations](#) that one should be aware about when using `SameSite` attribute to protect against CSRF attacks.

Setting the `SameSite` attribute to `Strict` provides a stronger defense but can confuse users. Consider a user that stays logged into a social media site hosted at <https://social.example.com>. The user receives an email at <https://email.example.org> that includes a link to the social media site. If the user clicks on the link, they would rightfully expect to be authenticated to the social media site. However, if the `SameSite` attribute is `Strict` the cookie would not be sent and so the user would not be authenticated.

Note

We could improve the protection and usability of `SameSite` protection against CSRF attacks by implementing [gh-7537](#).

Another obvious consideration is that in order for the `SameSite` attribute to protect users, the browser must support the `SameSite` attribute. Most modern browsers do [support the SameSite attribute](#). However, older browsers that are still in use may not.

For this reason, it is generally recommended to use the `SameSite` attribute as a defense in depth rather than the sole protection against CSRF attacks.

When to use CSRF protection

When should you use CSRF protection? Our recommendation is to use CSRF protection for any request that could be processed by a browser by normal users. If you are only creating a service that is used by non-browser clients, you will likely want to disable CSRF protection.

CSRF protection and JSON

A common question is "do I need to protect JSON requests made by javascript?" The short answer is, it depends. However, you must be very careful as there are CSRF exploits that can impact JSON requests. For example, a malicious user can create a [CSRF with JSON using the following form](#):

```
<form action="https://bank.example.com/transfer" method="post" enctype="text/plain">
  <input name="{\"amount\":100,\"routingNumber\":\"evilsRoutingNumber\",\"account\":\"evilsAccountNumber\",
    \"ignore_me\":\" value='test'}" type='hidden'>
  <input type="submit"
    value="Win Money!" />
</form>
```

Example 5.7 CSRF with JSON form

This will produce the following JSON structure

```
{
  "amount": 100,
  "routingNumber": "evilsRoutingNumber",
  "account": "evilsAccountNumber",
  "ignore_me": "=test"
}
```

Example 5.8 CSRF with JSON request

If an application were not validating the Content-Type, then it would be exposed to this exploit. Depending on the setup, a Spring MVC application that validates the Content-Type could still be exploited by updating the URL suffix to end with `.json` as shown below:

```
<form action="https://bank.example.com/transfer.json" method="post" enctype="text/plain">
  <input name="{\"amount\":100,\"routingNumber\":\"evilsRoutingNumber\",\"account\":\"evilsAccountNumber\",
    \"ignore_me\":\" value='test'}" type='hidden'>
  <input type="submit"
    value="Win Money!" />
</form>
```

Example 5.9 CSRF with JSON Spring MVC form

CSRF and Stateless Browser Applications

What if my application is stateless? That doesn't necessarily mean you are protected. In fact, if a user does not need to perform any actions in the web browser for a given request, they are likely still vulnerable to CSRF attacks.

For example, consider an application that uses a custom cookie that contains all the state within it for authentication instead of the JSESSIONID. When the CSRF attack is made the custom cookie will be sent with the request in the same manner that the JSESSIONID cookie was sent in our previous example. This application will be vulnerable to CSRF attacks.

Applications that use basic authentication are also vulnerable to CSRF attacks. The application is vulnerable since the browser will automatically include the username and password in any requests in the same manner that the JSESSIONID cookie was sent in our previous example.

CSRF Considerations

There are a few special considerations to consider when implementing protection against CSRF attacks.

Logging In

In order to protect against [forging log in requests](#) the log in HTTP request should be protected against CSRF attacks. Protecting against forging log in requests is necessary so that a malicious user cannot read a victim's sensitive information. The attack is executed by:

- A malicious user performs a CSRF log in using the malicious user's credentials. The victim is now authenticated as the malicious user.
- The malicious user then tricks the victim to visit the compromised website and enter sensitive information
- The information is associated to the malicious user's account so the malicious user can log in with their own credentials and view the victim's sensitive information

A possible complication to ensuring log in HTTP requests are protected against CSRF attacks is that the user might experience a session timeout that causes the request to be rejected. A session timeout is surprising to users who do not expect to need to have a session in order to log in. For more information refer to the section called "CSRF and Session Timeouts".

Logging Out

In order to protect against forging log out requests, the log out HTTP request should be protected against CSRF attacks. Protecting against forging log out requests is necessary so a malicious user cannot read a victim's sensitive information. For details on the attack refer to [this blog post](#).

A possible complication to ensuring log out HTTP requests are protected against CSRF attacks is that the user might experience a session timeout that causes the request to be rejected. A session timeout is surprising to users who do not expect to need to have a session in order to log out. For more information refer to the section called "CSRF and Session Timeouts".

CSRF and Session Timeouts

More often than not, the expected CSRF token is stored in the session. This means that as soon as the session expires the server will not find an expected CSRF token and reject the HTTP request. There are a number of options to solve timeouts each of which come with trade offs.

- The best way to mitigate the timeout is by using JavaScript to request a CSRF token on form submission. The form is then updated with the CSRF token and submitted.
- Another option is to have some JavaScript that lets the user know their session is about to expire. The user can click a button to continue and refresh the session.
- Finally, the expected CSRF token could be stored in a cookie. This allows the expected CSRF token to outlive the session.

One might ask why the expected CSRF token isn't stored in a cookie by default. This is because there are known exploits in which headers (i.e. specify the cookies) can be set by another domain. This is the same reason Ruby on Rails [no longer skips CSRF checks when the header X-Requested-With is present](#). See [this webappsec.org thread](#) for details on how to perform the exploit. Another disadvantage is that by removing the state (i.e. the timeout) you lose the ability to forcibly terminate the token if it is compromised.

Multipart (file upload)

Protecting multipart requests (file uploads) from CSRF attacks causes a [chicken and the egg](#) problem. In order to prevent a CSRF attack from occurring, the body of the HTTP request must be read to obtain

actual CSRF token. However, reading the body means that the file will be uploaded which means an external site can upload a file.

There are two options to using CSRF protection with multipart/form-data. Each option has its trade-offs.

- [Place CSRF Token in the Body](#)
- [Place CSRF Token in the URL](#)

Note

Before you integrate Spring Security's CSRF protection with multipart file upload, ensure that you can upload without the CSRF protection first. More information about using multipart forms with Spring can be found within the [1.1.11. Multipart Resolver](#) section of the Spring reference and the [MultipartFilter javadoc](#).

Place CSRF Token in the Body

The first option is to include the actual CSRF token in the body of the request. By placing the CSRF token in the body, the body will be read before authorization is performed. This means that anyone can place temporary files on your server. However, only authorized users will be able to submit a File that is processed by your application. In general, this is the recommended approach because the temporary file upload should have a negligible impact on most servers.

Include CSRF Token in URL

If allowing unauthorized users to upload temporary files is not acceptable, an alternative is to include the expected CSRF token as a query parameter in the action attribute of the form. The disadvantage to this approach is that query parameters can be leaked. More generally, it is considered best practice to place sensitive data within the body or headers to ensure it is not leaked. Additional information can be found in [RFC 2616 Section 15.1.3 Encoding Sensitive Information in URI's](#).

HiddenHttpMethodFilter

In some applications a form parameter can be used to override the HTTP method. For example, the form below could be used to treat the HTTP method as a `delete` rather than a `post`.

```
<form action="/process"
  method="post">
  <!-- ... -->
  <input type="hidden"
    name="_method"
    value="delete"/>
</form>
```

Example 5.10 CSRF Hidden HTTP Method Form

Overriding the HTTP method occurs in a filter. That filter must be placed before Spring Security's support. Note that overriding only happens on a `post`, so this is actually unlikely to cause any real problems. However, it is still best practice to ensure it is placed before Spring Security's filters.

6. Project Modules

In Spring Security 3.0, the codebase was sub-divided into separate jars which more clearly separate different functionality areas and third-party dependencies. If you use Maven to build your project, these are the modules you should add to your `pom.xml`. Even if you do not use Maven, we recommend that you consult the `pom.xml` files to get an idea of third-party dependencies and versions. Another good idea is to examine the libraries that are included in the sample applications.

6.1 Core — `spring-security-core.jar`

This module contains core authentication and access-control classes and interfaces, remoting support, and basic provisioning APIs. It is required by any application that uses Spring Security. It supports standalone applications, remote clients, method (service layer) security, and JDBC user provisioning. It contains the following top-level packages:

- `org.springframework.security.core`
- `org.springframework.security.access`
- `org.springframework.security.authentication`
- `org.springframework.security.provisioning`

6.2 Remoting — `spring-security-remoting.jar`

This module provides integration with Spring Remoting. You do not need this unless you are writing a remote client that uses Spring Remoting. The main package is `org.springframework.security.remoting`.

6.3 Web — `spring-security-web.jar`

This module contains filters and related web-security infrastructure code. It contains anything with a servlet API dependency. You need it if you require Spring Security web authentication services and URL-based access-control. The main package is `org.springframework.security.web`.

6.4 Config — `spring-security-config.jar`

This module contains the security namespace parsing code and Java configuration code. You need it if you use the Spring Security XML namespace for configuration or Spring Security's Java Configuration support. The main package is `org.springframework.security.config`. None of the classes are intended for direct use in an application.

6.5 LDAP — `spring-security-ldap.jar`

This module provides LDAP authentication and provisioning code. It is required if you need to use LDAP authentication or manage LDAP user entries. The top-level package is `org.springframework.security.ldap`.

6.6 OAuth 2.0 Core — `spring-security-oauth2-core.jar`

`spring-security-oauth2-core.jar` contains core classes and interfaces that provide support for the OAuth 2.0 Authorization Framework and for OpenID Connect Core 1.0. It is required by applications

that use OAuth 2.0 or OpenID Connect Core 1.0, such as client, resource server, and authorization server. The top-level package is `org.springframework.security.oauth2.core`.

6.7 OAuth 2.0 Client — `spring-security-oauth2-client.jar`

`spring-security-oauth2-client.jar` contains Spring Security's client support for OAuth 2.0 Authorization Framework and OpenID Connect Core 1.0. It is required by applications that use OAuth 2.0 Login or OAuth Client support. The top-level package is `org.springframework.security.oauth2.client`.

6.8 OAuth 2.0 JOSE — `spring-security-oauth2-jose.jar`

`spring-security-oauth2-jose.jar` contains Spring Security's support for the JOSE (Javascript Object Signing and Encryption) framework. The JOSE framework is intended to provide a method to securely transfer claims between parties. It is built from a collection of specifications:

- JSON Web Token (JWT)
- JSON Web Signature (JWS)
- JSON Web Encryption (JWE)
- JSON Web Key (JWK)

It contains the following top-level packages:

- `org.springframework.security.oauth2.jwt`
- `org.springframework.security.oauth2.jose`

6.9 OAuth 2.0 Resource Server — `spring-security-oauth2-resource-server.jar`

`spring-security-oauth2-resource-server.jar` contains Spring Security's support for OAuth 2.0 Resource Servers. It is used to protect APIs via OAuth 2.0 Bearer Tokens. The top-level package is `org.springframework.security.oauth2.server.resource`.

6.10 ACL — `spring-security-acl.jar`

This module contains a specialized domain object ACL implementation. It is used to apply security to specific domain object instances within your application. The top-level package is `org.springframework.security.acls`.

6.11 CAS — `spring-security-cas.jar`

This module contains Spring Security's CAS client integration. You should use it if you want to use Spring Security web authentication with a CAS single sign-on server. The top-level package is `org.springframework.security.cas`.

6.12 OpenID — `spring-security-openid.jar`

This module contains OpenID web authentication support. It is used to authenticate users against an external OpenID server. The top-level package is `org.springframework.security.openid`. It requires OpenID4Java.

6.13 Test — `spring-security-test.jar`

This module contains support for testing with Spring Security.

7. Samples

Spring Security includes many [samples](#) applications.

Part II. Servlet Applications

Spring Security integrates with the Servlet Container by using a standard Servlet `Filter`. This means it works with any application that runs in a Servlet Container. More concretely, you do not need to use Spring in your Servlet-based application to take advantage of Spring Security.

8. Hello Spring Security

This section covers a minimal Spring Security application that uses [Spring Boot](#), [Java Configuration](#), or [XML Configuration](#).

8.1 Hello Spring Security (Boot)

This section covers the minimum setup for how to use Spring Security with Spring Boot. For how to use Spring Security with Java Configuration, see Section 8.2, “Hello Spring Security (Java Configuration)”. For how to use Spring Security with XML Configuration, see Section 8.3, “Hello Spring Security (XML)”.

Note

The completed application can be found at [samples/boot/helloworld](#)

Updating Dependencies

The only step you need to do is update the dependencies by using [Maven](#) or [Gradle](#). For your convenience, you can download a minimal Spring Boot + Spring Security application by [clicking here](#).

Starting Hello Spring Security Boot

You can now [run the Spring Boot application](#) by using the Maven Plugin's `run` goal. The following example shows how to do so (and the beginning of the output from doing so):

```
$ ./mvn spring-boot:run
...
INFO 23689 --- [ restartedMain] .s.s.UserDetailsServiceAutoConfiguration :
Using generated security password: 8e557245-73e2-4286-969a-ff57fe326336
...
```

Example 8.1 Running Spring Boot Application

Spring Boot Auto Configuration

Spring Boot automatically:

- Enables Spring Security's default configuration, which creates a servlet `Filter` as a bean named `springSecurityFilterChain`. This bean is responsible for all the security (protecting the application URLs, validating submitted username and passwords, redirecting to the log in form, and so on) within your application.
- Creates a `UserDetailsService` bean with a username of `user` and a randomly generated password that is logged to the console.
- Registers the `Filter` with a bean named `springSecurityFilterChain` with the `Servlet` container for every request.

Spring Boot is not configuring much, but it does a lot. A summary of the features follows:

- Require an authenticated user for any interaction with the application
- Generate a default login form for you

- Let the user with a username of `user` and a password that is logged to the console to authenticate with form-based authentication (in the preceding example, the password is `8e557245-73e2-4286-969a-ff57fe326336`)
- Protects the password storage with BCrypt
- Lets the user log out
- [CSRF attack](#) prevention
- [Session Fixation](#) protection
- Security Header integration
 - [HTTP Strict Transport Security](#) for secure requests
 - [X-Content-Type-Options](#) integration
 - Cache Control (can be overridden later by your application to allow caching of your static resources)
 - [X-XSS-Protection](#) integration
 - X-Frame-Options integration to help prevent [Clickjacking](#)
- Integrate with the following Servlet API methods:
 - [HttpServletRequest#getRemoteUser\(\)](#)
 - [HttpServletRequest.html#getUserPrincipal\(\)](#)
 - [HttpServletRequest.html#isUserInRole\(java.lang.String\)](#)
 - [HttpServletRequest.html#login\(java.lang.String, java.lang.String\)](#)
 - [HttpServletRequest.html#logout\(\)](#)

8.2 Hello Spring Security (Java Configuration)

This section covers how to use Spring Security with Java Configuration. For how to use Spring Security with XML configuration, see Section 8.3, “Hello Spring Security (XML)”. For how to use Spring Security with Spring Boot configuration, see Section 8.1, “Hello Spring Security (Boot)”.

Note

You can find the completed application at [samples/javaconfig/helloworld](#).

Updating Dependencies

The first step is to update the dependencies by using [Maven](#) or [Gradle](#).

Minimal `@EnableWebSecurity` Configuration

The first step is to create our Spring Security Java configuration. The configuration creates a servlet `Filter` (known as the `springSecurityFilterChain`), which is responsible for all the security features (protecting the application URLs, validating submitted username and passwords, redirecting to

the log in form, and so on) within your application. The following example shows the most basic example of a Spring Security Java Configuration:

```
import org.springframework.context.annotation.*;
import org.springframework.security.config.annotation.web.configuration.*;
import org.springframework.security.core.userdetails.*;
import org.springframework.security.provisioning.*;

@EnableWebSecurity
public class WebSecurityConfig {

    // @formatter:off
    @Bean
    public UserDetailsService userDetailsService() {
        UserDetails user = User.withDefaultPasswordEncoder()
            .username("user")
            .password("password")
            .roles("USER")
            .build();
        return new InMemoryUserDetailsManager(user);
    }
    // @formatter:on
}
```

Example 8.2 WebSecurity.java

There really is not much to this configuration, but it does a lot. A summary of the features follows:

- Require an authenticated user for any interaction with the application
- Generate a default login form for you
- Lets the user with a username of `user` and a password of `password` authenticate with form-based authentication
- Protects the password storage with BCrypt
- Lets the user log out
- [CSRF attack](#) prevention
- [Session Fixation](#) protection
- Security Header integration
 - [HTTP Strict Transport Security](#) for secure requests
 - [X-Content-Type-Options](#) integration
 - Cache Control (can be overridden later by your application to allow caching of your static resources)
 - [X-XSS-Protection](#) integration
 - X-Frame-Options integration to help prevent [Clickjacking](#)
- Integrate with the following Servlet API methods:
 - [HttpServletRequest#getRemoteUser\(\)](#)
 - [HttpServletRequest.html#getUserPrincipal\(\)](#)
 - [HttpServletRequest.html#isUserInRole\(java.lang.String\)](#)

- [HttpServletRequest.html#login\(java.lang.String, java.lang.String\)](#)
- [HttpServletRequest.html#logout\(\)](#)

Using AbstractSecurityWebApplicationInitializer

The next step is to register the `springSecurityFilterChain` with the war. Spring Security provides a base class (`AbstractSecurityWebApplicationInitializer`) that leverages [Spring's WebApplicationInitializer support](#).

The following example shows an example configuration:

```
import org.springframework.security.web.context.*;

public class SecurityInitializer
    extends AbstractSecurityWebApplicationInitializer {

    public SecurityInitializer() {
        super(WebSecurityConfig.class);
    }
}
```

Example 8.3 SecurityInitializer.java

The `SecurityInitializer` does the following things:

- Adds a `ContextLoaderListener` that loads the [WebSecurityConfig](#).
- Finds the bean of type `Filter` named `springSecurityFilterChain` and registers it to process every URL in the application.

Note

If you are integrating with a Spring MVC application, be sure to configure the `DispatcherServlet` to load the configuration from the root `ApplicationContext`. The following example shows how to do so:

```
public class MvcInitializer extends
    AbstractAnnotationConfigDispatcherServletInitializer {

    // the Root Config is registered in SecurityInitializer
    @Override
    protected Class<?>[] getRootConfigClasses() {
        return null;
    }

    // the Spring MVC configuration should be added to SecurityInitializer constructor
    // i.e.
    // super(MvcConfig.class, WebSecurityConfig.class);
    @Override
    protected Class<?>[] getServletConfigClasses() {
        return null;
    }

    @Override
    protected String[] getServletMappings() {
        return new String[] { "/" };
    }
}
```

Example 8.4 MvcInitializer.java

8.3 Hello Spring Security (XML)

This section covers how to use Spring Security with XML Configuration. For how to use Spring Security with Java configuration, see Section 8.2, “Hello Spring Security (Java Configuration)”. For how to use Spring Security with Spring Boot configuration, see Section 8.1, “Hello Spring Security (Boot)”.

Updating Dependencies

The first step is to update the dependencies by using [Maven](#) or [Gradle](#).

Minimal `<http>` Configuration

In this section, we discuss how to use Spring Security with XML Configuration.

Note

The completed application can be found at [samples/xml/helloworld](#)

The first step is to create our Spring Security XML Configuration. The configuration creates a Servlet Filter (known as the `springSecurityFilterChain`), which is responsible for all the security (protecting the application URLs, validating submitted username and passwords, redirecting to the log in form, and so on) within your application. The following example shows the most basic example of a Spring Security XML Configuration:

```
<b:beans xmlns="http://www.springframework.org/schema/security"
  xmlns:b="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.springframework.org/schema/beans https://
www.springframework.org/schema/beans/spring-beans.xsd
  http://www.springframework.org/schema/security https://www.springframework.org/
schema/security/spring-security.xsd">
  <http />

  <user-service>
    <user name="user" password="{noop}password" authorities="ROLE_USER" />
  </user-service>
</b:beans>
```

Example 8.5 `src/main/webapp/WEB-INF/spring/security.xml`

There really is not much to this configuration, but it does a lot. A summary of the features follows:

- Require an authenticated user for any interaction with the application
- Generate a default login form for you
- Lets the user with a username of `user` and a password of `password` authenticate with form-based authentication
- Protects the password storage with BCrypt
- Lets the user to log out
- [CSRF attack](#) prevention
- [Session Fixation](#) protection

- Security Header integration
 - [HTTP Strict Transport Security](#) for secure requests
 - [X-Content-Type-Options](#) integration
 - Cache Control (can be overridden later by your application to allow caching of your static resources)
 - [X-XSS-Protection](#) integration
 - X-Frame-Options integration to help prevent [Clickjacking](#)
- Integrate with the following Servlet API methods:
 - [HttpServletRequest#getRemoteUser\(\)](#)
 - [HttpServletRequest.html#getUserPrincipal\(\)](#)
 - [HttpServletRequest.html#isUserInRole\(java.lang.String\)](#)
 - [HttpServletRequest.html#login\(java.lang.String, java.lang.String\)](#)
 - [HttpServletRequest.html#logout\(\)](#)

web.xml Configuration

The next step is to ensure that our Security configuration is being read in. To do so, we need to ensure a `ContextLoaderListener` is registered and the `contextConfigLocation` is including the configuration. The following example shows how to do so:

```

<?xml version="1.0" encoding="UTF-8"?>
<web-app version="3.0" xmlns="http://java.sun.com/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
  http://java.sun.com/xml/ns/javaee/web-app_3_0.xsd">

  <!--
    Loads the Spring configurations from contextConfigLocation
  -->
  <listener>
    <listener-class>org.springframework.web.context.ContextLoaderListener</listener-class>
  </listener>

  <!--
    The locations of the Spring Configuration. In this case, all configuration is
    in /WEB-INF/spring/
  -->
  <context-param>
    <param-name>contextConfigLocation</param-name>
    <param-value>
      /WEB-INF/spring/*.xml
    </param-value>
  </context-param>

  <!--
    DelegatingFilterProxy looks for a Spring bean by the name of filter (springSecurityFilterChain)
    and delegates
    all work to that Bean. This is how the Servlet Container can a Spring Bean to act as a Servlet
    Filter.
  -->
  <filter>
    <filter-name>springSecurityFilterChain</filter-name>
    <filter-class>org.springframework.web.filter.DelegatingFilterProxy</filter-class>
  </filter>
  <filter-mapping>
    <filter-name>springSecurityFilterChain</filter-name>
    <url-pattern>/*</url-pattern>
  </filter-mapping>

</web-app>

```

Example 8.6 `src/main/webapp/WEB-INF/web.xml`

Note

If you integrate with an existing Spring MVC application, be sure to configure the `DispatcherServlet` to load the configuration from the root `ApplicationContext`. The following example shows how to do so:

`src/main/webapp/WEB-INF/web.xml`.

```

<servlet>
  <servlet-name>spring</servlet-name>
  <servlet-class>org.springframework.web.servlet.DispatcherServlet</servlet-class>
  <!-- Load Spring MVC configuration from root ApplicationContext (context-param from above) -->
  >
  <init-param>
    <param-name>contextConfigLocation</param-name>
    <param-value></param-value>
  </init-param>
</servlet>

<servlet-mapping>
  <servlet-name>spring</servlet-name>
  <url-pattern>/*</url-pattern>
</servlet-mapping>

```

9. Architecture and Implementation

Once you are familiar with setting up and running some namespace-configuration based applications, you may wish to develop more of an understanding of how the framework actually works behind the namespace facade. Like most software, Spring Security has certain central interfaces, classes and conceptual abstractions that are commonly used throughout the framework. In this part of the reference guide we will look at some of these and see how they work together to support authentication and access-control within Spring Security.

9.1 Technical Overview

Runtime Environment

Spring Security 5.2.1.RELEASE requires a Java 8 Runtime Environment or higher. As Spring Security aims to operate in a self-contained manner, there is no need to place any special configuration files into your Java Runtime Environment. In particular, there is no need to configure a special Java Authentication and Authorization Service (JAAS) policy file or place Spring Security into common classpath locations.

Similarly, if you are using an EJB Container or Servlet Container there is no need to put any special configuration files anywhere, nor include Spring Security in a server classloader. All the required files will be contained within your application.

This design offers maximum deployment time flexibility, as you can simply copy your target artifact (be it a JAR, WAR or EAR) from one system to another and it will immediately work.

Core Components

As of Spring Security 3.0, the contents of the `spring-security-core` jar were stripped down to the bare minimum. It no longer contains any code related to web-application security, LDAP or namespace configuration. We'll take a look here at some of the Java types that you'll find in the core module. They represent the building blocks of the framework, so if you ever need to go beyond a simple namespace configuration then it's important that you understand what they are, even if you don't actually need to interact with them directly.

SecurityContextHolder, SecurityContext and Authentication Objects

The most fundamental object is `SecurityContextHolder`. This is where we store details of the present security context of the application, which includes details of the principal currently using the application. By default the `SecurityContextHolder` uses a `ThreadLocal` to store these details, which means that the security context is always available to methods in the same thread of execution, even if the security context is not explicitly passed around as an argument to those methods. Using a `ThreadLocal` in this way is quite safe if care is taken to clear the thread after the present principal's request is processed. Of course, Spring Security takes care of this for you automatically so there is no need to worry about it.

Some applications aren't entirely suitable for using a `ThreadLocal`, because of the specific way they work with threads. For example, a Swing client might want all threads in a Java Virtual Machine to use the same security context. `SecurityContextHolder` can be configured with a strategy on startup to specify how you would like the context to be stored. For a standalone application you would use the `SecurityContextHolder.MODE_GLOBAL` strategy. Other applications might want to have threads spawned by the secure thread also assume the same security identity. This is

achieved by using `SecurityContextHolder.MODE_INHERITABLETHREADLOCAL`. You can change the mode from the default `SecurityContextHolder.MODE_THREADLOCAL` in two ways. The first is to set a system property, the second is to call a static method on `SecurityContextHolder`. Most applications won't need to change from the default, but if you do, take a look at the JavaDoc for `SecurityContextHolder` to learn more.

Obtaining information about the current user

Inside the `SecurityContextHolder` we store details of the principal currently interacting with the application. Spring Security uses an `Authentication` object to represent this information. You won't normally need to create an `Authentication` object yourself, but it is fairly common for users to query the `Authentication` object. You can use the following code block - from anywhere in your application - to obtain the name of the currently authenticated user, for example:

```
Object principal = SecurityContextHolder.getContext().getAuthentication().getPrincipal();

if (principal instanceof UserDetails) {
    String username = ((UserDetails)principal).getUsername();
} else {
    String username = principal.toString();
}
```

The object returned by the call to `getContext()` is an instance of the `SecurityContext` interface. This is the object that is kept in thread-local storage. As we'll see below, most authentication mechanisms within Spring Security return an instance of `UserDetails` as the principal.

The UserDetailsService

Another item to note from the above code fragment is that you can obtain a principal from the `Authentication` object. The principal is just an `Object`. Most of the time this can be cast into a `UserDetails` object. `UserDetails` is a core interface in Spring Security. It represents a principal, but in an extensible and application-specific way. Think of `UserDetails` as the adapter between your own user database and what Spring Security needs inside the `SecurityContextHolder`. Being a representation of something from your own user database, quite often you will cast the `UserDetails` to the original object that your application provided, so you can call business-specific methods (like `getEmail()`, `getEmployeeNumber()` and so on).

By now you're probably wondering, so when do I provide a `UserDetails` object? How do I do that? I thought you said this thing was declarative and I didn't need to write any Java code - what gives? The short answer is that there is a special interface called `UserDetailsService`. The only method on this interface accepts a `String`-based username argument and returns a `UserDetails`:

```
UserDetails loadUserByUsername(String username) throws UsernameNotFoundException;
```

This is the most common approach to loading information for a user within Spring Security and you will see it used throughout the framework whenever information on a user is required.

On successful authentication, `UserDetails` is used to build the `Authentication` object that is stored in the `SecurityContextHolder` (more on this [below](#)). The good news is that we provide a number of `UserDetailsService` implementations, including one that uses an in-memory map (`InMemoryDaoImpl`) and another that uses JDBC (`JdbcDaoImpl`). Most users tend to write their own, though, with their implementations often simply sitting on top of an existing Data Access Object (DAO) that represents their employees, customers, or other users of the application. Remember the advantage that whatever your `UserDetailsService` returns can always be obtained from the `SecurityContextHolder` using the above code fragment.

Note

There is often some confusion about `UserDetailsService`. It is purely a DAO for user data and performs no other function other than to supply that data to other components within the framework. In particular, it *does not* authenticate the user, which is done by the `AuthenticationManager`. In many cases it makes more sense to [implement `AuthenticationProvider`](#) directly if you require a custom authentication process.

GrantedAuthority

Besides the principal, another important method provided by `Authentication` is `getAuthorities()`. This method provides an array of `GrantedAuthority` objects. A `GrantedAuthority` is, not surprisingly, an authority that is granted to the principal. Such authorities are usually "roles", such as `ROLE_ADMINISTRATOR` or `ROLE_HR_SUPERVISOR`. These roles are later on configured for web authorization, method authorization and domain object authorization. Other parts of Spring Security are capable of interpreting these authorities, and expect them to be present. `GrantedAuthority` objects are usually loaded by the `UserDetailsService`.

Usually the `GrantedAuthority` objects are application-wide permissions. They are not specific to a given domain object. Thus, you wouldn't likely have a `GrantedAuthority` to represent a permission to `Employee` object number 54, because if there are thousands of such authorities you would quickly run out of memory (or, at the very least, cause the application to take a long time to authenticate a user). Of course, Spring Security is expressly designed to handle this common requirement, but you'd instead use the project's domain object security capabilities for this purpose.

Summary

Just to recap, the major building blocks of Spring Security that we've seen so far are:

- `SecurityContextHolder`, to provide access to the `SecurityContext`.
- `SecurityContext`, to hold the `Authentication` and possibly request-specific security information.
- `Authentication`, to represent the principal in a Spring Security-specific manner.
- `GrantedAuthority`, to reflect the application-wide permissions granted to a principal.
- `UserDetails`, to provide the necessary information to build an `Authentication` object from your application's DAOs or other source of security data.
- `UserDetailsService`, to create a `UserDetails` when passed in a `String`-based username (or certificate ID or the like).

Now that you've gained an understanding of these repeatedly-used components, let's take a closer look at the process of authentication.

Authentication

Spring Security can participate in many different authentication environments. While we recommend people use Spring Security for authentication and not integrate with existing Container Managed Authentication, it is nevertheless supported - as is integrating with your own proprietary authentication system.

What is authentication in Spring Security?

Let's consider a standard authentication scenario that everyone is familiar with.

1. A user is prompted to log in with a username and password.
2. The system (successfully) verifies that the password is correct for the username.
3. The context information for that user is obtained (their list of roles and so on).
4. A security context is established for the user
5. The user proceeds, potentially to perform some operation which is potentially protected by an access control mechanism which checks the required permissions for the operation against the current security context information.

The first four items constitute the authentication process so we'll take a look at how these take place within Spring Security.

1. The username and password are obtained and combined into an instance of `UsernamePasswordAuthenticationToken` (an instance of the `Authentication` interface, which we saw earlier).
2. The token is passed to an instance of `AuthenticationManager` for validation.
3. The `AuthenticationManager` returns a fully populated `Authentication` instance on successful authentication.
4. The security context is established by calling `SecurityContextHolder.getContext().setAuthentication(...)`, passing in the returned authentication object.

From that point on, the user is considered to be authenticated. Let's look at some code as an example.

```

import org.springframework.security.authentication.*;
import org.springframework.security.core.*;
import org.springframework.security.core.authority.SimpleGrantedAuthority;
import org.springframework.security.core.context.SecurityContextHolder;

public class AuthenticationExample {
    private static AuthenticationManager am = new SampleAuthenticationManager();

    public static void main(String[] args) throws Exception {
        BufferedReader in = new BufferedReader(new InputStreamReader(System.in));

        while(true) {
            System.out.println("Please enter your username:");
            String name = in.readLine();
            System.out.println("Please enter your password:");
            String password = in.readLine();
            try {
                Authentication request = new UsernamePasswordAuthenticationToken(name, password);
                Authentication result = am.authenticate(request);
                SecurityContextHolder.getContext().setAuthentication(result);
                break;
            } catch(AuthenticationException e) {
                System.out.println("Authentication failed: " + e.getMessage());
            }
        }
        System.out.println("Successfully authenticated. Security context contains: " +
            SecurityContextHolder.getContext().getAuthentication());
    }
}

class SampleAuthenticationManager implements AuthenticationManager {
    static final List<GrantedAuthority> AUTHORITIES = new ArrayList<GrantedAuthority>();

    static {
        AUTHORITIES.add(new SimpleGrantedAuthority("ROLE_USER"));
    }

    public Authentication authenticate(Authentication auth) throws AuthenticationException {
        if (auth.getName().equals(auth.getCredentials())) {
            return new UsernamePasswordAuthenticationToken(auth.getName(),
                auth.getCredentials(), AUTHORITIES);
        }
        throw new BadCredentialsException("Bad Credentials");
    }
}

```

Here we have written a little program that asks the user to enter a username and password and performs the above sequence. The `AuthenticationManager` which we've implemented here will authenticate any user whose username and password are the same. It assigns a single role to every user. The output from the above will be something like:

```

Please enter your username:
bob
Please enter your password:
password
Authentication failed: Bad Credentials
Please enter your username:
bob
Please enter your password:
bob
Successfully authenticated. Security context contains: \
org.springframework.security.authentication.UsernamePasswordAuthenticationToken@441d0230: \
Principal: bob; Password: [PROTECTED]; \
Authenticated: true; Details: null; \
Granted Authorities: ROLE_USER

```

Note that you don't normally need to write any code like this. The process will normally occur internally, in a web authentication filter for example. We've just included the code here to show that the question

of what actually constitutes authentication in Spring Security has quite a simple answer. A user is authenticated when the `SecurityContextHolder` contains a fully populated `Authentication` object.

Setting the `SecurityContextHolder` Contents Directly

In fact, Spring Security doesn't mind how you put the `Authentication` object inside the `SecurityContextHolder`. The only critical requirement is that the `SecurityContextHolder` contains an `Authentication` which represents a principal before the `AbstractSecurityInterceptor` (which we'll see more about later) needs to authorize a user operation.

You can (and many users do) write their own filters or MVC controllers to provide interoperability with authentication systems that are not based on Spring Security. For example, you might be using Container-Managed Authentication which makes the current user available from a `ThreadLocal` or JNDI location. Or you might work for a company that has a legacy proprietary authentication system, which is a corporate "standard" over which you have little control. In situations like this it's quite easy to get Spring Security to work, and still provide authorization capabilities. All you need to do is write a filter (or equivalent) that reads the third-party user information from a location, build a Spring Security-specific `Authentication` object, and put it into the `SecurityContextHolder`. In this case you also need to think about things which are normally taken care of automatically by the built-in authentication infrastructure. For example, you might need to pre-emptively create an HTTP session to [cache the context between requests](#), before you write the response to the client ¹.

If you're wondering how the `AuthenticationManager` is implemented in a real world example, we'll look at that in the [core services chapter](#).

Authentication in a Web Application

Now let's explore the situation where you are using Spring Security in a web application (without `web.xml` security enabled). How is a user authenticated and the security context established?

Consider a typical web application's authentication process:

1. You visit the home page, and click on a link.
2. A request goes to the server, and the server decides that you've asked for a protected resource.
3. As you're not presently authenticated, the server sends back a response indicating that you must authenticate. The response will either be an HTTP response code, or a redirect to a particular web page.
4. Depending on the authentication mechanism, your browser will either redirect to the specific web page so that you can fill out the form, or the browser will somehow retrieve your identity (via a BASIC authentication dialogue box, a cookie, a X.509 certificate etc.).
5. The browser will send back a response to the server. This will either be an HTTP POST containing the contents of the form that you filled out, or an HTTP header containing your authentication details.
6. Next the server will decide whether or not the presented credentials are valid. If they're valid, the next step will happen. If they're invalid, usually your browser will be asked to try again (so you return to step two above).

¹It isn't possible to create a session once the response has been committed.

7. The original request that you made to cause the authentication process will be retried. Hopefully you've authenticated with sufficient granted authorities to access the protected resource. If you have sufficient access, the request will be successful. Otherwise, you'll receive back an HTTP error code 403, which means "forbidden".

Spring Security has distinct classes responsible for most of the steps described above. The main participants (in the order that they are used) are the `ExceptionTranslationFilter`, an `AuthenticationEntryPoint` and an "authentication mechanism", which is responsible for calling the `AuthenticationManager` which we saw in the previous section.

ExceptionTranslationFilter

`ExceptionTranslationFilter` is a Spring Security filter that has responsibility for detecting any Spring Security exceptions that are thrown. Such exceptions will generally be thrown by an `AbstractSecurityInterceptor`, which is the main provider of authorization services. We will discuss `AbstractSecurityInterceptor` in the next section, but for now we just need to know that it produces Java exceptions and knows nothing about HTTP or how to go about authenticating a principal. Instead the `ExceptionTranslationFilter` offers this service, with specific responsibility for either returning error code 403 (if the principal has been authenticated and therefore simply lacks sufficient access - as per step seven above), or launching an `AuthenticationEntryPoint` (if the principal has not been authenticated and therefore we need to go commence step three).

AuthenticationEntryPoint

The `AuthenticationEntryPoint` is responsible for step three in the above list. As you can imagine, each web application will have a default authentication strategy (well, this can be configured like nearly everything else in Spring Security, but let's keep it simple for now). Each major authentication system will have its own `AuthenticationEntryPoint` implementation, which typically performs one of the actions described in step 3.

Authentication Mechanism

Once your browser submits your authentication credentials (either as an HTTP form post or HTTP header) there needs to be something on the server that "collects" these authentication details. By now we're at step six in the above list. In Spring Security we have a special name for the function of collecting authentication details from a user agent (usually a web browser), referring to it as the "authentication mechanism". Examples are form-base login and Basic authentication. Once the authentication details have been collected from the user agent, an `Authentication` "request" object is built and then presented to the `AuthenticationManager`.

After the authentication mechanism receives back the fully-populated `Authentication` object, it will deem the request valid, put the `Authentication` into the `SecurityContextHolder`, and cause the original request to be retried (step seven above). If, on the other hand, the `AuthenticationManager` rejected the request, the authentication mechanism will ask the user agent to retry (step two above).

Storing the SecurityContext between requests

Depending on the type of application, there may need to be a strategy in place to store the security context between user operations. In a typical web application, a user logs in once and is subsequently identified by their session id. The server caches the principal information for the duration session. In Spring Security, the responsibility for storing the `SecurityContext` between requests falls to the `SecurityContextPersistenceFilter`, which by default stores the context as an `HttpSession` attribute between HTTP requests. It restores the context to the `SecurityContextHolder` for each

request and, crucially, clears the `SecurityContextHolder` when the request completes. You shouldn't interact directly with the `HttpSession` for security purposes. There is simply no justification for doing so - always use the `SecurityContextHolder` instead.

Many other types of application (for example, a stateless RESTful web service) do not use HTTP sessions and will re-authenticate on every request. However, it is still important that the `SecurityContextPersistenceFilter` is included in the chain to make sure that the `SecurityContextHolder` is cleared after each request.

Note

In an application which receives concurrent requests in a single session, the same `SecurityContext` instance will be shared between threads. Even though a `ThreadLocal` is being used, it is the same instance that is retrieved from the `HttpSession` for each thread. This has implications if you wish to temporarily change the context under which a thread is running. If you just use `SecurityContextHolder.getContext()`, and call `setAuthentication(anAuthentication)` on the returned context object, then the `Authentication` object will change in *all* concurrent threads which share the same `SecurityContext` instance. You can customize the behaviour of `SecurityContextPersistenceFilter` to create a completely new `SecurityContext` for each request, preventing changes in one thread from affecting another. Alternatively you can create a new instance just at the point where you temporarily change the context. The method `SecurityContextHolder.createEmptyContext()` always returns a new context instance.

Access-Control (Authorization) in Spring Security

The main interface responsible for making access-control decisions in Spring Security is the `AccessDecisionManager`. It has a `decide` method which takes an `Authentication` object representing the principal requesting access, a "secure object" (see below) and a list of security metadata attributes which apply for the object (such as a list of roles which are required for access to be granted).

Security and AOP Advice

If you're familiar with AOP, you'd be aware there are different types of advice available: before, after, throws and around. An around advice is very useful, because an advisor can elect whether or not to proceed with a method invocation, whether or not to modify the response, and whether or not to throw an exception. Spring Security provides an around advice for method invocations as well as web requests. We achieve an around advice for method invocations using Spring's standard AOP support and we achieve an around advice for web requests using a standard `Filter`.

For those not familiar with AOP, the key point to understand is that Spring Security can help you protect method invocations as well as web requests. Most people are interested in securing method invocations on their services layer. This is because the services layer is where most business logic resides in current-generation Java EE applications. If you just need to secure method invocations in the services layer, Spring's standard AOP will be adequate. If you need to secure domain objects directly, you will likely find that AspectJ is worth considering.

You can elect to perform method authorization using AspectJ or Spring AOP, or you can elect to perform web request authorization using filters. You can use zero, one, two or three of these approaches together. The mainstream usage pattern is to perform some web request authorization, coupled with some Spring AOP method invocation authorization on the services layer.

Secure Objects and the `AbstractSecurityInterceptor`

So what *is* a "secure object" anyway? Spring Security uses the term to refer to any object that can have security (such as an authorization decision) applied to it. The most common examples are method invocations and web requests.

Each supported secure object type has its own interceptor class, which is a subclass of `AbstractSecurityInterceptor`. Importantly, by the time the `AbstractSecurityInterceptor` is called, the `SecurityContextHolder` will contain a valid `Authentication` if the principal has been authenticated.

`AbstractSecurityInterceptor` provides a consistent workflow for handling secure object requests, typically:

1. Look up the "configuration attributes" associated with the present request
2. Submitting the secure object, current `Authentication` and configuration attributes to the `AccessDecisionManager` for an authorization decision
3. Optionally change the `Authentication` under which the invocation takes place
4. Allow the secure object invocation to proceed (assuming access was granted)
5. Call the `AfterInvocationManager` if configured, once the invocation has returned. If the invocation raised an exception, the `AfterInvocationManager` will not be invoked.

What are Configuration Attributes?

A "configuration attribute" can be thought of as a `String` that has special meaning to the classes used by `AbstractSecurityInterceptor`. They are represented by the interface `ConfigAttribute` within the framework. They may be simple role names or have more complex meaning, depending on the how sophisticated the `AccessDecisionManager` implementation is. The `AbstractSecurityInterceptor` is configured with a `SecurityMetadataSource` which it uses to look up the attributes for a secure object. Usually this configuration will be hidden from the user. Configuration attributes will be entered as annotations on secured methods or as access attributes on secured URLs. For example, when we saw something like `<intercept-url pattern='/secure/**' access='ROLE_A,ROLE_B' />` in the namespace introduction, this is saying that the configuration attributes `ROLE_A` and `ROLE_B` apply to web requests matching the given pattern. In practice, with the default `AccessDecisionManager` configuration, this means that anyone who has a `GrantedAuthority` matching either of these two attributes will be allowed access. Strictly speaking though, they are just attributes and the interpretation is dependent on the `AccessDecisionManager` implementation. The use of the prefix `ROLE_` is a marker to indicate that these attributes are roles and should be consumed by Spring Security's `RoleVoter`. This is only relevant when a voter-based `AccessDecisionManager` is in use. We'll see how the `AccessDecisionManager` is implemented in the [authorization chapter](#).

`RunAsManager`

Assuming `AccessDecisionManager` decides to allow the request, the `AbstractSecurityInterceptor` will normally just proceed with the request. Having said that, on rare occasions users may want to replace the `Authentication` inside the `SecurityContext` with a different `Authentication`, which is handled by the `AccessDecisionManager` calling a `RunAsManager`. This might be useful in reasonably unusual situations, such as if a services layer

method needs to call a remote system and present a different identity. Because Spring Security automatically propagates security identity from one server to another (assuming you're using a properly-configured RMI or HttpInvoker remoting protocol client), this may be useful.

AfterInvocationManager

Following the secure object invocation proceeding and then returning - which may mean a method invocation completing or a filter chain proceeding - the `AbstractSecurityInterceptor` gets one final chance to handle the invocation. At this stage the `AbstractSecurityInterceptor` is interested in possibly modifying the return object. We might want this to happen because an authorization decision couldn't be made "on the way in" to a secure object invocation. Being highly pluggable, `AbstractSecurityInterceptor` will pass control to an `AfterInvocationManager` to actually modify the object if needed. This class can even entirely replace the object, or throw an exception, or not change it in any way as it chooses. The after-invocation checks will only be executed if the invocation is successful. If an exception occurs, the additional checks will be skipped.

`AbstractSecurityInterceptor` and its related objects are shown in Figure 9.1, "Security interceptors and the "secure object" model"




Figure 9.1. Security interceptors and the "secure object" model

Extending the Secure Object Model

Only developers contemplating an entirely new way of intercepting and authorizing requests would need to use secure objects directly. For example, it would be possible to build a new secure object to secure calls to a messaging system. Anything that requires security and also provides a way of intercepting a call (like the AOP around advice semantics) is capable of being made into a secure object. Having said that, most Spring applications will simply use the three currently supported secure object types (AOP Alliance `MethodInvocation`, AspectJ `JoinPoint` and web request `FilterInvocation`) with complete transparency.

9.2 Core Services

Now that we have a high-level overview of the Spring Security architecture and its core classes, let's take a closer look at one or two of the core interfaces and their implementations, in particular the `AuthenticationManager`, `UserDetailsService` and the `AccessDecisionManager`. These crop up regularly throughout the remainder of this document so it's important you know how they are configured and how they operate.

The `AuthenticationManager`, `ProviderManager` and `AuthenticationProvider`

The `AuthenticationManager` is just an interface, so the implementation can be anything we choose, but how does it work in practice? What if we need to check multiple authentication databases or a combination of different authentication services such as a database and an LDAP server?

The default implementation in Spring Security is called `ProviderManager` and rather than handling the authentication request itself, it delegates to a list of configured `AuthenticationProvider`s, each of which is queried in turn to see if it can perform the authentication. Each provider will either throw an exception or return a fully populated `Authentication` object. Remember our good friends, `UserDetails` and `UserDetailsService`? If not, head back to the previous chapter and refresh your memory. The most common approach to verifying an authentication request is to load the corresponding `UserDetails` and check the loaded password against the one that has been entered by the user. This is the approach used by the `DaoAuthenticationProvider` (see below). The loaded `UserDetails` object - and particularly the `GrantedAuthority`s it contains - will be used when building the fully populated `Authentication` object which is returned from a successful authentication and stored in the `SecurityContext`.

If you are using the namespace, an instance of `ProviderManager` is created and maintained internally, and you add providers to it by using the namespace authentication provider elements (see [the namespace chapter](#)). In this case, you should not declare a `ProviderManager` bean in your application context. However, if you are not using the namespace then you would declare it like so:

```
<bean id="authenticationManager"
      class="org.springframework.security.authentication.ProviderManager">
  <constructor-arg>
    <list>
      <ref local="daoAuthenticationProvider"/>
      <ref local="anonymousAuthenticationProvider"/>
      <ref local="ldapAuthenticationProvider"/>
    </list>
  </constructor-arg>
</bean>
```

In the above example we have three providers. They are tried in the order shown (which is implied by the use of a `List`), with each provider able to attempt authentication, or skip authentication by simply returning `null`. If all implementations return `null`, the `ProviderManager` will throw a

`ProviderNotFoundException`. If you're interested in learning more about chaining providers, please refer to the `ProviderManager` Javadoc.

Authentication mechanisms such as a web form-login processing filter are injected with a reference to the `ProviderManager` and will call it to handle their authentication requests. The providers you require will sometimes be interchangeable with the authentication mechanisms, while at other times they will depend on a specific authentication mechanism. For example, `DaoAuthenticationProvider` and `LdapAuthenticationProvider` are compatible with any mechanism which submits a simple username/password authentication request and so will work with form-based logins or HTTP Basic authentication. On the other hand, some authentication mechanisms create an authentication request object which can only be interpreted by a single type of `AuthenticationProvider`. An example of this would be JA-SIG CAS, which uses the notion of a service ticket and so can therefore only be authenticated by a `CasAuthenticationProvider`. You needn't be too concerned about this, because if you forget to register a suitable provider, you'll simply receive a `ProviderNotFoundException` when an attempt to authenticate is made.

Erasing Credentials on Successful Authentication

By default (from Spring Security 3.1 onwards) the `ProviderManager` will attempt to clear any sensitive credentials information from the `Authentication` object which is returned by a successful authentication request. This prevents information like passwords being retained longer than necessary.

This may cause issues when you are using a cache of user objects, for example, to improve performance in a stateless application. If the `Authentication` contains a reference to an object in the cache (such as a `UserDetails` instance) and this has its credentials removed, then it will no longer be possible to authenticate against the cached value. You need to take this into account if you are using a cache. An obvious solution is to make a copy of the object first, either in the cache implementation or in the `AuthenticationProvider` which creates the returned `Authentication` object. Alternatively, you can disable the `eraseCredentialsAfterAuthentication` property on `ProviderManager`. See the Javadoc for more information.

DaoAuthenticationProvider

The simplest `AuthenticationProvider` implemented by Spring Security is `DaoAuthenticationProvider`, which is also one of the earliest supported by the framework. It leverages a `UserDetailsService` (as a DAO) in order to lookup the username, password and `GrantedAuthority`s. It authenticates the user simply by comparing the password submitted in a `UsernamePasswordAuthenticationToken` against the one loaded by the `UserDetailsService`. Configuring the provider is quite simple:

```
<bean id="daoAuthenticationProvider"
      class="org.springframework.security.authentication.dao.DaoAuthenticationProvider">
  <property name="userDetailsService" ref="inMemoryDaoImpl"/>
  <property name="passwordEncoder" ref="passwordEncoder"/>
</bean>
```

The `PasswordEncoder` is optional. A `PasswordEncoder` provides encoding and decoding of passwords presented in the `UserDetails` object that is returned from the configured `UserDetailsService`. This will be discussed in more detail [below](#).

UserDetailsService Implementations

As mentioned in the earlier in this reference guide, most authentication providers take advantage of the `UserDetails` and `UserDetailsService` interfaces. Recall that the contract for `UserDetailsService` is a single method:

```
UserDetails loadUserByUsername(String username) throws UsernameNotFoundException;
```

The returned `UserDetails` is an interface that provides getters that guarantee non-null provision of authentication information such as the username, password, granted authorities and whether the user account is enabled or disabled. Most authentication providers will use a `UserDetailsService`, even if the username and password are not actually used as part of the authentication decision. They may use the returned `UserDetails` object just for its `GrantedAuthority` information, because some other system (like LDAP or X.509 or CAS etc) has undertaken the responsibility of actually validating the credentials.

Given `UserDetailsService` is so simple to implement, it should be easy for users to retrieve authentication information using a persistence strategy of their choice. Having said that, Spring Security does include a couple of useful base implementations, which we'll look at below.

In-Memory Authentication

Is easy to use create a custom `UserDetailsService` implementation that extracts information from a persistence engine of choice, but many applications do not require such complexity. This is particularly true if you're building a prototype application or just starting integrating Spring Security, when you don't really want to spend time configuring databases or writing `UserDetailsService` implementations. For this sort of situation, a simple option is to use the `user-service` element from the security namespace:

```
<user-service id="userDetailsService">
  <!-- Password is prefixed with {noop} to indicate to DelegatingPasswordEncoder that
  NoOpPasswordEncoder should be used. This is not safe for production, but makes reading
  in samples easier. Normally passwords should be hashed using BCrypt -->
  <user name="jimi" password="{noop}jimispasword" authorities="ROLE_USER, ROLE_ADMIN" />
  <user name="bob" password="{noop}bobspasword" authorities="ROLE_USER" />
</user-service>
```

This also supports the use of an external properties file:

```
<user-service id="userDetailsService" properties="users.properties"/>
```

The properties file should contain entries in the form

```
username=password,grantedAuthority[,grantedAuthority][,enabled|disabled]
```

For example

```
jimi=jimispasword,ROLE_USER,ROLE_ADMIN,enabled
bob=bobspasword,ROLE_USER,enabled
```

JdbcDaoImpl

Spring Security also includes a `UserDetailsService` that can obtain authentication information from a JDBC data source. Internally Spring JDBC is used, so it avoids the complexity of a fully-featured object relational mapper (ORM) just to store user details. If your application does use an ORM tool, you might prefer to write a custom `UserDetailsService` to reuse the mapping files you've probably already created. Returning to `JdbcDaoImpl`, an example configuration is shown below:

```
<bean id="dataSource" class="org.springframework.jdbc.datasource.DriverManagerDataSource">
<property name="driverClassName" value="org.hsqldb.jdbcDriver"/>
<property name="url" value="jdbc:hsqldb:hsqldb://localhost:9001"/>
<property name="username" value="sa"/>
<property name="password" value="" />
</bean>

<bean id="userDetailsService"
      class="org.springframework.security.core.userdetails.jdbc.JdbcDaoImpl">
<property name="dataSource" ref="dataSource"/>
</bean>
```

You can use different relational database management systems by modifying the `DriverManagerDataSource` shown above. You can also use a global data source obtained from JNDI, as with any other Spring configuration.

Authority Groups

By default, `JdbcDaoImpl` loads the authorities for a single user with the assumption that the authorities are mapped directly to users (see the [database schema appendix](#)). An alternative approach is to partition the authorities into groups and assign groups to the user. Some people prefer this approach as a means of administering user rights. See the `JdbcDaoImpl` Javadoc for more information on how to enable the use of group authorities. The group schema is also included in the appendix.

10. Authentication

10.1 In-Memory Authentication

We have already seen an example of configuring in-memory authentication for a single user. Below is an example to configure multiple users:

```
@Bean
public UserDetailsService userDetailsService() throws Exception {
    // ensure the passwords are encoded properly
    UserBuilder users = User.withDefaultPasswordEncoder();
    InMemoryUserDetailsManager manager = new InMemoryUserDetailsManager();
    manager.createUser(users.username("user").password("password").roles("USER").build());
    manager.createUser(users.username("admin").password("password").roles("USER", "ADMIN").build());
    return manager;
}
```

10.2 JDBC Authentication

You can find the updates to support JDBC based authentication. The example below assumes that you have already defined a `DataSource` within your application. The [jdbc-javaconfig](#) sample provides a complete example of using JDBC based authentication.

```
@Autowired
private DataSource dataSource;

@Autowired
public void configureGlobal(AuthenticationManagerBuilder auth) throws Exception {
    // ensure the passwords are encoded properly
    UserBuilder users = User.withDefaultPasswordEncoder();
    auth
        .jdbcAuthentication()
            .dataSource(dataSource)
            .withDefaultSchema()
            .withUser(users.username("user").password("password").roles("USER"))
            .withUser(users.username("admin").password("password").roles("USER", "ADMIN"));
}
```

10.3 LDAP Authentication

Overview

LDAP is often used by organizations as a central repository for user information and as an authentication service. It can also be used to store the role information for application users.

There are many different scenarios for how an LDAP server may be configured so Spring Security's LDAP provider is fully configurable. It uses separate strategy interfaces for authentication and role retrieval and provides default implementations which can be configured to handle a wide range of situations.

You should be familiar with LDAP before trying to use it with Spring Security. The following link provides a good introduction to the concepts involved and a guide to setting up a directory using the free LDAP server OpenLDAP: <http://www.zytrax.com/books/ldap/>. Some familiarity with the JNDI APIs used to access LDAP from Java may also be useful. We don't use any third-party LDAP libraries (Mozilla, JLDAP etc.) in the LDAP provider, but extensive use is made of Spring LDAP, so some familiarity with that project may be useful if you plan on adding your own customizations.

When using LDAP authentication, it is important to ensure that you configure LDAP connection pooling properly. If you are unfamiliar with how to do this, you can refer to the [Java LDAP documentation](#).

Using LDAP with Spring Security

LDAP authentication in Spring Security can be roughly divided into the following stages.

- Obtaining the unique LDAP "Distinguished Name", or DN, from the login name. This will often mean performing a search in the directory, unless the exact mapping of usernames to DNs is known in advance. So a user might enter the name "joe" when logging in, but the actual name used to authenticate to LDAP will be the full DN, such as `uid=joe,ou=users,dc=spring,dc=io`.
- Authenticating the user, either by "binding" as that user or by performing a remote "compare" operation of the user's password against the password attribute in the directory entry for the DN.
- Loading the list of authorities for the user.

The exception is when the LDAP directory is just being used to retrieve user information and authenticate against it locally. This may not be possible as directories are often set up with limited read access for attributes such as user passwords.

We will look at some configuration scenarios below. For full information on available configuration options, please consult the security namespace schema (information from which should be available in your XML editor).

10.4 Configuring an LDAP Server

The first thing you need to do is configure the server against which authentication should take place. This is done using the `<ldap-server>` element from the security namespace. This can be configured to point at an external LDAP server, using the `url` attribute:

```
<ldap-server url="ldap://springframework.org:389/dc=springframework,dc=org" />
```

Note

`spring-security` provides integration with `apacheds` and `unboundid` as a embedded ldap servers. You can choose between them using the attribute `mode` in `ldap-server`.

Using an Embedded Test Server

The `<ldap-server>` element can also be used to create an embedded server, which can be very useful for testing and demonstrations. In this case you use it without the `url` attribute:

```
<ldap-server root="dc=springframework,dc=org"/>
```

Here we've specified that the root DIT of the directory should be `"dc=springframework,dc=org"`, which is the default. Used this way, the namespace parser will create an embedded Apache Directory server and scan the classpath for any LDIF files, which it will attempt to load into the server. You can customize this behaviour using the `ldif` attribute, which defines an LDIF resource to be loaded:

```
<ldap-server ldif="classpath:users.ldif" />
```

This makes it a lot easier to get up and running with LDAP, since it can be inconvenient to work all the time with an external server. It also insulates the user from the complex bean configuration needed to wire up an Apache Directory server. Using plain Spring Beans the configuration would be much more

cluttered. You must have the necessary Apache Directory dependency jars available for your application to use. These can be obtained from the LDAP sample application.

Using Bind Authentication

This is the most common LDAP authentication scenario.

```
<ldap-authentication-provider user-dn-pattern="uid={0},ou=people"/>
```

This simple example would obtain the DN for the user by substituting the user login name in the supplied pattern and attempting to bind as that user with the login password. This is OK if all your users are stored under a single node in the directory. If instead you wished to configure an LDAP search filter to locate the user, you could use the following:

```
<ldap-authentication-provider user-search-filter="(uid={0})"
  user-search-base="ou=people"/>
```

If used with the server definition above, this would perform a search under the DN `ou=people,dc=springframework,dc=org` using the value of the `user-search-filter` attribute as a filter. Again the user login name is substituted for the parameter in the filter name, so it will search for an entry with the `uid` attribute equal to the user name. If `user-search-base` isn't supplied, the search will be performed from the root.

Loading Authorities

How authorities are loaded from groups in the LDAP directory is controlled by the following attributes.

- `group-search-base`. Defines the part of the directory tree under which group searches should be performed.
- `group-role-attribute`. The attribute which contains the name of the authority defined by the group entry. Defaults to `cn`.
- `group-search-filter`. The filter which is used to search for group membership. The default is `uniqueMember={0}`, corresponding to the `groupOfUniqueNames` LDAP class³. In this case, the substituted parameter is the full distinguished name of the user. The parameter `{1}` can be used if you want to filter on the login name.

So if we used the following configuration

```
<ldap-authentication-provider user-dn-pattern="uid={0},ou=people"
  group-search-base="ou=groups" />
```

and authenticated successfully as user "ben", the subsequent loading of authorities would perform a search under the directory entry `ou=groups,dc=springframework,dc=org`, looking for entries which contain the attribute `uniqueMember` with value `uid=ben,ou=people,dc=springframework,dc=org`. By default the authority names will have the prefix `ROLE_` prepended. You can change this using the `role-prefix` attribute. If you don't want any prefix, use `role-prefix="none"`. For more information on loading authorities, see the Javadoc for the `DefaultLdapAuthoritiesPopulator` class.

10.5 Implementation Classes

The namespace configuration options we've used above are simple to use and much more concise than using Spring beans explicitly. There are situations when you may need to know how to configure Spring Security LDAP directly in your application context. You may wish to customize the behaviour of

some of the classes, for example. If you're happy using namespace configuration then you can skip this section and the next one.

The main LDAP provider class, `LdapAuthenticationProvider`, doesn't actually do much itself but delegates the work to two other beans, an `LdapAuthenticator` and an `LdapAuthoritiesPopulator` which are responsible for authenticating the user and retrieving the user's set of `GrantedAuthority`s respectively.

LdapAuthenticator Implementations

The authenticator is also responsible for retrieving any required user attributes. This is because the permissions on the attributes may depend on the type of authentication being used. For example, if binding as the user, it may be necessary to read them with the user's own permissions.

There are currently two authentication strategies supplied with Spring Security:

- Authentication directly to the LDAP server ("bind" authentication).
- Password comparison, where the password supplied by the user is compared with the one stored in the repository. This can either be done by retrieving the value of the password attribute and checking it locally or by performing an LDAP "compare" operation, where the supplied password is passed to the server for comparison and the real password value is never retrieved.

Common Functionality

Before it is possible to authenticate a user (by either strategy), the distinguished name (DN) has to be obtained from the login name supplied to the application. This can be done either by simple pattern-matching (by setting the `setUserDnPatterns` array property) or by setting the `userSearch` property. For the DN pattern-matching approach, a standard Java pattern format is used, and the login name will be substituted for the parameter `{0}`. The pattern should be relative to the DN that the configured `SpringSecurityContextSource` will bind to (see the section on [connecting to the LDAP server](#) for more information on this). For example, if you are using an LDAP server with the URL `ldap://monkeymachine.co.uk/dc=springframework,dc=org`, and have a pattern `uid={0},ou=greatapes`, then a login name of "gorilla" will map to a DN `uid=gorilla,ou=greatapes,dc=springframework,dc=org`. Each configured DN pattern will be tried in turn until a match is found. For information on using a search, see the section on [search objects](#) below. A combination of the two approaches can also be used - the patterns will be checked first and if no matching DN is found, the search will be used.

BindAuthenticator

The class `BindAuthenticator` in the package `org.springframework.security.ldap.authentication` implements the bind authentication strategy. It simply attempts to bind as the user.

PasswordComparisonAuthenticator

The class `PasswordComparisonAuthenticator` implements the password comparison authentication strategy.

Connecting to the LDAP Server

The beans discussed above have to be able to connect to the server. They both have to be supplied with a `SpringSecurityContextSource` which is an extension of Spring LDAP's `ContextSource`. Unless you have special requirements, you will usually configure a

`DefaultSpringSecurityContextSource` bean, which can be configured with the URL of your LDAP server and optionally with the username and password of a "manager" user which will be used by default when binding to the server (instead of binding anonymously). For more information read the Javadoc for this class and for Spring LDAP's `AbstractContextSource`.

LDAP Search Objects

Often a more complicated strategy than simple DN-matching is required to locate a user entry in the directory. This can be encapsulated in an `LdapUserSearch` instance which can be supplied to the authenticator implementations, for example, to allow them to locate a user. The supplied implementation is `FilterBasedLdapUserSearch`.

FilterBasedLdapUserSearch

This bean uses an LDAP filter to match the user object in the directory. The process is explained in the Javadoc for the corresponding search method on the [JDK DirContext class](#). As explained there, the search filter can be supplied with parameters. For this class, the only valid parameter is `{0}` which will be replaced with the user's login name.

LdapAuthoritiesPopulator

After authenticating the user successfully, the `LdapAuthenticationProvider` will attempt to load a set of authorities for the user by calling the configured `LdapAuthoritiesPopulator` bean. The `DefaultLdapAuthoritiesPopulator` is an implementation which will load the authorities by searching the directory for groups of which the user is a member (typically these will be `groupOfNames` or `groupOfUniqueNames` entries in the directory). Consult the Javadoc for this class for more details on how it works.

If you want to use LDAP only for authentication, but load the authorities from a difference source (such as a database) then you can provide your own implementation of this interface and inject that instead.

Spring Bean Configuration

A typical configuration, using some of the beans we've discussed here, might look like this:

```
<bean id="contextSource"
    class="org.springframework.security.ldap.DefaultSpringSecurityContextSource">
    <constructor-arg value="ldap://monkeymachine:389/dc=springframework,dc=org"/>
    <property name="userDn" value="cn=manager,dc=springframework,dc=org"/>
    <property name="password" value="password"/>
</bean>

<bean id="ldapAuthProvider"
    class="org.springframework.security.ldap.authentication.LdapAuthenticationProvider">
    <constructor-arg>
        <bean class="org.springframework.security.ldap.authentication.BindAuthenticator">
            <constructor-arg ref="contextSource"/>
            <property name="userDnPatterns">
                <list><value>uid={0},ou=people</value></list>
            </property>
        </bean>
    </constructor-arg>
    <constructor-arg>
        <bean class="org.springframework.security.ldap.userdetails.DefaultLdapAuthoritiesPopulator">
            <constructor-arg ref="contextSource"/>
            <constructor-arg value="ou=groups"/>
            <property name="groupRoleAttribute" value="ou"/>
        </bean>
    </constructor-arg>
</bean>
```

This would set up the provider to access an LDAP server with URL `ldap://monkeymachine:389/dc=springframework,dc=org`. Authentication will be performed by attempting to bind with the DN `uid=<user-login-name>,ou=people,dc=springframework,dc=org`. After successful authentication, roles will be assigned to the user by searching under the DN `ou=groups,dc=springframework,dc=org` with the default filter (`member=<user's-DN>`). The role name will be taken from the "ou" attribute of each match.

To configure a user search object, which uses the filter (`uid=<user-login-name>`) for use instead of the DN-pattern (or in addition to it), you would configure the following bean

```
<bean id="userSearch"
      class="org.springframework.security.ldap.search.FilterBasedLdapUserSearch">
  <constructor-arg index="0" value="" />
  <constructor-arg index="1" value="(uid={0})" />
  <constructor-arg index="2" ref="contextSource" />
</bean>
```

and use it by setting the `BindAuthenticator` bean's `userSearch` property. The authenticator would then call the search object to obtain the correct user's DN before attempting to bind as this user.

LDAP Attributes and Customized UserDetails

The net result of an authentication using `LdapAuthenticationProvider` is the same as a normal Spring Security authentication using the standard `UserDetailsService` interface. A `UserDetails` object is created and stored in the returned `Authentication` object. As with using a `UserDetailsService`, a common requirement is to be able to customize this implementation and add extra properties. When using LDAP, these will normally be attributes from the user entry. The creation of the `UserDetails` object is controlled by the provider's `UserDetailsContextMapper` strategy, which is responsible for mapping user objects to and from LDAP context data:

```
public interface UserDetailsContextMapper {

    UserDetails mapUserFromContext(DirContextOperations ctx, String username,
                                   Collection<GrantedAuthority> authorities);

    void mapUserToContext(UserDetails user, DirContextAdapter ctx);
}
```

Only the first method is relevant for authentication. If you provide an implementation of this interface and inject it into the `LdapAuthenticationProvider`, you have control over exactly how the `UserDetails` object is created. The first parameter is an instance of Spring LDAP's `DirContextOperations` which gives you access to the LDAP attributes which were loaded during authentication. The `username` parameter is the name used to authenticate and the final parameter is the collection of authorities loaded for the user by the configured `LdapAuthoritiesPopulator`.

The way the context data is loaded varies slightly depending on the type of authentication you are using. With the `BindAuthenticator`, the context returned from the bind operation will be used to read the attributes, otherwise the data will be read using the standard context obtained from the configured `ContextSource` (when a search is configured to locate the user, this will be the data returned by the search object).

10.6 Active Directory Authentication

Active Directory supports its own non-standard authentication options, and the normal usage pattern doesn't fit too cleanly with the standard `LdapAuthenticationProvider`. Typically authentication is performed using the domain username (in the form `user@domain`), rather than using an LDAP

distinguished name. To make this easier, Spring Security 3.1 has an authentication provider which is customized for a typical Active Directory setup.

ActiveDirectoryLdapAuthenticationProvider

Configuring `ActiveDirectoryLdapAuthenticationProvider` is quite straightforward. You just need to supply the domain name and an LDAP URL supplying the address of the server⁵. An example configuration would then look like this:

```
<bean id="adAuthenticationProvider"
    class="org.springframework.security.ldap.authentication.ad.ActiveDirectoryLdapAuthenticationProvider">
    <constructor-arg value="mydomain.com" />
    <constructor-arg value="ldap://adserver.mydomain.com/" />
</bean>
```

Note that there is no need to specify a separate `ContextSource` in order to define the server location - the bean is completely self-contained. A user named "Sharon", for example, would then be able to authenticate by entering either the username `sharon` or the full Active Directory userPrincipalName, namely `sharon@mydomain.com`. The user's directory entry will then be located, and the attributes returned for possible use in customizing the created `UserDetails` object (a `UserDetailsContextMapper` can be injected for this purpose, as described above). All interaction with the directory takes place with the identity of the user themselves. There is no concept of a "manager" user.

By default, the user authorities are obtained from the `memberOf` attribute values of the user entry. The authorities allocated to the user can again be customized using a `UserDetailsContextMapper`. You can also inject a `GrantedAuthoritiesMapper` into the provider instance to control the authorities which end up in the `Authentication` object.

Active Directory Error Codes

By default, a failed result will cause a standard Spring Security `BadCredentialsException`. If you set the property `convertSubErrorCodesToExceptions` to `true`, the exception messages will be parsed to attempt to extract the Active Directory-specific error code and raise a more specific exception. Check the class Javadoc for more information.

10.7 LDAP Java Configuration

You can find the updates to support LDAP based authentication. The [ldap-javaconfig](#) sample provides a complete example of using LDAP based authentication.

```
@Autowired
private DataSource dataSource;

@Autowired
public void configureGlobal(AuthenticationManagerBuilder auth) throws Exception {
    auth
        .ldapAuthentication()
            .userDnPatterns("uid={0},ou=people")
            .groupSearchBase("ou=groups");
}
```

The example above uses the following LDIF and an embedded Apache DS LDAP instance.

⁵It is also possible to obtain the server's IP address using a DNS lookup. This is not currently supported, but hopefully will be in a future version.

users.ldif.

```

dn: ou=groups,dc=springframework,dc=org
objectclass: top
objectclass: organizationalUnit
ou: groups

dn: ou=people,dc=springframework,dc=org
objectclass: top
objectclass: organizationalUnit
ou: people

dn: uid=admin,ou=people,dc=springframework,dc=org
objectclass: top
objectclass: person
objectclass: organizationalPerson
objectclass: inetOrgPerson
cn: Rod Johnson
sn: Johnson
uid: admin
userPassword: password

dn: uid=user,ou=people,dc=springframework,dc=org
objectclass: top
objectclass: person
objectclass: organizationalPerson
objectclass: inetOrgPerson
cn: Dianne Emu
sn: Emu
uid: user
userPassword: password

dn: cn=user,ou=groups,dc=springframework,dc=org
objectclass: top
objectclass: groupOfNames
cn: user
uniqueMember: uid=admin,ou=people,dc=springframework,dc=org
uniqueMember: uid=user,ou=people,dc=springframework,dc=org

dn: cn=admin,ou=groups,dc=springframework,dc=org
objectclass: top
objectclass: groupOfNames
cn: admin
uniqueMember: uid=admin,ou=people,dc=springframework,dc=org

```

10.8 AuthenticationProvider

AuthenticationProvider Java Configuration

You can define custom authentication by exposing a custom `AuthenticationProvider` as a bean. For example, the following will customize authentication assuming that `SpringAuthenticationProvider` implements `AuthenticationProvider`:

Note

This is only used if the `AuthenticationManagerBuilder` has not been populated

```

@Bean
public SpringAuthenticationProvider springAuthenticationProvider() {
    return new SpringAuthenticationProvider();
}

```

AuthenticationProvider XML Configuration

In practice you will need a more scalable source of user information than a few names added to the application context file. Most likely you will want to store your user information in something like a database or an LDAP server. LDAP namespace configuration is dealt with in the [LDAP chapter](#), so we won't cover it here. If you have a custom implementation of Spring Security's `UserDetailsService`, called "myUserDetailsService" in your application context, then you can authenticate against this using

```
<authentication-manager>
  <authentication-provider user-service-ref='myUserDetailsService' />
</authentication-manager>
```

If you want to use a database, then you can use

```
<authentication-manager>
<authentication-provider>
  <jdbc-user-service data-source-ref="securityDataSource" />
</authentication-provider>
</authentication-manager>
```

Where "securityDataSource" is the name of a `DataSource` bean in the application context, pointing at a database containing the standard Spring Security [user data tables](#). Alternatively, you could configure a Spring Security `JdbcDaoImpl` bean and point at that using the `user-service-ref` attribute:

```
<authentication-manager>
<authentication-provider user-service-ref='myUserDetailsService' />
</authentication-manager>

<beans:bean id="myUserDetailsService"
  class="org.springframework.security.core.userdetails.jdbc.JdbcDaoImpl">
<beans:property name="dataSource" ref="dataSource" />
</beans:bean>
```

You can also use standard `AuthenticationProvider` beans as follows

```
<authentication-manager>
  <authentication-provider ref='myAuthenticationProvider' />
</authentication-manager>
```

where `myAuthenticationProvider` is the name of a bean in your application context which implements `AuthenticationProvider`. You can use multiple `authentication-provider` elements, in which case the providers will be queried in the order they are declared. See Section 10.11, "The Authentication Manager and the Namespace" for more information on how the Spring Security `AuthenticationManager` is configured using the namespace.

10.9 UserDetailsService

You can define custom authentication by exposing a custom `UserDetailsService` as a bean. For example, the following will customize authentication assuming that `SpringDataUserDetailsService` implements `UserDetailsService`:

Note

This is only used if the `AuthenticationManagerBuilder` has not been populated and no `AuthenticationProviderBean` is defined.

```
@Bean
public SpringDataUserDetailsService springDataUserDetailsService() {
    return new SpringDataUserDetailsService();
}
```

You can also customize how passwords are encoded by exposing a `PasswordEncoder` as a bean. For example, if you use `bcrypt` you can add a bean definition as shown below:

```
@Bean
public BCryptPasswordEncoder passwordEncoder() {
    return new BCryptPasswordEncoder();
}
```

10.10 Password Encoding

Spring Security's `PasswordEncoder` interface is used to perform a one way transformation of a password to allow the password to be stored securely. Given `PasswordEncoder` is a one way transformation, it is not intended when the password transformation needs to be two way (i.e. storing credentials used to authenticate to a database). Typically `PasswordEncoder` is used for storing a password that needs to be compared to a user provided password at the time of authentication.

Password History

Throughout the years the standard mechanism for storing passwords has evolved. In the beginning passwords were stored in plain text. The passwords were assumed to be safe because the data store the passwords were saved in required credentials to access it. However, malicious users were able to find ways to get large "data dumps" of usernames and passwords using attacks like SQL Injection. As more and more user credentials became public security experts realized we needed to do more to protect users passwords.

Developers were then encouraged to store passwords after running them through a one way hash such as SHA-256. When a user tried to authenticate, the hashed password would be compared to the hash of the password that they typed. This meant that the system only needed to store the one way hash of the password. If a breach occurred, then only the one way hashes of the passwords were exposed. Since the hashes were one way and it was computationally difficult to guess the passwords given the hash, it would not be worth the effort to figure out each password in the system. To defeat this new system malicious users decided to create lookup tables known as [Rainbow Tables](#). Rather than doing the work of guessing each password every time, they computed the password once and stored it in a lookup table.

To mitigate the effectiveness of Rainbow Tables, developers were encouraged to use salted passwords. Instead of using just the password as input to the hash function, random bytes (known as salt) would be generated for every users' password. The salt and the user's password would be ran through the hash function which produced a unique hash. The salt would be stored alongside the user's password in clear text. Then when a user tried to authenticate, the hashed password would be compared to the hash of the stored salt and the password that they typed. The unique salt meant that Rainbow Tables were no longer effective because the hash was different for every salt and password combination.

In modern times we realize that cryptographic hashes (like SHA-256) are no longer secure. The reason is that with modern hardware we can perform billions of hash calculations a second. This means that we can crack each password individually with ease.

Developers are now encouraged to leverage adaptive one-way functions to store a password. Validation of passwords with adaptive one-way functions are intentionally resource (i.e. CPU, memory, etc) intensive. An adaptive one-way function allows configuring a "work factor" which can grow as hardware

gets better. It is recommended that the "work factor" be tuned to take about 1 second to verify a password on your system. This trade off is to make it difficult for attackers to crack the password, but not so costly it puts excessive burden on your own system. Spring Security has attempted to provide a good starting point for the "work factor", but users are encouraged to customize the "work factor" for their own system since the performance will vary drastically from system to system. Examples of adaptive one-way functions that should be used include [bcrypt](#), [PBKDF2](#), [scrypt](#), and [Argon2](#).

Because adaptive one-way functions are intentionally resource intensive, validating a username and password for every request will degrade performance of an application significantly. There is nothing Spring Security (or any other library) can do to speed up the validation of the password since security is gained by making the validation resource intensive. Users are encouraged to exchange the long term credentials (i.e. username and password) for a short term credential (i.e. session, OAuth Token, etc). The short term credential can be validated quickly without any loss in security.

DelegatingPasswordEncoder

Prior to Spring Security 5.0 the default `PasswordEncoder` was `NoOpPasswordEncoder` which required plain text passwords. Based upon the [Password History](#) section you might expect that the default `PasswordEncoder` is now something like `BCryptPasswordEncoder`. However, this ignores three real world problems:

- There are many applications using old password encodings that cannot easily migrate
- The best practice for password storage will change again.
- As a framework Spring Security cannot make breaking changes frequently

Instead Spring Security introduces `DelegatingPasswordEncoder` which solves all of the problems by:

- Ensuring that passwords are encoded using the current password storage recommendations
- Allowing for validating passwords in modern and legacy formats
- Allowing for upgrading the encoding in the future

You can easily construct an instance of `DelegatingPasswordEncoder` using `PasswordEncoderFactories`.

```
PasswordEncoder passwordEncoder =
    PasswordEncoderFactories.createDelegatingPasswordEncoder();
```

Alternatively, you may create your own custom instance. For example:

```
String idForEncode = "bcrypt";
Map encoders = new HashMap<>();
encoders.put(idForEncode, new BCryptPasswordEncoder());
encoders.put("noop", NoOpPasswordEncoder.getInstance());
encoders.put("pbkdf2", new Pbkdf2PasswordEncoder());
encoders.put("scrypt", new SCryptPasswordEncoder());
encoders.put("sha256", new StandardPasswordEncoder());

PasswordEncoder passwordEncoder =
    new DelegatingPasswordEncoder(idForEncode, encoders);
```

Password Storage Format

The general format for a password is:


```
{id}encodedPassword
```

Such that `id` is an identifier used to look up which `PasswordEncoder` should be used and `encodedPassword` is the original encoded password for the selected `PasswordEncoder`. The `id` must be at the beginning of the password, start with `{` and end with `}`. If the `id` cannot be found, the `id` will be null. For example, the following might be a list of passwords encoded using different `id`. All of the original passwords are "password".

```
{bcrypt}$2a$10$dXJ3SW6G7P50lGmMkkmwe.20cQQubK3.HZWzG3YB1t1Ry.fqvM/BG ❶
{noop}password ❷
{pbkdf2}5d923b44a6d129f3ddf3e3c8d29412723dcbde72445e8ef6bf3b508fbf17fa4ed4d6b99ca763d8dc ❸
{sCrypt}$e0801$8bWJaSu2IKSn9Z9kM+TPXfOc/9bdYSrN1oD9qfVThWEwdRTnO7re7Ei+fUZrJ68k9lTyuTeUp4of4g24hHnazw==
$0AOec05+bXxvuu/1qZ6NUR+xQYvYv7BeL1QxwRpY5Pc= ❹
{sha256}97cde38028ad898ebc02e690819fa220e88c62e0699403e94fff291cfffaf8410849f27605abcbc0 ❺
```

- ❶ The first password would have a `PasswordEncoder` `id` of `bcrypt` and `encodedPassword` of `$2a$10$dXJ3SW6G7P50lGmMkkmwe.20cQQubK3.HZWzG3YB1t1Ry.fqvM/BG`. When matching it would delegate to `BCryptPasswordEncoder`
- ❷ The second password would have a `PasswordEncoder` `id` of `noop` and `encodedPassword` of `password`. When matching it would delegate to `NoOpPasswordEncoder`
- ❸ The third password would have a `PasswordEncoder` `id` of `pbkdf2` and `encodedPassword` of `5d923b44a6d129f3ddf3e3c8d29412723dcbde72445e8ef6bf3b508fbf17fa4ed4d6b99ca763d8dc`. When matching it would delegate to `Pbkdf2PasswordEncoder`
- ❹ The fourth password would have a `PasswordEncoder` `id` of `sCrypt` and `encodedPassword` of `$e0801$8bWJaSu2IKSn9Z9kM+TPXfOc/9bdYSrN1oD9qfVThWEwdRTnO7re7Ei+fUZrJ68k9lTyuTeUp4of4g24hHnazw==$0AOec05+bXxvuu/1qZ6NUR+xQYvYv7BeL1QxwRpY5Pc=`. When matching it would delegate to `SCryptPasswordEncoder`
- ❺ The final password would have a `PasswordEncoder` `id` of `sha256` and `encodedPassword` of `97cde38028ad898ebc02e690819fa220e88c62e0699403e94fff291cfffaf8410849f27605abcbc0`. When matching it would delegate to `StandardPasswordEncoder`

Note

Some users might be concerned that the storage format is provided for a potential hacker. This is not a concern because the storage of the password does not rely on the algorithm being a secret. Additionally, most formats are easy for an attacker to figure out without the prefix. For example, `BCrypt` passwords often start with `$2a$`.

Password Encoding

The `idForEncode` passed into the constructor determines which `PasswordEncoder` will be used for encoding passwords. In the `DelegatingPasswordEncoder` we constructed above, that means that the result of encoding `password` would be delegated to `BCryptPasswordEncoder` and be prefixed with `{bcrypt}`. The end result would look like:

```
{bcrypt}$2a$10$dXJ3SW6G7P50lGmMkkmwe.20cQQubK3.HZWzG3YB1t1Ry.fqvM/BG
```

Password Matching

Matching is done based upon the `{id}` and the mapping of the `id` to the `PasswordEncoder` provided in the constructor. Our example in the section called "Password Storage Format" provides a working example of how this is done. By default, the result of invoking

`matches(CharSequence, String)` with a password and an `id` that is not mapped (including a null `id`) will result in an `IllegalArgumentException`. This behavior can be customized using `DelegatingPasswordEncoder.setDefaultPasswordEncoderForMatches(PasswordEncoder)`.

By using the `id` we can match on any password encoding, but encode passwords using the most modern password encoding. This is important, because unlike encryption, password hashes are designed so that there is no simple way to recover the plaintext. Since there is no way to recover the plaintext, it makes it difficult to migrate the passwords. While it is simple for users to migrate `NoOpPasswordEncoder`, we chose to include it by default to make it simple for the getting started experience.

Getting Started Experience

If you are putting together a demo or a sample, it is a bit cumbersome to take time to hash the passwords of your users. There are convenience mechanisms to make this easier, but this is still not intended for production.

```
User user = User.withDefaultPasswordEncoder()
    .username("user")
    .password("password")
    .roles("user")
    .build();
System.out.println(user.getPassword());
// {bcrypt}$2a$10$dXJ3SW6G7P50lGmMkkmwe.20cQQubK3.HZWzG3YB1t1Ry.fqvM/BG
```

If you are creating multiple users, you can also reuse the builder.

```
UserBuilder users = User.withDefaultPasswordEncoder();
User user = users
    .username("user")
    .password("password")
    .roles("USER")
    .build();
User admin = users
    .username("admin")
    .password("password")
    .roles("USER", "ADMIN")
    .build();
```

This does hash the password that is stored, but the passwords are still exposed in memory and in the compiled source code. Therefore, it is still not considered secure for a production environment. For production, you should hash your passwords externally.

Troubleshooting

The following error occurs when one of the passwords that are stored has no `id` as described in the section called “Password Storage Format”.

```
java.lang.IllegalArgumentException: There is no PasswordEncoder mapped for the id "null"
    at org.springframework.security.crypto.password.DelegatingPasswordEncoder
    $UnmappedIdPasswordEncoder.matches(DelegatingPasswordEncoder.java:233)
    at
    org.springframework.security.crypto.password.DelegatingPasswordEncoder.matches(DelegatingPasswordEncoder.java:196)
```

The easiest way to resolve the error is to switch to explicitly provide the `PasswordEncoder` that your passwords are encoded with. The easiest way to resolve it is to figure out how your passwords are currently being stored and explicitly provide the correct `PasswordEncoder`. If you are migrating from Spring Security 4.2.x you can revert to the previous behavior by exposing a `NoOpPasswordEncoder` bean. For example, if you are using Java Configuration, you can create a configuration that looks like:

Warning

Reverting to `NoOpPasswordEncoder` is not considered to be secure. You should instead migrate to using `DelegatingPasswordEncoder` to support secure password encoding.

```
@Bean
public static NoOpPasswordEncoder passwordEncoder() {
    return NoOpPasswordEncoder.getInstance();
}
```

if you are using XML configuration, you can expose a `PasswordEncoder` with the id `passwordEncoder`:

```
<b:bean id="passwordEncoder"
        class="org.springframework.security.crypto.password.NoOpPasswordEncoder" factory-
        method="getInstance"/>
```

Alternatively, you can prefix all of your passwords with the correct id and continue to use `DelegatingPasswordEncoder`. For example, if you are using `BCrypt`, you would migrate your password from something like:

```
$2a$10$dXJ3SW6G7P501GmMkkmwe.20cQQubK3.HZWzG3YB1t1Ry.fqvM/BG
```

to

```
{bcrypt}$2a$10$dXJ3SW6G7P501GmMkkmwe.20cQQubK3.HZWzG3YB1t1Ry.fqvM/BG
```

For a complete listing of the mappings refer to the Javadoc on [PasswordEncoderFactories](#).

BCryptPasswordEncoder

The `BCryptPasswordEncoder` implementation uses the widely supported [bcrypt](#) algorithm to hash the passwords. In order to make it more resistant to password cracking, `bcrypt` is deliberately slow. Like other adaptive one-way functions, it should be tuned to take about 1 second to verify a password on your system.

```
// Create an encoder with strength 16
BCryptPasswordEncoder encoder = new BCryptPasswordEncoder(16);
String result = encoder.encode("myPassword");
assertTrue(encoder.matches("myPassword", result));
```

Argon2PasswordEncoder

The `Argon2PasswordEncoder` implementation uses the [Argon2](#) algorithm to hash the passwords. `Argon2` is the winner of the [Password Hashing Competition](#). In order to defeat password cracking on custom hardware, `Argon2` is a deliberately slow algorithm that requires large amounts of memory. Like other adaptive one-way functions, it should be tuned to take about 1 second to verify a password on your system. The current implementation of the `Argon2PasswordEncoder` requires `BouncyCastle`.

```
// Create an encoder with all the defaults
Argon2PasswordEncoder encoder = new Argon2PasswordEncoder();
String result = encoder.encode("myPassword");
assertTrue(encoder.matches("myPassword", result));
```

Pbkdf2PasswordEncoder

The `Pbkdf2PasswordEncoder` implementation uses the [PBKDF2](#) algorithm to hash the passwords. In order to defeat password cracking `PBKDF2` is a deliberately slow algorithm. Like other adaptive one-

way functions, it should be tuned to take about 1 second to verify a password on your system. This algorithm is a good choice when FIPS certification is required.

```
// Create an encoder with all the defaults
Pbkdf2PasswordEncoder encoder = new Pbkdf2PasswordEncoder();
String result = encoder.encode("myPassword");
assertTrue(encoder.matches("myPassword", result));
```

SCryptPasswordEncoder

The `SCryptPasswordEncoder` implementation uses [scrypt](#) algorithm to hash the passwords. In order to defeat password cracking on custom hardware scrypt is a deliberately slow algorithm that requires large amounts of memory. Like other adaptive one-way functions, it should be tuned to take about 1 second to verify a password on your system.

```
// Create an encoder with all the defaults
SCryptPasswordEncoder encoder = new SCryptPasswordEncoder();
String result = encoder.encode("myPassword");
assertTrue(encoder.matches("myPassword", result));
```

Other PasswordEncoders

There are a significant number of other `PasswordEncoder` implementations that exist entirely for backward compatibility. They are all deprecated to indicate that they are no longer considered secure. However, there are no plans to remove them since it is difficult to migrate existing legacy systems.

Password Encoder XML Configuration

Passwords should always be encoded using a secure hashing algorithm designed for the purpose (not a standard algorithm like SHA or MD5). This is supported by the `<password-encoder>` element. With bcrypt encoded passwords, the original authentication provider configuration would look like this:

```
<beans:bean name="bcryptEncoder"
  class="org.springframework.security.crypto.bcrypt.BCryptPasswordEncoder"/>

<authentication-manager>
<authentication-provider>
  <password-encoder ref="bcryptEncoder"/>
  <user-service>
  <user name="jimi" password="$2a$10$ddEWZU18aU0GdZPPpy7wbu82dvEw/pBpbRvDQRqA41y6mK1CoH00m"
    authorities="ROLE_USER, ROLE_ADMIN" />
  <user name="bob" password="$2a$10$/e1FpMBnAYYig6KRR5bvOOYeZr1ie1hSogJryg9qDlhza4oCw1Qka"
    authorities="ROLE_USER" />
  </user-service>
</authentication-provider>
</authentication-manager>
```

bcrypt is a good choice for most cases, unless you have a legacy system which forces you to use a different algorithm. If you are using a simple hashing algorithm or, even worse, storing plain text passwords, then you should consider migrating to a more secure option like bcrypt.

10.11 The Authentication Manager and the Namespace

The main interface which provides authentication services in Spring Security is the `AuthenticationManager`. This is usually an instance of Spring Security's `ProviderManager` class, which you may already be familiar with if you've used the framework before. If not, it will be covered later, in the [technical overview chapter](#). The bean instance is registered using the `authentication-manager` namespace element. You can't use a custom `AuthenticationManager` if you are using

either HTTP or method security through the namespace, but this should not be a problem as you have full control over the `AuthenticationProvider`s that are used.

You may want to register additional `AuthenticationProvider` beans with the `ProviderManager` and you can do this using the `<authentication-provider>` element with the `ref` attribute, where the value of the attribute is the name of the provider bean you want to add. For example:

```
<authentication-manager>
<authentication-provider ref="casAuthenticationProvider"/>
</authentication-manager>

<bean id="casAuthenticationProvider"
      class="org.springframework.security.cas.authentication.CasAuthenticationProvider">
  ...
</bean>
```

Another common requirement is that another bean in the context may require a reference to the `AuthenticationManager`. You can easily register an alias for the `AuthenticationManager` and use this name elsewhere in your application context.

```
<security:authentication-manager alias="authenticationManager">
  ...
</security:authentication-manager>

<bean id="customizedFormLoginFilter"
      class="com.somecompany.security.web.CustomFormLoginFilter">
  <property name="authenticationManager" ref="authenticationManager"/>
  ...
</bean>
```

10.12 Session Management

HTTP session related functionality is handled by a combination of the `SessionManagementFilter` and the `SessionAuthenticationStrategy` interface, which the filter delegates to. Typical usage includes session-fixation protection attack prevention, detection of session timeouts and restrictions on how many sessions an authenticated user may have open concurrently.

Detecting Timeouts

You can configure Spring Security to detect the submission of an invalid session ID and redirect the user to an appropriate URL. This is achieved through the `session-management` element:

```
<http>
  ...
  <session-management invalid-session-url="/invalidSession.htm" />
</http>
```

Note that if you use this mechanism to detect session timeouts, it may falsely report an error if the user logs out and then logs back in without closing the browser. This is because the session cookie is not cleared when you invalidate the session and will be resubmitted even if the user has logged out. You may be able to explicitly delete the `JSESSIONID` cookie on logging out, for example by using the following syntax in the logout handler:

```
<http>
  <logout delete-cookies="JSESSIONID" />
</http>
```

Unfortunately this can't be guaranteed to work with every servlet container, so you will need to test it in your environment

Note

=== If you are running your application behind a proxy, you may also be able to remove the session cookie by configuring the proxy server. For example, using Apache HTTPD's `mod_headers`, the following directive would delete the `JSESSIONID` cookie by expiring it in the response to a logout request (assuming the application is deployed under the path `/tutorial`):

```
<LocationMatch "/tutorial/logout">
Header always set Set-Cookie "JSESSIONID=;Path=/tutorial;Expires=Thu, 01 Jan 1970 00:00:00 GMT"
</LocationMatch>
```

===

Concurrent Session Control

If you wish to place constraints on a single user's ability to log in to your application, Spring Security supports this out of the box with the following simple additions. First you need to add the following listener to your `web.xml` file to keep Spring Security updated about session lifecycle events:

```
<listener>
<listener-class>
    org.springframework.security.web.session.HttpSessionEventPublisher
</listener-class>
</listener>
```

Then add the following lines to your application context:

```
<http>
...
<session-management>
    <concurrency-control max-sessions="1" />
</session-management>
</http>
```

This will prevent a user from logging in multiple times - a second login will cause the first to be invalidated. Often you would prefer to prevent a second login, in which case you can use

```
<http>
...
<session-management>
    <concurrency-control max-sessions="1" error-if-maximum-exceeded="true" />
</session-management>
</http>
```

The second login will then be rejected. By "rejected", we mean that the user will be sent to the `authentication-failure-url` if form-based login is being used. If the second authentication takes place through another non-interactive mechanism, such as "remember-me", an "unauthorized" (401) error will be sent to the client. If instead you want to use an error page, you can add the attribute `session-authentication-error-url` to the `session-management` element.

If you are using a customized authentication filter for form-based login, then you have to configure concurrent session control support explicitly. More details can be found in the [Session Management chapter](#).

Session Fixation Attack Protection

[Session fixation](#) attacks are a potential risk where it is possible for a malicious attacker to create a session by accessing a site, then persuade another user to log in with the same session (by sending

them a link containing the session identifier as a parameter, for example). Spring Security protects against this automatically by creating a new session or otherwise changing the session ID when a user logs in. If you don't require this protection, or it conflicts with some other requirement, you can control the behavior using the `session-fixation-protection` attribute on `<session-management>`, which has four options

- `none` - Don't do anything. The original session will be retained.
- `newSession` - Create a new "clean" session, without copying the existing session data (Spring Security-related attributes will still be copied).
- `migrateSession` - Create a new session and copy all existing session attributes to the new session. This is the default in Servlet 3.0 or older containers.
- `changeSessionId` - Do not create a new session. Instead, use the session fixation protection provided by the Servlet container (`HttpServletRequest#changeSessionId()`). This option is only available in Servlet 3.1 (Java EE 7) and newer containers. Specifying it in older containers will result in an exception. This is the default in Servlet 3.1 and newer containers.

When session fixation protection occurs, it results in a `SessionFixationProtectionEvent` being published in the application context. If you use `changeSessionId`, this protection will also result in any `javax.servlet.http.HttpSessionIdListeners` being notified, so use caution if your code listens for both events. See the [Session Management](#) chapter for additional information.

SessionManagementFilter

The `SessionManagementFilter` checks the contents of the `SecurityContextRepository` against the current contents of the `SecurityContextHolder` to determine whether a user has been authenticated during the current request, typically by a non-interactive authentication mechanism, such as pre-authentication or remember-me¹⁹. If the repository contains a security context, the filter does nothing. If it doesn't, and the thread-local `SecurityContext` contains a (non-anonymous) `Authentication` object, the filter assumes they have been authenticated by a previous filter in the stack. It will then invoke the configured `SessionAuthenticationStrategy`.

If the user is not currently authenticated, the filter will check whether an invalid session ID has been requested (because of a timeout, for example) and will invoke the configured `InvalidSessionStrategy`, if one is set. The most common behaviour is just to redirect to a fixed URL and this is encapsulated in the standard implementation `SimpleRedirectInvalidSessionStrategy`. The latter is also used when configuring an invalid session URL through the namespace, [as described earlier](#).

SessionAuthenticationStrategy

`SessionAuthenticationStrategy` is used by both `SessionManagementFilter` and `AbstractAuthenticationProcessingFilter`, so if you are using a customized form-login class, for example, you will need to inject it into both of these. In this case, a typical configuration, combining the namespace and custom beans might look like this:

¹⁹Authentication by mechanisms which perform a redirect after authenticating (such as form-login) will not be detected by `SessionManagementFilter`, as the filter will not be invoked during the authenticating request. Session-management functionality has to be handled separately in these cases.

```

<http>
<custom-filter position="FORM_LOGIN_FILTER" ref="myAuthFilter" />
<session-management session-authentication-strategy-ref="sas"/>
</http>

<beans:bean id="myAuthFilter" class=
"org.springframework.security.web.authentication.UsernamePasswordAuthenticationFilter">
  <beans:property name="sessionAuthenticationStrategy" ref="sas" />
  ...
</beans:bean>

<beans:bean id="sas" class=
"org.springframework.security.web.authentication.session.SessionFixationProtectionStrategy" />

```

Note that the use of the default, `SessionFixationProtectionStrategy` may cause issues if you are storing beans in the session which implement `HttpSessionBindingListener`, including Spring session-scoped beans. See the Javadoc for this class for more information.

Concurrency Control

Spring Security is able to prevent a principal from concurrently authenticating to the same application more than a specified number of times. Many ISVs take advantage of this to enforce licensing, whilst network administrators like this feature because it helps prevent people from sharing login names. You can, for example, stop user "Batman" from logging onto the web application from two different sessions. You can either expire their previous login or you can report an error when they try to log in again, preventing the second login. Note that if you are using the second approach, a user who has not explicitly logged out (but who has just closed their browser, for example) will not be able to log in again until their original session expires.

Concurrency control is supported by the namespace, so please check the earlier namespace chapter for the simplest configuration. Sometimes you need to customize things though.

The implementation uses a specialized version of `SessionAuthenticationStrategy`, called `ConcurrentSessionControlAuthenticationStrategy`.

Note

Previously the concurrent authentication check was made by the `ProviderManager`, which could be injected with a `ConcurrentSessionController`. The latter would check if the user was attempting to exceed the number of permitted sessions. However, this approach required that an HTTP session be created in advance, which is undesirable. In Spring Security 3, the user is first authenticated by the `AuthenticationManager` and once they are successfully authenticated, a session is created and the check is made whether they are allowed to have another session open.

To use concurrent session support, you'll need to add the following to `web.xml`:

```

<listener>
  <listener-class>
    org.springframework.security.web.session.HttpSessionEventPublisher
  </listener-class>
</listener>

```

In addition, you will need to add the `ConcurrentSessionFilter` to your `FilterChainProxy`. The `ConcurrentSessionFilter` requires two constructor arguments, `sessionRegistry`, which generally points to an instance of `SessionRegistryImpl`, and `sessionInformationExpiredStrategy`, which defines the strategy to apply when a session has

expired. A configuration using the namespace to create the `FilterChainProxy` and other default beans might look like this:

```
<http>
<custom-filter position="CONCURRENT_SESSION_FILTER" ref="concurrencyFilter" />
<custom-filter position="FORM_LOGIN_FILTER" ref="myAuthFilter" />

<session-management session-authentication-strategy-ref="sas"/>
</http>

<beans:bean id="redirectSessionInformationExpiredStrategy"
class="org.springframework.security.web.session.SimpleRedirectSessionInformationExpiredStrategy">
<beans:constructor-arg name="invalidSessionUrl" value="/session-expired.htm" />
</beans:bean>

<beans:bean id="concurrencyFilter"
class="org.springframework.security.web.session.ConcurrentSessionFilter">
<beans:constructor-arg name="sessionRegistry" ref="sessionRegistry" />
<beans:constructor-
arg name="sessionInformationExpiredStrategy" ref="redirectSessionInformationExpiredStrategy" />
</beans:bean>

<beans:bean id="myAuthFilter" class=
"org.springframework.security.web.authentication.UsernamePasswordAuthenticationFilter">
<beans:property name="sessionAuthenticationStrategy" ref="sas" />
<beans:property name="authenticationManager" ref="authenticationManager" />
</beans:bean>

<beans:bean id="sas" class="org.springframework.security.web.authentication.session.CompositeSessionAuthenticationStrategy">
<beans:constructor-arg>
<beans:list>

<beans:bean class="org.springframework.security.web.authentication.session.ConcurrentSessionControlAuthenticationStrategy">
<beans:constructor-arg ref="sessionRegistry"/>
<beans:property name="maximumSessions" value="1" />
<beans:property name="exceptionIfMaximumExceeded" value="true" />
</beans:bean>

<beans:bean class="org.springframework.security.web.authentication.session.SessionFixationProtectionStrategy">
</beans:bean>

<beans:bean class="org.springframework.security.web.authentication.session.RegisterSessionAuthenticationStrategy">
<beans:constructor-arg ref="sessionRegistry"/>
</beans:bean>
</beans:list>
</beans:constructor-arg>
</beans:bean>

<beans:bean id="sessionRegistry"
class="org.springframework.security.core.session.SessionRegistryImpl" />
```

Adding the listener to `web.xml` causes an `ApplicationEvent` to be published to the Spring `ApplicationContext` every time a `HttpSession` commences or terminates. This is critical, as it allows the `SessionRegistryImpl` to be notified when a session ends. Without it, a user will never be able to log back in again once they have exceeded their session allowance, even if they log out of another session or it times out.

Querying the `SessionRegistry` for currently authenticated users and their sessions

Setting up concurrency-control, either through the namespace or using plain beans has the useful side effect of providing you with a reference to the `SessionRegistry` which you can use directly within your application, so even if you don't want to restrict the number of sessions a user may have, it may be worth setting up the infrastructure anyway. You can set the `maximumSession` property to `-1` to allow unlimited sessions. If you're using the namespace, you can set an alias for the internally-created

`SessionRegistry` using the `session-registry-alias` attribute, providing a reference which you can inject into your own beans.

The `getAllPrincipals()` method supplies you with a list of the currently authenticated users. You can list a user's sessions by calling the `getAllSessions(Object principal, boolean includeExpiredSessions)` method, which returns a list of `SessionInformation` objects. You can also expire a user's session by calling `expireNow()` on a `SessionInformation` instance. When the user returns to the application, they will be prevented from proceeding. You may find these methods useful in an administration application, for example. Have a look at the Javadoc for more information.

10.13 Remember-Me Authentication

Overview

Remember-me or persistent-login authentication refers to web sites being able to remember the identity of a principal between sessions. This is typically accomplished by sending a cookie to the browser, with the cookie being detected during future sessions and causing automated login to take place. Spring Security provides the necessary hooks for these operations to take place, and has two concrete remember-me implementations. One uses hashing to preserve the security of cookie-based tokens and the other uses a database or other persistent storage mechanism to store the generated tokens.

Note that both implementations require a `UserDetailsService`. If you are using an authentication provider which doesn't use a `UserDetailsService` (for example, the LDAP provider) then it won't work unless you also have a `UserDetailsService` bean in your application context.

Simple Hash-Based Token Approach

This approach uses hashing to achieve a useful remember-me strategy. In essence a cookie is sent to the browser upon successful interactive authentication, with the cookie being composed as follows:

```
base64(username + ":" + expirationTime + ":" +
md5Hex(username + ":" + expirationTime + ":" password + ":" + key))
```

username:	As identifiable to the <code>UserDetailsService</code>
password:	That matches the one in the retrieved <code>UserDetails</code>
expirationTime:	The date and time when the remember-me token expires, expressed in milliseconds
key:	A private key to prevent modification of the remember-me token

As such the remember-me token is valid only for the period specified, and provided that the username, password and key does not change. Notably, this has a potential security issue in that a captured remember-me token will be usable from any user agent until such time as the token expires. This is the same issue as with digest authentication. If a principal is aware a token has been captured, they can easily change their password and immediately invalidate all remember-me tokens on issue. If more significant security is needed you should use the approach described in the next section. Alternatively remember-me services should simply not be used at all.

If you are familiar with the topics discussed in the chapter on [namespace configuration](#), you can enable remember-me authentication just by adding the `<remember-me>` element:

```
<http>
...
<remember-me key="myAppKey"/>
</http>
```

The `UserDetailsService` will normally be selected automatically. If you have more than one in your application context, you need to specify which one should be used with the `user-service-ref` attribute, where the value is the name of your `UserDetailsService` bean.

Persistent Token Approach

This approach is based on the article http://jaspan.com/improved_persistent_login_cookie_best_practice with some minor modifications²⁰. To use the this approach with namespace configuration, you would supply a `datasource` reference:

```
<http>
...
<remember-me data-source-ref="someDataSource"/>
</http>
```

The database should contain a `persistent_logins` table, created using the following SQL (or equivalent):

```
create table persistent_logins (username varchar(64) not null,
                             series varchar(64) primary key,
                             token varchar(64) not null,
                             last_used timestamp not null)
```

Remember-Me Interfaces and Implementations

Remember-me is used with `UsernamePasswordAuthenticationFilter`, and is implemented via hooks in the `AbstractAuthenticationProcessingFilter` superclass. It is also used within `BasicAuthenticationFilter`. The hooks will invoke a concrete `RememberMeServices` at the appropriate times. The interface looks like this:

```
Authentication autoLogin(HttpServletRequest request, HttpServletResponse response);

void loginFail(HttpServletRequest request, HttpServletResponse response);

void loginSuccess(HttpServletRequest request, HttpServletResponse response,
                  Authentication successfulAuthentication);
```

Please refer to the Javadoc for a fuller discussion on what the methods do, although note at this stage that `AbstractAuthenticationProcessingFilter` only calls the `loginFail()` and `loginSuccess()` methods. The `autoLogin()` method is called by `RememberMeAuthenticationFilter` whenever the `SecurityContextHolder` does not contain an `Authentication`. This interface therefore provides the underlying remember-me implementation with sufficient notification of authentication-related events, and delegates to the implementation whenever a candidate web request might contain a cookie and wish to be remembered. This design allows any number of remember-me implementation strategies. We've seen above that Spring Security provides two implementations. We'll look at these in turn.

TokenBasedRememberMeServices

This implementation supports the simpler approach described in the section called "Simple Hash-Based Token Approach". `TokenBasedRememberMeServices` generates a `RememberMeAuthenticationToken`, which is processed by `RememberMeAuthenticationProvider`. A key is shared between this authentication provider

²⁰Essentially, the username is not included in the cookie, to prevent exposing a valid login name unnecessarily. There is a discussion on this in the comments section of this article.

and the `TokenBasedRememberMeServices`. In addition, `TokenBasedRememberMeServices` requires a `UserDetailsService` from which it can retrieve the username and password for signature comparison purposes, and generate the `RememberMeAuthenticationToken` to contain the correct `GrantedAuthority`s. Some sort of logout command should be provided by the application that invalidates the cookie if the user requests this. `TokenBasedRememberMeServices` also implements Spring Security's `LogoutHandler` interface so can be used with `LogoutFilter` to have the cookie cleared automatically.

The beans required in an application context to enable remember-me services are as follows:

```
<bean id="rememberMeFilter" class=
"org.springframework.security.web.authentication.rememberme.RememberMeAuthenticationFilter">
<property name="rememberMeServices" ref="rememberMeServices"/>
<property name="authenticationManager" ref="theAuthenticationManager" />
</bean>

<bean id="rememberMeServices" class=
"org.springframework.security.web.authentication.rememberme.TokenBasedRememberMeServices">
<property name="userDetailsService" ref="myUserDetailsService"/>
<property name="key" value="springRocks"/>
</bean>

<bean id="rememberMeAuthenticationProvider" class=
"org.springframework.security.authentication.RememberMeAuthenticationProvider">
<property name="key" value="springRocks"/>
</bean>
```

Don't forget to add your `RememberMeServices` implementation to your `UsernamePasswordAuthenticationFilter.setRememberMeServices()` property, include the `RememberMeAuthenticationProvider` in your `AuthenticationManager.setProviders()` list, and add `RememberMeAuthenticationFilter` into your `FilterChainProxy` (typically immediately after your `UsernamePasswordAuthenticationFilter`).

PersistentTokenBasedRememberMeServices

This class can be used in the same way as `TokenBasedRememberMeServices`, but it additionally needs to be configured with a `PersistentTokenRepository` to store the tokens. There are two standard implementations.

- `InMemoryTokenRepositoryImpl` which is intended for testing only.
- `JdbcTokenRepositoryImpl` which stores the tokens in a database.

The database schema is described above in the section called "Persistent Token Approach".

10.14 OpenID Support

The namespace supports [OpenID](#) login either instead of, or in addition to normal form-based login, with a simple change:

```
<http>
<intercept-url pattern="/**" access="ROLE_USER" />
<openid-login />
</http>
```

You should then register yourself with an OpenID provider (such as [myopenid.com](#)), and add the user information to your in-memory `<user-service>`:

```
<user name="https://jimi.hendrix.myopenid.com/" authorities="ROLE_USER" />
```

You should be able to login using the `myopenid.com` site to authenticate. It is also possible to select a specific `UserDetailsService` bean for use OpenID by setting the `user-service-ref` attribute on the `openid-login` element. See the previous section on [authentication providers](#) for more information. Note that we have omitted the password attribute from the above user configuration, since this set of user data is only being used to load the authorities for the user. A random password will be generated internally, preventing you from accidentally using this user data as an authentication source elsewhere in your configuration.

Attribute Exchange

Support for OpenID [attribute exchange](#). As an example, the following configuration would attempt to retrieve the email and full name from the OpenID provider, for use by the application:

```
<openid-login>
<attribute-exchange>
  <openid-attribute name="email" type="https://axschema.org/contact/email" required="true"/>
  <openid-attribute name="name" type="https://axschema.org/namePerson"/>
</attribute-exchange>
</openid-login>
```

The "type" of each OpenID attribute is a URI, determined by a particular schema, in this case <https://axschema.org/>. If an attribute must be retrieved for successful authentication, the `required` attribute can be set. The exact schema and attributes supported will depend on your OpenID provider. The attribute values are returned as part of the authentication process and can be accessed afterwards using the following code:

```
OpenIDAuthenticationToken token =
    (OpenIDAuthenticationToken)SecurityContextHolder.getContext().getAuthentication();
List<OpenIDAttribute> attributes = token.getAttributes();
```

The `OpenIDAttribute` contains the attribute type and the retrieved value (or values in the case of multi-valued attributes). We'll see more about how the `SecurityContextHolder` class is used when we look at core Spring Security components in the [technical overview](#) chapter. Multiple attribute exchange configurations are also supported, if you wish to use multiple identity providers. You can supply multiple `attribute-exchange` elements, using an `identifier-matcher` attribute on each. This contains a regular expression which will be matched against the OpenID identifier supplied by the user. See the OpenID sample application in the codebase for an example configuration, providing different attribute lists for the Google, Yahoo and MyOpenID providers.

10.15 Anonymous Authentication

Overview

It's generally considered good security practice to adopt a "deny-by-default" where you explicitly specify what is allowed and disallow everything else. Defining what is accessible to unauthenticated users is a similar situation, particularly for web applications. Many sites require that users must be authenticated for anything other than a few URLs (for example the home and login pages). In this case it is easiest to define access configuration attributes for these specific URLs rather than have for every secured resource. Put differently, sometimes it is nice to say `ROLE_SOMETHING` is required by default and only allow certain exceptions to this rule, such as for login, logout and home pages of an application. You could also omit these pages from the filter chain entirely, thus bypassing the access control checks, but

this may be undesirable for other reasons, particularly if the pages behave differently for authenticated users.

This is what we mean by anonymous authentication. Note that there is no real conceptual difference between a user who is "anonymously authenticated" and an unauthenticated user. Spring Security's anonymous authentication just gives you a more convenient way to configure your access-control attributes. Calls to servlet API calls such as `getCallerPrincipal`, for example, will still return null even though there is actually an anonymous authentication object in the `SecurityContextHolder`.

There are other situations where anonymous authentication is useful, such as when an auditing interceptor queries the `SecurityContextHolder` to identify which principal was responsible for a given operation. Classes can be authored more robustly if they know the `SecurityContextHolder` always contains an `Authentication` object, and never null.

Configuration

Anonymous authentication support is provided automatically when using the HTTP configuration Spring Security 3.0 and can be customized (or disabled) using the `<anonymous>` element. You don't need to configure the beans described here unless you are using traditional bean configuration.

Three classes that together provide the anonymous authentication feature. `AnonymousAuthenticationToken` is an implementation of `Authentication`, and stores the `GrantedAuthority`s which apply to the anonymous principal. There is a corresponding `AnonymousAuthenticationProvider`, which is chained into the `ProviderManager` so that `AnonymousAuthenticationToken`s are accepted. Finally, there is an `AnonymousAuthenticationFilter`, which is chained after the normal authentication mechanisms and automatically adds an `AnonymousAuthenticationToken` to the `SecurityContextHolder` if there is no existing `Authentication` held there. The definition of the filter and authentication provider appears as follows:

```
<bean id="anonymousAuthFilter"
      class="org.springframework.security.web.authentication.AnonymousAuthenticationFilter">
  <property name="key" value="foobar"/>
  <property name="userAttribute" value="anonymousUser,ROLE_ANONYMOUS"/>
</bean>

<bean id="anonymousAuthenticationProvider"
      class="org.springframework.security.authentication.AnonymousAuthenticationProvider">
  <property name="key" value="foobar"/>
</bean>
```

The `key` is shared between the filter and authentication provider, so that tokens created by the former are accepted by the latter²³. The `userAttribute` is expressed in the form of `usernameInTheAuthenticationToken, grantedAuthority[, grantedAuthority]`. This is the same syntax as used after the equals sign for the `userMap` property of `InMemoryDaoImpl`.

As explained earlier, the benefit of anonymous authentication is that all URI patterns can have security applied to them. For example:

²³The use of the `key` property should not be regarded as providing any real security here. It is merely a book-keeping exercise. If you are sharing a `ProviderManager` which contains an `AnonymousAuthenticationProvider` in a scenario where it is possible for an authenticating client to construct the `Authentication` object (such as with RMI invocations), then a malicious client could submit an `AnonymousAuthenticationToken` which it had created itself (with chosen username and authority list). If the `key` is guessable or can be found out, then the token would be accepted by the anonymous provider. This isn't a problem with normal usage but if you are using RMI you would be best to use a customized `ProviderManager` which omits the anonymous provider rather than sharing the one you use for your HTTP authentication mechanisms.

```

<bean id="filterSecurityInterceptor"
  class="org.springframework.security.web.access.intercept.FilterSecurityInterceptor">
  <property name="authenticationManager" ref="authenticationManager"/>
  <property name="accessDecisionManager" ref="httpRequestAccessDecisionManager"/>
  <property name="securityMetadata">
    <security:filter-security-metadata-source>
    <security:intercept-url pattern='/index.jsp' access='ROLE_ANONYMOUS,ROLE_USER' />
    <security:intercept-url pattern='/hello.htm' access='ROLE_ANONYMOUS,ROLE_USER' />
    <security:intercept-url pattern='/logout.jsp' access='ROLE_ANONYMOUS,ROLE_USER' />
    <security:intercept-url pattern='/login.jsp' access='ROLE_ANONYMOUS,ROLE_USER' />
    <security:intercept-url pattern='/**' access='ROLE_USER' />
    </security:filter-security-metadata-source> " +
  </property>
</bean>

```

AuthenticationTrustResolver

Rounding out the anonymous authentication discussion is the `AuthenticationTrustResolver` interface, with its corresponding `AuthenticationTrustResolverImpl` implementation. This interface provides an `isAnonymous(Authentication)` method, which allows interested classes to take into account this special type of authentication status. The `ExceptionTranslationFilter` uses this interface in processing `AccessDeniedExceptions`. If an `AccessDeniedException` is thrown, and the authentication is of an anonymous type, instead of throwing a 403 (forbidden) response, the filter will instead commence the `AuthenticationEntryPoint` so the principal can authenticate properly. This is a necessary distinction, otherwise principals would always be deemed "authenticated" and never be given an opportunity to login via form, basic, digest or some other normal authentication mechanism.

You will often see the `ROLE_ANONYMOUS` attribute in the above interceptor configuration replaced with `IS_AUTHENTICATED_ANONYMOUSLY`, which is effectively the same thing when defining access controls. This is an example of the use of the `AuthenticatedVoter` which we will see in the [authorization chapter](#). It uses an `AuthenticationTrustResolver` to process this particular configuration attribute and grant access to anonymous users. The `AuthenticatedVoter` approach is more powerful, since it allows you to differentiate between anonymous, remember-me and fully-authenticated users. If you don't need this functionality though, then you can stick with `ROLE_ANONYMOUS`, which will be processed by Spring Security's standard `RoleVoter`.

10.16 Pre-Authentication Scenarios

There are situations where you want to use Spring Security for authorization, but the user has already been reliably authenticated by some external system prior to accessing the application. We refer to these situations as "pre-authenticated" scenarios. Examples include X.509, Siteminder and authentication by the Java EE container in which the application is running. When using pre-authentication, Spring Security has to

- Identify the user making the request.
- Obtain the authorities for the user.

The details will depend on the external authentication mechanism. A user might be identified by their certificate information in the case of X.509, or by an HTTP request header in the case of Siteminder. If relying on container authentication, the user will be identified by calling the `getUserPrincipal()` method on the incoming HTTP request. In some cases, the external mechanism may supply role/authority information for the user but in others the authorities must be obtained from a separate source, such as a `UserDetailsService`.

Pre-Authentication Framework Classes

Because most pre-authentication mechanisms follow the same pattern, Spring Security has a set of classes which provide an internal framework for implementing pre-authenticated authentication providers. This removes duplication and allows new implementations to be added in a structured fashion, without having to write everything from scratch. You don't need to know about these classes if you want to use something like [X.509 authentication](#), as it already has a namespace configuration option which is simpler to use and get started with. If you need to use explicit bean configuration or are planning on writing your own implementation then an understanding of how the provided implementations work will be useful. You will find classes under the `org.springframework.security.web.authentication.preauth`. We just provide an outline here so you should consult the Javadoc and source where appropriate.

AbstractPreAuthenticatedProcessingFilter

This class will check the current contents of the security context and, if empty, it will attempt to extract user information from the HTTP request and submit it to the `AuthenticationManager`. Subclasses override the following methods to obtain this information:

```
protected abstract Object getPreAuthenticatedPrincipal(HttpServletRequest request);
protected abstract Object getPreAuthenticatedCredentials(HttpServletRequest request);
```

After calling these, the filter will create a `PreAuthenticatedAuthenticationToken` containing the returned data and submit it for authentication. By "authentication" here, we really just mean further processing to perhaps load the user's authorities, but the standard Spring Security authentication architecture is followed.

Like other Spring Security authentication filters, the pre-authentication filter has an `authenticationDetailsSource` property which by default will create a `WebAuthenticationDetails` object to store additional information such as the session-identifier and originating IP address in the `details` property of the `Authentication` object. In cases where user role information can be obtained from the pre-authentication mechanism, the data is also stored in this property, with the details implementing the `GrantedAuthoritiesContainer` interface. This enables the authentication provider to read the authorities which were externally allocated to the user. We'll look at a concrete example next.

J2eeBasedPreAuthenticatedWebAuthenticationDetailsSource

If the filter is configured with an `authenticationDetailsSource` which is an instance of this class, the authority information is obtained by calling the `isUserInRole(String role)` method for each of a pre-determined set of "mappable roles". The class gets these from a configured `MappableAttributesRetriever`. Possible implementations include hard-coding a list in the application context and reading the role information from the `<security-role>` information in a `web.xml` file. The pre-authentication sample application uses the latter approach.

There is an additional stage where the roles (or attributes) are mapped to Spring Security `GrantedAuthority` objects using a configured `Attributes2GrantedAuthoritiesMapper`. The default will just add the usual `ROLE_` prefix to the names, but it gives you full control over the behaviour.

PreAuthenticatedAuthenticationProvider

The pre-authenticated provider has little more to do than load the `UserDetails` object for the user. It does this by delegating to an `AuthenticationUserDetailsService`. The latter is similar to the standard `UserDetailsService` but takes an `Authentication` object rather than just user name:


```
public interface AuthenticationUserDetailsService {
    UserDetails loadUserDetails(Authentication token) throws UsernameNotFoundException;
}
```

This interface may have also other uses but with pre-authentication it allows access to the authorities which were packaged in the `Authentication` object, as we saw in the previous section. The `PreAuthenticatedGrantedAuthoritiesUserDetailsService` class does this. Alternatively, it may delegate to a standard `UserDetailsService` via the `UserDetailsServiceWrapper` implementation.

Http403ForbiddenEntryPoint

The `AuthenticationEntryPoint` was discussed in the [technical overview](#) chapter. Normally it is responsible for kick-starting the authentication process for an unauthenticated user (when they try to access a protected resource), but in the pre-authenticated case this doesn't apply. You would only configure the `ExceptionHandlerFilter` with an instance of this class if you aren't using pre-authentication in combination with other authentication mechanisms. It will be called if the user is rejected by the `AbstractPreAuthenticatedProcessingFilter` resulting in a null authentication. It always returns a 403-forbidden response code if called.

Concrete Implementations

X.509 authentication is covered in its [own chapter](#). Here we'll look at some classes which provide support for other pre-authenticated scenarios.

Request-Header Authentication (Siteminder)

An external authentication system may supply information to the application by setting specific headers on the HTTP request. A well-known example of this is Siteminder, which passes the username in a header called `SM_USER`. This mechanism is supported by the class `RequestHeaderAuthenticationFilter` which simply extracts the username from the header. It defaults to using the name `SM_USER` as the header name. See the Javadoc for more details.

Tip

Note that when using a system like this, the framework performs no authentication checks at all and it is *extremely* important that the external system is configured properly and protects all access to the application. If an attacker is able to forge the headers in their original request without this being detected then they could potentially choose any username they wished.

Siteminder Example Configuration

A typical configuration using this filter would look like this:

```

<security:http>
<!-- Additional http configuration omitted -->
<security:custom-filter position="PRE_AUTH_FILTER" ref="siteminderFilter" />
</security:http>

<bean id="siteminderFilter" class="org.springframework.security.web.authentication.preauth.RequestHeaderAuthenticationFilter">
<property name="principalRequestHeader" value="SM_USER"/>
<property name="authenticationManager" ref="authenticationManager" />
</bean>

<bean id="preauthAuthProvider" class="org.springframework.security.web.authentication.preauth.PreAuthenticatedAuthenticationProvider">
<property name="preAuthenticatedUserDetailsService">
<bean id="userDetailsServiceWrapper"
class="org.springframework.security.core.userdetails.UserDetailsServiceWrapper">
<property name="userDetailsService" ref="userDetailsService"/>
</bean>
</property>
</bean>

<security:authentication-manager alias="authenticationManager">
<security:authentication-provider ref="preauthAuthProvider" />
</security:authentication-manager>

```

We've assumed here that the [security namespace](#) is being used for configuration. It's also assumed that you have added a `UserDetailsService` (called "userDetailsService") to your configuration to load the user's roles.

Java EE Container Authentication

The class `J2eePreAuthenticatedProcessingFilter` will extract the username from the `userPrincipal` property of the `HttpServletRequest`. Use of this filter would usually be combined with the use of Java EE roles as described above in the section called "J2eeBasedPreAuthenticatedWebAuthenticationDetailsSource".

There is a sample application in the codebase which uses this approach, so get hold of the code from [github](#) and have a look at the application context file if you are interested. The code is in the `samples/xml/preauth` directory.

10.17 Java Authentication and Authorization Service (JAAS) Provider

Overview

Spring Security provides a package able to delegate authentication requests to the Java Authentication and Authorization Service (JAAS). This package is discussed in detail below.

AbstractJaasAuthenticationProvider

The `AbstractJaasAuthenticationProvider` is the basis for the provided JAAS `AuthenticationProvider` implementations. Subclasses must implement a method that creates the `LoginContext`. The `AbstractJaasAuthenticationProvider` has a number of dependencies that can be injected into it that are discussed below.

JAAS CallbackHandler

Most JAAS `LoginModule`s require a callback of some sort. These callbacks are usually used to obtain the username and password from the user.

In a Spring Security deployment, Spring Security is responsible for this user interaction (via the authentication mechanism). Thus, by the time the authentication request is delegated through to JAAS, Spring Security's authentication mechanism will already have fully-populated an `Authentication` object containing all the information required by the JAAS `LoginModule`.

Therefore, the JAAS package for Spring Security provides two default callback handlers, `JaasNameCallbackHandler` and `JaasPasswordCallbackHandler`. Each of these callback handlers implement `JaasAuthenticationCallbackHandler`. In most cases these callback handlers can simply be used without understanding the internal mechanics.

For those needing full control over the callback behavior, internally `AbstractJaasAuthenticationProvider` wraps these `JaasAuthenticationCallbackHandler`s with an `InternalCallbackHandler`. The `InternalCallbackHandler` is the class that actually implements JAAS normal `CallbackHandler` interface. Any time that the JAAS `LoginModule` is used, it is passed a list of application context configured `InternalCallbackHandler`s. If the `LoginModule` requests a callback against the `InternalCallbackHandler`s, the callback is in-turn passed to the `JaasAuthenticationCallbackHandler`s being wrapped.

JAAS AuthorityGranter

JAAS works with principals. Even "roles" are represented as principals in JAAS. Spring Security, on the other hand, works with `Authentication` objects. Each `Authentication` object contains a single principal, and multiple `GrantedAuthority`s. To facilitate mapping between these different concepts, Spring Security's JAAS package includes an `AuthorityGranter` interface.

An `AuthorityGranter` is responsible for inspecting a JAAS principal and returning a set of `String`s, representing the authorities assigned to the principal. For each returned authority string, the `AbstractJaasAuthenticationProvider` creates a `JaasGrantedAuthority` (which implements Spring Security's `GrantedAuthority` interface) containing the authority string and the JAAS principal that the `AuthorityGranter` was passed. The `AbstractJaasAuthenticationProvider` obtains the JAAS principals by firstly successfully authenticating the user's credentials using the JAAS `LoginModule`, and then accessing the `LoginContext` it returns. A call to `LoginContext.getSubject().getPrincipals()` is made, with each resulting principal passed to each `AuthorityGranter` defined against the `AbstractJaasAuthenticationProvider.setAuthorityGranters(List)` property.

Spring Security does not include any production `AuthorityGranter`s given that every JAAS principal has an implementation-specific meaning. However, there is a `TestAuthorityGranter` in the unit tests that demonstrates a simple `AuthorityGranter` implementation.

DefaultJaasAuthenticationProvider

The `DefaultJaasAuthenticationProvider` allows a JAAS `Configuration` object to be injected into it as a dependency. It then creates a `LoginContext` using the injected JAAS `Configuration`. This means that `DefaultJaasAuthenticationProvider` is not bound any particular implementation of `Configuration` as `JaasAuthenticationProvider` is.

InMemoryConfiguration

In order to make it easy to inject a `Configuration` into `DefaultJaasAuthenticationProvider`, a default in-memory implementation named `InMemoryConfiguration` is provided. The implementation constructor accepts a `Map` where each key represents a login configuration name and the value represents an `Array` of `AppConfigurationEntry`s. `InMemoryConfiguration` also supports a

default Array of `AppConfigurationEntry` objects that will be used if no mapping is found within the provided Map. For details, refer to the class level javadoc of `InMemoryConfiguration`.

DefaultJaasAuthenticationProvider Example Configuration

While the Spring configuration for `InMemoryConfiguration` can be more verbose than the standard JAAS configuration files, using it in conjunction with `DefaultJaasAuthenticationProvider` is more flexible than `JaasAuthenticationProvider` since it not dependant on the default Configuration implementation.

An example configuration of `DefaultJaasAuthenticationProvider` using `InMemoryConfiguration` is provided below. Note that custom implementations of Configuration can easily be injected into `DefaultJaasAuthenticationProvider` as well.

```
<bean id="jaasAuthProvider"
class="org.springframework.security.authentication.jaas.DefaultJaasAuthenticationProvider">
<property name="configuration">
<bean class="org.springframework.security.authentication.jaas.memory.InMemoryConfiguration">
<constructor-arg>
<map>
<!--
SPRINGSECURITY is the default loginContextName
for AbstractJaasAuthenticationProvider
-->
<entry key="SPRINGSECURITY">
<array>
<bean class="javax.security.auth.login.AppConfigurationEntry">
<constructor-arg value="sample.SampleLoginModule" />
<constructor-arg>
<util:constant static-field=
"javax.security.auth.login.AppConfigurationEntry$LoginModuleControlFlag.REQUIRED"/>
</constructor-arg>
<constructor-arg>
<map></map>
</constructor-arg>
</bean>
</array>
</entry>
</map>
</constructor-arg>
</bean>
</property>
<property name="authorityGranters">
<list>
<!-- You will need to write your own implementation of AuthorityGranter -->
<bean class="org.springframework.security.authentication.jaas.TestAuthorityGranter"/>
</list>
</property>
</bean>
```

JaasAuthenticationProvider

The `JaasAuthenticationProvider` assumes the default Configuration is an instance of `ConfigFile`. This assumption is made in order to attempt to update the Configuration. The `JaasAuthenticationProvider` then uses the default Configuration to create the `LoginContext`.

Let's assume we have a JAAS login configuration file, `/WEB-INF/login.conf`, with the following contents:

```
JAASTest {
    sample.SampleLoginModule required;
};
```

Like all Spring Security beans, the `JaasAuthenticationProvider` is configured via the application context. The following definitions would correspond to the above JAAS login configuration file:

```
<bean id="jaasAuthenticationProvider"
class="org.springframework.security.authentication.jaas.JaasAuthenticationProvider">
<property name="loginConfig" value="/WEB-INF/login.conf"/>
<property name="loginContextName" value="JAASTest"/>
<property name="callbackHandlers">
<list>
<bean
class="org.springframework.security.authentication.jaas.JaasNameCallbackHandler"/>
<bean
class="org.springframework.security.authentication.jaas.JaasPasswordCallbackHandler"/>
</list>
</property>
<property name="authorityGranters">
<list>
<bean class="org.springframework.security.authentication.jaas.TestAuthorityGranter"/>
</list>
</property>
</bean>
```

Running as a Subject

If configured, the `JaasApiIntegrationFilter` will attempt to run as the Subject on the `JaasAuthenticationToken`. This means that the Subject can be accessed using:

```
Subject subject = Subject.getSubject(AccessController.getContext());
```

This integration can easily be configured using the [jaas-api-provision](#) attribute. This feature is useful when integrating with legacy or external API's that rely on the JAAS Subject being populated.

10.18 CAS Authentication

Overview

JA-SIG produces an enterprise-wide single sign on system known as CAS. Unlike other initiatives, JA-SIG's Central Authentication Service is open source, widely used, simple to understand, platform independent, and supports proxy capabilities. Spring Security fully supports CAS, and provides an easy migration path from single-application deployments of Spring Security through to multiple-application deployments secured by an enterprise-wide CAS server.

You can learn more about CAS at <https://www.apereo.org>. You will also need to visit this site to download the CAS Server files.

How CAS Works

Whilst the CAS web site contains documents that detail the architecture of CAS, we present the general overview again here within the context of Spring Security. Spring Security 3.x supports CAS 3. At the time of writing, the CAS server was at version 3.4.

Somewhere in your enterprise you will need to setup a CAS server. The CAS server is simply a standard WAR file, so there isn't anything difficult about setting up your server. Inside the WAR file you will customise the login and other single sign on pages displayed to users.

When deploying a CAS 3.4 server, you will also need to specify an `AuthenticationHandler` in the `deployerConfigContext.xml` included with CAS. The `AuthenticationHandler` has a simple method that returns a boolean as to whether a given set of `Credentials` is valid. Your `AuthenticationHandler` implementation will need to link into some type of backend

authentication repository, such as an LDAP server or database. CAS itself includes numerous `AuthenticationHandler`s out of the box to assist with this. When you download and deploy the server war file, it is set up to successfully authenticate users who enter a password matching their username, which is useful for testing.

Apart from the CAS server itself, the other key players are of course the secure web applications deployed throughout your enterprise. These web applications are known as "services". There are three types of services. Those that authenticate service tickets, those that can obtain proxy tickets, and those that authenticate proxy tickets. Authenticating a proxy ticket differs because the list of proxies must be validated and often times a proxy ticket can be reused.

Spring Security and CAS Interaction Sequence

The basic interaction between a web browser, CAS server and a Spring Security-secured service is as follows:

- The web user is browsing the service's public pages. CAS or Spring Security is not involved.
- The user eventually requests a page that is either secure or one of the beans it uses is secure. Spring Security's `ExceptionTranslationFilter` will detect the `AccessDeniedException` or `AuthenticationException`.
- Because the user's `Authentication` object (or lack thereof) caused an `AuthenticationException`, the `ExceptionTranslationFilter` will call the configured `AuthenticationEntryPoint`. If using CAS, this will be the `CasAuthenticationEntryPoint` class.
- The `CasAuthenticationEntryPoint` will redirect the user's browser to the CAS server. It will also indicate a service parameter, which is the callback URL for the Spring Security service (your application). For example, the URL to which the browser is redirected might be <https://my.company.com/cas/login?service=https%3A%2F%2Fserver3.company.com%2Fwebapp%2Flogin/cas>.
- After the user's browser redirects to CAS, they will be prompted for their username and password. If the user presents a session cookie which indicates they've previously logged on, they will not be prompted to login again (there is an exception to this procedure, which we'll cover later). CAS will use the `PasswordHandler` (or `AuthenticationHandler` if using CAS 3.0) discussed above to decide whether the username and password is valid.
- Upon successful login, CAS will redirect the user's browser back to the original service. It will also include a `ticket` parameter, which is an opaque string representing the "service ticket". Continuing our earlier example, the URL the browser is redirected to might be <https://server3.company.com/webapp/login/cas?ticket=ST-0-ER94xMJmn6pha35CQRoZ>.
- Back in the service web application, the `CasAuthenticationFilter` is always listening for requests to `/login/cas` (this is configurable, but we'll use the defaults in this introduction). The processing filter will construct a `UsernamePasswordAuthenticationToken` representing the service ticket. The principal will be equal to `CasAuthenticationFilter.CAS_STATEFUL_IDENTIFIER`, whilst the credentials will be the service ticket opaque value. This authentication request will then be handed to the configured `AuthenticationManager`.
- The `AuthenticationManager` implementation will be the `ProviderManager`, which is in turn configured with the `CasAuthenticationProvider`. The `CasAuthenticationProvider`

only responds to `UsernamePasswordAuthenticationToken`s containing the CAS-specific principal (such as `CasAuthenticationFilter.CAS_STATEFUL_IDENTIFIER`) and `CasAuthenticationToken`s (discussed later).

- `CasAuthenticationProvider` will validate the service ticket using a `TicketValidator` implementation. This will typically be a `Cas20ServiceTicketValidator` which is one of the classes included in the CAS client library. In the event the application needs to validate proxy tickets, the `Cas20ProxyTicketValidator` is used. The `TicketValidator` makes an HTTPS request to the CAS server in order to validate the service ticket. It may also include a proxy callback URL, which is included in this example: <https://my.company.com/cas/proxyValidate?service=https%3A%2F%2Fserver3.company.com%2Fwebapp%2Flogin/cas&ticket=ST-0-ER94xMJmn6pha35CQRoZ&pgtUrl=https://server3.company.com/webapp/login/cas/proxyreceptor>.
- Back on the CAS server, the validation request will be received. If the presented service ticket matches the service URL the ticket was issued to, CAS will provide an affirmative response in XML indicating the username. If any proxy was involved in the authentication (discussed below), the list of proxies is also included in the XML response.
- [OPTIONAL] If the request to the CAS validation service included the proxy callback URL (in the `pgtUrl` parameter), CAS will include a `pgtIou` string in the XML response. This `pgtIou` represents a proxy-granting ticket IOU. The CAS server will then create its own HTTPS connection back to the `pgtUrl`. This is to mutually authenticate the CAS server and the claimed service URL. The HTTPS connection will be used to send a proxy granting ticket to the original web application. For example, <https://server3.company.com/webapp/login/cas/proxyreceptor?pgtIou=PGTIOU-0-R0zIgrl4pdAQwBvJWO3vnNpevwqStbSGcq3vKB2SqSFFRnjPHt&pgtId=PGT-1-si9YkkHLrtACBo64rmsi3v2nf7cpCResXg5MpESZFArbaZiOKH>.
- The `Cas20TicketValidator` will parse the XML received from the CAS server. It will return to the `CasAuthenticationProvider` a `TicketResponse`, which includes the username (mandatory), proxy list (if any were involved), and proxy-granting ticket IOU (if the proxy callback was requested).
- Next `CasAuthenticationProvider` will call a configured `CasProxyDecider`. The `CasProxyDecider` indicates whether the proxy list in the `TicketResponse` is acceptable to the service. Several implementations are provided with Spring Security: `RejectProxyTickets`, `AcceptAnyCasProxy` and `NamedCasProxyDecider`. These names are largely self-explanatory, except `NamedCasProxyDecider` which allows a `List` of trusted proxies to be provided.
- `CasAuthenticationProvider` will next request a `AuthenticationUserDetailsService` to load the `GrantedAuthority` objects that apply to the user contained in the `Assertion`.
- If there were no problems, `CasAuthenticationProvider` constructs a `CasAuthenticationToken` including the details contained in the `TicketResponse` and the `GrantedAuthority`s.
- Control then returns to `CasAuthenticationFilter`, which places the created `CasAuthenticationToken` in the security context.
- The user's browser is redirected to the original page that caused the `AuthenticationException` (or a [custom destination](#) depending on the configuration).

It's good that you're still here! Let's now look at how this is configured

Configuration of CAS Client

The web application side of CAS is made easy due to Spring Security. It is assumed you already know the basics of using Spring Security, so these are not covered again below. We'll assume a namespace based configuration is being used and add in the CAS beans as required. Each section builds upon the previous section. A full [CAS sample application](#) can be found in the Spring Security Samples.

Service Ticket Authentication

This section describes how to setup Spring Security to authenticate Service Tickets. Often times this is all a web application requires. You will need to add a `ServiceProperties` bean to your application context. This represents your CAS service:

```
<bean id="serviceProperties"
      class="org.springframework.security.cas.ServiceProperties">
  <property name="service"
    value="https://localhost:8443/cas-sample/login/cas"/>
  <property name="sendRenew" value="false"/>
</bean>
```

The `service` must equal a URL that will be monitored by the `CasAuthenticationFilter`. The `sendRenew` defaults to false, but should be set to true if your application is particularly sensitive. What this parameter does is tell the CAS login service that a single sign on login is unacceptable. Instead, the user will need to re-enter their username and password in order to gain access to the service.

The following beans should be configured to commence the CAS authentication process (assuming you're using a namespace configuration):

```
<security:http entry-point-ref="casEntryPoint">
  ...
  <security:custom-filter position="CAS_FILTER" ref="casFilter" />
</security:http>

<bean id="casFilter"
      class="org.springframework.security.cas.web.CasAuthenticationFilter">
  <property name="authenticationManager" ref="authenticationManager"/>
</bean>

<bean id="casEntryPoint"
      class="org.springframework.security.cas.web.CasAuthenticationEntryPoint">
  <property name="loginUrl" value="https://localhost:9443/cas/login"/>
  <property name="serviceProperties" ref="serviceProperties"/>
</bean>
```

For CAS to operate, the `ExceptionHandlerFilter` must have its `authenticationEntryPoint` property set to the `CasAuthenticationEntryPoint` bean. This can easily be done using [entry-point-ref](#) as is done in the example above. The `CasAuthenticationEntryPoint` must refer to the `ServiceProperties` bean (discussed above), which provides the URL to the enterprise's CAS login server. This is where the user's browser will be redirected.

The `CasAuthenticationFilter` has very similar properties to the `UsernamePasswordAuthenticationFilter` (used for form-based logins). You can use these properties to customize things like behavior for authentication success and failure.

Next you need to add a `CasAuthenticationProvider` and its collaborators:


```

<security:authentication-manager alias="authenticationManager">
<security:authentication-provider ref="casAuthenticationProvider" />
</security:authentication-manager>

<bean id="casAuthenticationProvider"
    class="org.springframework.security.cas.authentication.CasAuthenticationProvider">
<property name="authenticationUserService">
    <bean class="org.springframework.security.core.userdetails.UserDetailsByNameServiceWrapper">
    <constructor-arg ref="userService" />
    </bean>
</property>
<property name="serviceProperties" ref="serviceProperties" />
<property name="ticketValidator">
    <bean class="org.jasig.cas.client.validation.Cas20ServiceTicketValidator">
    <constructor-arg index="0" value="https://localhost:9443/cas" />
    </bean>
</property>
<property name="key" value="an_id_for_this_auth_provider_only"/>
</bean>

<security:user-service id="userService">
<!-- Password is prefixed with {noop} to indicate to DelegatingPasswordEncoder that
NoOpPasswordEncoder should be used.
This is not safe for production, but makes reading
in samples easier.
Normally passwords should be hashed using BCrypt -->
<security:user name="joe" password="{noop}joe" authorities="ROLE_USER" />
...
</security:user-service>

```

The `CasAuthenticationProvider` uses a `UserDetailsService` instance to load the authorities for a user, once they have been authenticated by CAS. We've shown a simple in-memory setup here. Note that the `CasAuthenticationProvider` does not actually use the password for authentication, but it does use the authorities.

The beans are all reasonably self-explanatory if you refer back to the [How CAS Works](#) section.

This completes the most basic configuration for CAS. If you haven't made any mistakes, your web application should happily work within the framework of CAS single sign on. No other parts of Spring Security need to be concerned about the fact CAS handled authentication. In the following sections we will discuss some (optional) more advanced configurations.

Single Logout

The CAS protocol supports Single Logout and can be easily added to your Spring Security configuration. Below are updates to the Spring Security configuration that handle Single Logout

```

<security:http entry-point-ref="casEntryPoint">
  ...
<security:logout logout-success-url="/cas-logout.jsp"/>
<security:custom-filter ref="requestSingleLogoutFilter" before="LOGOUT_FILTER"/>
<security:custom-filter ref="singleLogoutFilter" before="CAS_FILTER"/>
</security:http>

<!-- This filter handles a Single Logout Request from the CAS Server -->
<bean id="singleLogoutFilter" class="org.jasig.cas.client.session.SingleSignOutFilter"/>

<!-- This filter redirects to the CAS Server to signal Single Logout should be performed -->
<bean id="requestSingleLogoutFilter"
  class="org.springframework.security.web.authentication.logout.LogoutFilter">
  <constructor-arg value="https://localhost:9443/cas/logout"/>
  <constructor-arg>
    <bean class=
      "org.springframework.security.web.authentication.logout.SecurityContextLogoutHandler"/>
  </constructor-arg>
  <property name="filterProcessesUrl" value="/logout/cas"/>
</bean>

```

The `logout` element logs the user out of the local application, but does not terminate the session with the CAS server or any other applications that have been logged into. The `requestSingleLogoutFilter` filter will allow the URL of `/spring_security_cas_logout` to be requested to redirect the application to the configured CAS Server logout URL. Then the CAS Server will send a Single Logout request to all the services that were signed into. The `singleLogoutFilter` handles the Single Logout request by looking up the `HttpSession` in a static `Map` and then invalidating it.

It might be confusing why both the `logout` element and the `singleLogoutFilter` are needed. It is considered best practice to logout locally first since the `SingleSignOutFilter` just stores the `HttpSession` in a static `Map` in order to call `invalidate` on it. With the configuration above, the flow of logout would be:

- The user requests `/logout` which would log the user out of the local application and send the user to the logout success page.
- The logout success page, `/cas-logout.jsp`, should instruct the user to click a link pointing to `/logout/cas` in order to logout out of all applications.
- When the user clicks the link, the user is redirected to the CAS single logout URL (<https://localhost:9443/cas/logout>).
- On the CAS Server side, the CAS single logout URL then submits single logout requests to all the CAS Services. On the CAS Service side, JASIG's `SingleSignOutFilter` processes the logout request by invalidating the original session.

The next step is to add the following to your `web.xml`

```

<filter>
<filter-name>characterEncodingFilter</filter-name>
<filter-class>
    org.springframework.web.filter.CharacterEncodingFilter
</filter-class>
</filter>
<init-param>
    <param-name>encoding</param-name>
    <param-value>UTF-8</param-value>
</init-param>
</filter>
<filter-mapping>
<filter-name>characterEncodingFilter</filter-name>
<url-pattern>/*</url-pattern>
</filter-mapping>
<listener>
<listener-class>
    org.jasig.cas.client.session.SingleSignOutHttpSessionListener
</listener-class>
</listener>

```

When using the `SingleSignOutFilter` you might encounter some encoding issues. Therefore it is recommended to add the `CharacterEncodingFilter` to ensure that the character encoding is correct when using the `SingleSignOutFilter`. Again, refer to JASIG's documentation for details. The `SingleSignOutHttpSessionListener` ensures that when an `HttpSession` expires, the mapping used for single logout is removed.

Authenticating to a Stateless Service with CAS

This section describes how to authenticate to a service using CAS. In other words, this section discusses how to setup a client that uses a service that authenticates with CAS. The next section describes how to setup a stateless service to Authenticate using CAS.

Configuring CAS to Obtain Proxy Granting Tickets

In order to authenticate to a stateless service, the application needs to obtain a proxy granting ticket (PGT). This section describes how to configure Spring Security to obtain a PGT building upon the `thencas-st[Service Ticket Authentication]` configuration.

The first step is to include a `ProxyGrantingTicketStorage` in your Spring Security configuration. This is used to store PGT's that are obtained by the `CasAuthenticationFilter` so that they can be used to obtain proxy tickets. An example configuration is shown below

```

<!--
NOTE: In a real application you should not use an in memory implementation.
You will also want to ensure to clean up expired tickets by calling
ProxyGrantingTicketStorage.cleanup()
-->
<bean id="pgtStorage" class="org.jasig.cas.client.proxy.ProxyGrantingTicketStorageImpl"/>

```

The next step is to update the `CasAuthenticationProvider` to be able to obtain proxy tickets. To do this replace the `Cas20ServiceTicketValidator` with a `Cas20ProxyTicketValidator`. The `proxyCallbackUrl` should be set to a URL that the application will receive PGT's at. Last, the configuration should also reference the `ProxyGrantingTicketStorage` so it can use a PGT to obtain proxy tickets. You can find an example of the configuration changes that should be made below.

```

<bean id="casAuthenticationProvider"
      class="org.springframework.security.cas.authentication.CasAuthenticationProvider">
  ...
  <property name="ticketValidator">
    <bean class="org.jasig.cas.client.validation.Cas20ProxyTicketValidator">
      <constructor-arg value="https://localhost:9443/cas"/>
      <property name="proxyCallbackUrl"
                value="https://localhost:8443/cas-sample/login/cas/proxyreceptor"/>
      <property name="proxyGrantingTicketStorage" ref="pgtStorage"/>
    </bean>
  </property>
</bean>

```

The last step is to update the `CasAuthenticationFilter` to accept PGT and to store them in the `ProxyGrantingTicketStorage`. It is important the `proxyReceptorUrl` matches the `proxyCallbackUrl` of the `Cas20ProxyTicketValidator`. An example configuration is shown below.

```

<bean id="casFilter"
      class="org.springframework.security.cas.web.CasAuthenticationFilter">
  ...
  <property name="proxyGrantingTicketStorage" ref="pgtStorage"/>
  <property name="proxyReceptorUrl" value="/login/cas/proxyreceptor"/>
</bean>

```

Calling a Stateless Service Using a Proxy Ticket

Now that Spring Security obtains PGTs, you can use them to create proxy tickets which can be used to authenticate to a stateless service. The [CAS sample application](#) contains a working example in the `ProxyTicketSampleServlet`. Example code can be found below:

```

protected void doGet(HttpServletRequest request, HttpServletResponse response)
    throws ServletException, IOException {
  // NOTE: The CasAuthenticationToken can also be obtained using
  // SecurityContextHolder.getContext().getAuthentication()
  final CasAuthenticationToken token = (CasAuthenticationToken) request.getUserPrincipal();
  // proxyTicket could be reused to make calls to the CAS service even if the
  // target url differs
  final String proxyTicket = token.getAssertion().getPrincipal().getProxyTicketFor(targetUrl);

  // Make a remote call using the proxy ticket
  final String serviceUrl = targetUrl+"?ticket="+URLEncoder.encode(proxyTicket, "UTF-8");
  String proxyResponse = CommonUtils.getResponseFromServer(serviceUrl, "UTF-8");
  ...
}

```

Proxy Ticket Authentication

The `CasAuthenticationProvider` distinguishes between stateful and stateless clients. A stateful client is considered any that submits to the `filterProcessUrl` of the `CasAuthenticationFilter`. A stateless client is any that presents an authentication request to `CasAuthenticationFilter` on a URL other than the `filterProcessUrl`.

Because remoting protocols have no way of presenting themselves within the context of an `HttpSession`, it isn't possible to rely on the default practice of storing the security context in the session between requests. Furthermore, because the CAS server invalidates a ticket after it has been validated by the `TicketValidator`, presenting the same proxy ticket on subsequent requests will not work.

One obvious option is to not use CAS at all for remoting protocol clients. However, this would eliminate many of the desirable features of CAS. As a middle-ground, the `CasAuthenticationProvider` uses a `StatelessTicketCache`. This is used solely for stateless clients which use a

principal equal to `CasAuthenticationFilter.CAS_STATELESS_IDENTIFIER`. What happens is the `CasAuthenticationProvider` will store the resulting `CasAuthenticationToken` in the `StatelessTicketCache`, keyed on the proxy ticket. Accordingly, remoting protocol clients can present the same proxy ticket and the `CasAuthenticationProvider` will not need to contact the CAS server for validation (aside from the first request). Once authenticated, the proxy ticket could be used for URLs other than the original target service.

This section builds upon the previous sections to accommodate proxy ticket authentication. The first step is to specify to authenticate all artifacts as shown below.

```
<bean id="serviceProperties"
      class="org.springframework.security.cas.ServiceProperties">
  ...
  <property name="authenticateAllArtifacts" value="true"/>
</bean>
```

The next step is to specify `serviceProperties` and the `authenticationDetailsSource` for the `CasAuthenticationFilter`. The `serviceProperties` property instructs the `CasAuthenticationFilter` to attempt to authenticate all artifacts instead of only ones present on the `filterProcessUrl`. The `ServiceAuthenticationDetailsSource` creates a `ServiceAuthenticationDetails` that ensures the current URL, based upon the `HttpServletRequest`, is used as the service URL when validating the ticket. The method for generating the service URL can be customized by injecting a custom `AuthenticationDetailsSource` that returns a custom `ServiceAuthenticationDetails`.

```
<bean id="casFilter"
      class="org.springframework.security.cas.web.CasAuthenticationFilter">
  ...
  <property name="serviceProperties" ref="serviceProperties"/>
  <property name="authenticationDetailsSource">
    <bean class=
      "org.springframework.security.cas.web.authentication.ServiceAuthenticationDetailsSource">
      <constructor-arg ref="serviceProperties"/>
    </bean>
  </property>
</bean>
```

You will also need to update the `CasAuthenticationProvider` to handle proxy tickets. To do this replace the `Cas20ServiceTicketValidator` with a `Cas20ProxyTicketValidator`. You will need to configure the `statelessTicketCache` and which proxies you want to accept. You can find an example of the updates required to accept all proxies below.

```

<bean id="casAuthenticationProvider"
      class="org.springframework.security.cas.authentication.CasAuthenticationProvider">
  ...
  <property name="ticketValidator">
    <bean class="org.jasig.cas.client.validation.Cas20ProxyTicketValidator">
      <constructor-arg value="https://localhost:9443/cas"/>
      <property name="acceptAnyProxy" value="true"/>
    </bean>
  </property>
  <property name="statelessTicketCache">
    <bean class="org.springframework.security.cas.authentication.EhCacheBasedTicketCache">
      <property name="cache">
        <bean class="net.sf.ehcache.Cache"
              init-method="initialise" destroy-method="dispose">
          <constructor-arg value="casTickets"/>
          <constructor-arg value="50"/>
          <constructor-arg value="true"/>
          <constructor-arg value="false"/>
          <constructor-arg value="3600"/>
          <constructor-arg value="900"/>
        </bean>
      </property>
    </bean>
  </property>
</bean>

```

10.19 X.509 Authentication

Overview

The most common use of X.509 certificate authentication is in verifying the identity of a server when using SSL, most commonly when using HTTPS from a browser. The browser will automatically check that the certificate presented by a server has been issued (ie digitally signed) by one of a list of trusted certificate authorities which it maintains.

You can also use SSL with "mutual authentication"; the server will then request a valid certificate from the client as part of the SSL handshake. The server will authenticate the client by checking that its certificate is signed by an acceptable authority. If a valid certificate has been provided, it can be obtained through the servlet API in an application. Spring Security X.509 module extracts the certificate using a filter. It maps the certificate to an application user and loads that user's set of granted authorities for use with the standard Spring Security infrastructure.

You should be familiar with using certificates and setting up client authentication for your servlet container before attempting to use it with Spring Security. Most of the work is in creating and installing suitable certificates and keys. For example, if you're using Tomcat then read the instructions here <https://tomcat.apache.org/tomcat-6.0-doc/ssl-howto.html>. It's important that you get this working before trying it out with Spring Security

Adding X.509 Authentication to Your Web Application

Enabling X.509 client authentication is very straightforward. Just add the `<x509 />` element to your http security namespace configuration.

```

<http>
  ...
  <x509 subject-principal-regex="CN=(.*)", user-service-ref="userService"/>;
</http>

```

The element has two optional attributes:

- `subject-principal-regex`. The regular expression used to extract a username from the certificate's subject name. The default value is shown above. This is the username which will be passed to the `UserDetailsService` to load the authorities for the user.
- `user-service-ref`. This is the bean Id of the `UserDetailsService` to be used with X.509. It isn't needed if there is only one defined in your application context.

The `subject-principal-regex` should contain a single group. For example the default expression `"CN=(.*?)"` matches the common name field. So if the subject name in the certificate is `"CN=Jimi Hendrix, OU=..."`, this will give a user name of `"Jimi Hendrix"`. The matches are case insensitive. So `"emailAddress=(.*?)"` will match `"EMAILADDRESS=jimi@hendrix.org,CN=..."` giving a user name `"jimi@hendrix.org"`. If the client presents a certificate and a valid username is successfully extracted, then there should be a valid `Authentication` object in the security context. If no certificate is found, or no corresponding user could be found then the security context will remain empty. This means that you can easily use X.509 authentication with other options such as a form-based login.

Setting up SSL in Tomcat

There are some pre-generated certificates in the `samples/certificate` directory in the Spring Security project. You can use these to enable SSL for testing if you don't want to generate your own. The file `server.jks` contains the server certificate, private key and the issuing certificate authority certificate. There are also some client certificate files for the users from the sample applications. You can install these in your browser to enable SSL client authentication.

To run tomcat with SSL support, drop the `server.jks` file into the tomcat `conf` directory and add the following connector to the `server.xml` file

```
<Connector port="8443" protocol="HTTP/1.1" SSLEnabled="true" scheme="https" secure="true"
  clientAuth="true" sslProtocol="TLS"
  keystoreFile="${catalina.home}/conf/server.jks"
  keystoreType="JKS" keystorePass="password"
  truststoreFile="${catalina.home}/conf/server.jks"
  truststoreType="JKS" truststorePass="password"
/>
```

`clientAuth` can also be set to `want` if you still want SSL connections to succeed even if the client doesn't provide a certificate. Clients which don't present a certificate won't be able to access any objects secured by Spring Security unless you use a non-X.509 authentication mechanism, such as form authentication.

10.20 Run-As Authentication Replacement

Overview

The `AbstractSecurityInterceptor` is able to temporarily replace the `Authentication` object in the `SecurityContext` and `SecurityContextHolder` during the secure object callback phase. This only occurs if the original `Authentication` object was successfully processed by the `AuthenticationManager` and `AccessDecisionManager`. The `RunAsManager` will indicate the replacement `Authentication` object, if any, that should be used during the `SecurityInterceptorCallback`.

By temporarily replacing the `Authentication` object during the secure object callback phase, the secured invocation will be able to call other objects which require different authentication and

authorization credentials. It will also be able to perform any internal security checks for specific `GrantedAuthority` objects. Because Spring Security provides a number of helper classes that automatically configure remoting protocols based on the contents of the `SecurityContextHolder`, these run-as replacements are particularly useful when calling remote web services

Configuration

A `RunAsManager` interface is provided by Spring Security:

```
Authentication buildRunAs(Authentication authentication, Object object,
    List<ConfigAttribute> config);

boolean supports(ConfigAttribute attribute);

boolean supports(Class clazz);
```

The first method returns the `Authentication` object that should replace the existing `Authentication` object for the duration of the method invocation. If the method returns null, it indicates no replacement should be made. The second method is used by the `AbstractSecurityInterceptor` as part of its startup validation of configuration attributes. The `supports(Class)` method is called by a security interceptor implementation to ensure the configured `RunAsManager` supports the type of secure object that the security interceptor will present.

One concrete implementation of a `RunAsManager` is provided with Spring Security. The `RunAsManagerImpl` class returns a replacement `RunAsUserToken` if any `ConfigAttribute` starts with `RUN_AS_`. If any such `ConfigAttribute` is found, the replacement `RunAsUserToken` will contain the same principal, credentials and granted authorities as the original `Authentication` object, along with a new `SimpleGrantedAuthority` for each `RUN_AS_ ConfigAttribute`. Each new `SimpleGrantedAuthority` will be prefixed with `ROLE_`, followed by the `RUN_AS ConfigAttribute`. For example, a `RUN_AS_SERVER` will result in the replacement `RunAsUserToken` containing a `ROLE_RUN_AS_SERVER` granted authority.

The replacement `RunAsUserToken` is just like any other `Authentication` object. It needs to be authenticated by the `AuthenticationManager`, probably via delegation to a suitable `AuthenticationProvider`. The `RunAsImplAuthenticationProvider` performs such authentication. It simply accepts as valid any `RunAsUserToken` presented.

To ensure malicious code does not create a `RunAsUserToken` and present it for guaranteed acceptance by the `RunAsImplAuthenticationProvider`, the hash of a key is stored in all generated tokens. The `RunAsManagerImpl` and `RunAsImplAuthenticationProvider` is created in the bean context with the same key:

```
<bean id="runAsManager"
    class="org.springframework.security.access.intercept.RunAsManagerImpl">
<property name="key" value="my_run_as_password"/>
</bean>

<bean id="runAsAuthenticationProvider"
    class="org.springframework.security.access.intercept.RunAsImplAuthenticationProvider">
<property name="key" value="my_run_as_password"/>
</bean>
```

By using the same key, each `RunAsUserToken` can be validated it was created by an approved `RunAsManagerImpl`. The `RunAsUserToken` is immutable after creation for security reasons

10.21 Form Login

Form Login Java Configuration

You might be wondering where the login form came from when you were prompted to log in, since we made no mention of any HTML files or JSPs. Since Spring Security's default configuration does not explicitly set a URL for the login page, Spring Security generates one automatically, based on the features that are enabled and using standard values for the URL which processes the submitted login, the default target URL the user will be sent to after logging in and so on.

While the automatically generated log in page is convenient to get up and running quickly, most applications will want to provide their own login page. When we want to change the default configuration, we can customize the `WebSecurityConfigurerAdapter` that we mentioned earlier by extending it like so:

```
public class WebSecurityConfig extends WebSecurityConfigurerAdapter {
    // ...
}
```

And then override the `configure` method as seen below:

```
protected void configure(HttpSecurity http) throws Exception {
    http
        .authorizeRequests(authorizeRequests ->
            authorizeRequests
                .anyRequest().authenticated()
        )
        .formLogin(formLogin ->
            formLogin
                .loginPage("/login") ❶
                .permitAll()         ❷
        );
}
```

- ❶ The updated configuration specifies the location of the log in page.
- ❷ We must grant all users (i.e. unauthenticated users) access to our log in page. The `formLogin().permitAll()` method allows granting access to all users for all URLs associated with form based log in.

An example log in page implemented with JSPs for our current configuration can be seen below:

Note

The login page below represents our current configuration. We could easily update our configuration if some of the defaults do not meet our needs.

```

<c:url value="/login" var="loginUrl"/>
<form action="{loginUrl}" method="post">
  <c:if test="{param.error != null}">
    <p>
      Invalid username and password.
    </p>
  </c:if>
  <c:if test="{param.logout != null}">
    <p>
      You have been logged out.
    </p>
  </c:if>
  <p>
    <label for="username">Username</label>
    <input type="text" id="username" name="username"/>
  </p>
  <p>
    <label for="password">Password</label>
    <input type="password" id="password" name="password"/>
  </p>
  <input type="hidden"
    name="{_csrf.parameterName}"
    value="{_csrf.token}"/>
  <button type="submit" class="btn">Log in</button>
</form>

```

- ❶ A POST to the /login URL will attempt to authenticate the user
- ❷ If the query parameter `error` exists, authentication was attempted and failed
- ❸ If the query parameter `logout` exists, the user was successfully logged out
- ❹ The username must be present as the HTTP parameter named `username`
- ❺ The password must be present as the HTTP parameter named `password`
- ❻ We must the section called “Include the CSRF Token” To learn more read the the section called “Cross Site Request Forgery (CSRF)” section of the reference

Form Login XML Configuration

Form and Basic Login Options

You might be wondering where the login form came from when you were prompted to log in, since we made no mention of any HTML files or JSPs. In fact, since we didn’t explicitly set a URL for the login page, Spring Security generates one automatically, based on the features that are enabled and using standard values for the URL which processes the submitted login, the default target URL the user will be sent to after logging in and so on. However, the namespace offers plenty of support to allow you to customize these options. For example, if you want to supply your own login page, you could use:

```

<http>
<intercept-url pattern="/login.jsp*" access="IS_AUTHENTICATED_ANONYMOUSLY"/>
<intercept-url pattern="/**" access="ROLE_USER" />
<form-login login-page="/login.jsp"/>
</http>

```

Also note that we’ve added an extra `intercept-url` element to say that any requests for the login page should be available to anonymous users²⁷ and also the [AuthenticatedVoter](#) class for more details on how the value `IS_AUTHENTICATED_ANONYMOUSLY` is processed.]. Otherwise the request would be matched by the pattern `/**` and it wouldn’t be possible to access the login page itself! This is a common configuration error and will result in an infinite loop in the application. Spring Security will emit a warning in the log if your login page appears to be secured. It is also possible to have all requests matching a

²⁷See the chapter on Section 10.15, “Anonymous Authentication”

particular pattern bypass the security filter chain completely, by defining a separate `http` element for the pattern like this:

```
<http pattern="/css/**" security="none"/>
<http pattern="/login.jsp*" security="none"/>

<http use-expressions="false">
<intercept-url pattern="/**" access="ROLE_USER" />
<form-login login-page="/login.jsp"/>
</http>
```

From Spring Security 3.1 it is now possible to use multiple `http` elements to define separate security filter chain configurations for different request patterns. If the `pattern` attribute is omitted from an `http` element, it matches all requests. Creating an unsecured pattern is a simple example of this syntax, where the pattern is mapped to an empty filter chain²⁸. We'll look at this new syntax in more detail in the chapter on the [Security Filter Chain](#).

It's important to realise that these unsecured requests will be completely oblivious to any Spring Security web-related configuration or additional attributes such as `requires-channel`, so you will not be able to access information on the current user or call secured methods during the request. Use `access='IS_AUTHENTICATED_ANONYMOUSLY'` as an alternative if you still want the security filter chain to be applied.

If you want to use basic authentication instead of form login, then change the configuration to

```
<http use-expressions="false">
<intercept-url pattern="/**" access="ROLE_USER" />
<http-basic />
</http>
```

Basic authentication will then take precedence and will be used to prompt for a login when a user attempts to access a protected resource. Form login is still available in this configuration if you wish to use it, for example through a login form embedded in another web page.

10.22 Basic and Digest Authentication

Basic and digest authentication are alternative authentication mechanisms which are popular in web applications. Basic authentication is often used with stateless clients which pass their credentials on each request. It's quite common to use it in combination with form-based authentication where an application is used through both a browser-based user interface and as a web-service. However, basic authentication transmits the password as plain text so it should only really be used over an encrypted transport layer such as HTTPS.

BasicAuthenticationFilter

`BasicAuthenticationFilter` is responsible for processing basic authentication credentials presented in HTTP headers. This can be used for authenticating calls made by Spring remoting protocols (such as Hessian and Burlap), as well as normal browser user agents (such as Firefox and Internet Explorer). The standard governing HTTP Basic Authentication is defined by RFC 1945, Section 11, and `BasicAuthenticationFilter` conforms with this RFC. Basic Authentication is an attractive approach to authentication, because it is very widely deployed in user agents and implementation is extremely simple (it's just a Base64 encoding of the username:password, specified in an HTTP header).

²⁸The use of multiple `<http>` elements is an important feature, allowing the namespace to simultaneously support both stateful and stateless paths within the same application, for example. The previous syntax, using the attribute `filters="none"` on an `intercept-url` element is incompatible with this change and is no longer supported in 3.1.

Configuration

To implement HTTP Basic Authentication, you need to add a `BasicAuthenticationFilter` to your filter chain. The application context should contain `BasicAuthenticationFilter` and its required collaborator:

```
<bean id="basicAuthenticationFilter"
class="org.springframework.security.web.authentication.www.BasicAuthenticationFilter">
<property name="authenticationManager" ref="authenticationManager"/>
<property name="authenticationEntryPoint" ref="authenticationEntryPoint"/>
</bean>

<bean id="authenticationEntryPoint"
class="org.springframework.security.web.authentication.www.BasicAuthenticationEntryPoint">
<property name="realmName" value="Name Of Your Realm"/>
</bean>
```

The configured `AuthenticationManager` processes each authentication request. If authentication fails, the configured `AuthenticationEntryPoint` will be used to retry the authentication process. Usually you will use the filter in combination with a `BasicAuthenticationEntryPoint`, which returns a 401 response with a suitable header to retry HTTP Basic authentication. If authentication is successful, the resulting `Authentication` object will be placed into the `SecurityContextHolder` as usual.

If the authentication event was successful, or authentication was not attempted because the HTTP header did not contain a supported authentication request, the filter chain will continue as normal. The only time the filter chain will be interrupted is if authentication fails and the `AuthenticationEntryPoint` is called.

10.23 DigestAuthenticationFilter

`DigestAuthenticationFilter` is capable of processing digest authentication credentials presented in HTTP headers. Digest Authentication attempts to solve many of the weaknesses of Basic authentication, specifically by ensuring credentials are never sent in clear text across the wire. Many user agents support Digest Authentication, including Mozilla Firefox and Internet Explorer. The standard governing HTTP Digest Authentication is defined by RFC 2617, which updates an earlier version of the Digest Authentication standard prescribed by RFC 2069. Most user agents implement RFC 2617. Spring Security's `DigestAuthenticationFilter` is compatible with the "auth" quality of protection (qop) prescribed by RFC 2617, which also provides backward compatibility with RFC 2069. Digest Authentication is a more attractive option if you need to use unencrypted HTTP (i.e. no TLS/HTTPS) and wish to maximise security of the authentication process. Indeed Digest Authentication is a mandatory requirement for the WebDAV protocol, as noted by RFC 2518 Section 17.1.

Note

You should not use Digest in modern applications because it is not considered secure. The most obvious problem is that you must store your passwords in plaintext, encrypted, or an MD5 format. All of these storage formats are considered insecure. Instead, you should use a one way adaptive password hash (i.e. bcrypt, PBKDF2, SCrypt, etc).

Central to Digest Authentication is a "nonce". This is a value the server generates. Spring Security's nonce adopts the following format:

```
base64(expirationTime + ":" + md5Hex(expirationTime + ":" + key))
expirationTime: The date and time when the nonce expires, expressed in milliseconds
key:           A private key to prevent modification of the nonce token
```

The `DigestAuthenticationEntryPoint` has a property specifying the key used for generating the nonce tokens, along with a `nonceValiditySeconds` property for determining the expiration time (default 300, which equals five minutes). Whist ever the nonce is valid, the digest is computed by concatenating various strings including the username, password, nonce, URI being requested, a client-generated nonce (merely a random value which the user agent generates each request), the realm name etc, then performing an MD5 hash. Both the server and user agent perform this digest computation, resulting in different hash codes if they disagree on an included value (eg password). In Spring Security implementation, if the server-generated nonce has merely expired (but the digest was otherwise valid), the `DigestAuthenticationEntryPoint` will send a `"stale=true"` header. This tells the user agent there is no need to disturb the user (as the password and username etc is correct), but simply to try again using a new nonce.

An appropriate value for the `nonceValiditySeconds` parameter of `DigestAuthenticationEntryPoint` depends on your application. Extremely secure applications should note that an intercepted authentication header can be used to impersonate the principal until the `expirationTime` contained in the nonce is reached. This is the key principle when selecting an appropriate setting, but it would be unusual for immensely secure applications to not be running over TLS/HTTPS in the first instance.

Because of the more complex implementation of Digest Authentication, there are often user agent issues. For example, Internet Explorer fails to present an “opaque” token on subsequent requests in the same session. Spring Security filters therefore encapsulate all state information into the “nonce” token instead. In our testing, Spring Security’s implementation works reliably with Mozilla Firefox and Internet Explorer, correctly handling nonce timeouts etc.

Configuration

Now that we’ve reviewed the theory, let’s see how to use it. To implement HTTP Digest Authentication, it is necessary to define `DigestAuthenticationFilter` in the filter chain. The application context will need to define the `DigestAuthenticationFilter` and its required collaborators:

```
<bean id="digestFilter" class=
    "org.springframework.security.web.authentication.www.DigestAuthenticationFilter">
  <property name="userDetailsService" ref="jdbcDaoImpl"/>
  <property name="authenticationEntryPoint" ref="digestEntryPoint"/>
  <property name="userCache" ref="userCache"/>
</bean>

<bean id="digestEntryPoint" class=
    "org.springframework.security.web.authentication.www.DigestAuthenticationEntryPoint">
  <property name="realmName" value="Contacts Realm via Digest Authentication"/>
  <property name="key" value="acegi"/>
  <property name="nonceValiditySeconds" value="10"/>
</bean>
```

The configured `UserDetailsService` is needed because `DigestAuthenticationFilter` must have direct access to the clear text password of a user. Digest Authentication will NOT work if you are using encoded passwords in your DAO²⁹. The DAO collaborator, along with the `UserCache`, are typically shared directly with a `DaoAuthenticationProvider`. The `authenticationEntryPoint`

²⁹It is possible to encode the password in the format `HEX(MD5(username:realm:password))` provided the `DigestAuthenticationFilter.passwordAlreadyEncoded` is set to `true`. However, other password encodings will not work with digest authentication.

property must be `DigestAuthenticationEntryPoint`, so that `DigestAuthenticationFilter` can obtain the correct `realmName` and `key` for digest calculations.

Like `BasicAuthenticationFilter`, if authentication is successful an `Authentication` request token will be placed into the `SecurityContextHolder`. If the authentication event was successful, or authentication was not attempted because the HTTP header did not contain a Digest Authentication request, the filter chain will continue as normal. The only time the filter chain will be interrupted is if authentication fails and the `AuthenticationEntryPoint` is called, as discussed in the previous paragraph.

Digest Authentication's RFC offers a range of additional features to further increase security. For example, the nonce can be changed on every request. Despite this, Spring Security implementation was designed to minimise the complexity of the implementation (and the doubtless user agent incompatibilities that would emerge), and avoid needing to store server-side state. You are invited to review RFC 2617 if you wish to explore these features in more detail. As far as we are aware, Spring Security's implementation does comply with the minimum standards of this RFC.

10.24 Handling Logouts

Logout Java Configuration

When using the [WebSecurityConfigurerAdapter](#), logout capabilities are automatically applied. The default is that accessing the URL `/logout` will log the user out by:

- Invalidating the HTTP Session
- Cleaning up any RememberMe authentication that was configured
- Clearing the `SecurityContextHolder`
- Redirect to `/login?logout`

Similar to configuring login capabilities, however, you also have various options to further customize your logout requirements:

```
protected void configure(HttpSecurity http) throws Exception {
    http
        .logout(logout ->                                ❶
            logout
                .logoutUrl("/my/logout")                  ❷
                .logoutSuccessUrl("/my/index")           ❸
                .logoutSuccessHandler(logoutSuccessHandler) ❹
                .invalidateHttpSession(true)              ❺
                .addLogoutHandler(logoutHandler)          ❻
                .deleteCookies(cookieNamesToClear)       ❼
            )
        ...
}
```

- ❶ Provides logout support. This is automatically applied when using `WebSecurityConfigurerAdapter`.
- ❷ The URL that triggers log out to occur (default is `/logout`). If CSRF protection is enabled (default), then the request must also be a POST. For more information, please consult the [JavaDoc](#).
- ❸ The URL to redirect to after logout has occurred. The default is `/login?logout`. For more information, please consult the [JavaDoc](#).
- ❹ Let's you specify a custom `LogoutSuccessHandler`. If this is specified, `logoutSuccessUrl()` is ignored. For more information, please consult the [JavaDoc](#).

- ⑤ Specify whether to invalidate the `HttpSession` at the time of logout. This is **true** by default. Configures the `SecurityContextLogoutHandler` under the covers. For more information, please consult the [JavaDoc](#).
- ⑥ Adds a `LogoutHandler`. `SecurityContextLogoutHandler` is added as the last `LogoutHandler` by default.
- ⑦ Allows specifying the names of cookies to be removed on logout success. This is a shortcut for adding a `CookieClearingLogoutHandler` explicitly.

Note

=== Logouts can of course also be configured using the XML Namespace notation. Please see the documentation for the [logout element](#) in the Spring Security XML Namespace section for further details. ===

Generally, in order to customize logout functionality, you can add [LogoutHandler](#) and/or [LogoutSuccessHandler](#) implementations. For many common scenarios, these handlers are applied under the covers when using the fluent API.

Logout XML Configuration

The `logout` element adds support for logging out by navigating to a particular URL. The default logout URL is `/logout`, but you can set it to something else using the `logout-url` attribute. More information on other available attributes may be found in the namespace appendix.

LogoutHandler

Generally, [LogoutHandler](#) implementations indicate classes that are able to participate in logout handling. They are expected to be invoked to perform necessary clean-up. As such they should not throw exceptions. Various implementations are provided:

- [PersistentTokenBasedRememberMeServices](#)
- [TokenBasedRememberMeServices](#)
- [CookieClearingLogoutHandler](#)
- [CsrfLogoutHandler](#)
- [SecurityContextLogoutHandler](#)
- [HeaderWriterLogoutHandler](#)

Please see the section called “Remember-Me Interfaces and Implementations” for details.

Instead of providing `LogoutHandler` implementations directly, the fluent API also provides shortcuts that provide the respective `LogoutHandler` implementations under the covers. E.g. `deleteCookies()` allows specifying the names of one or more cookies to be removed on logout success. This is a shortcut compared to adding a `CookieClearingLogoutHandler`.

LogoutSuccessHandler

The `LogoutSuccessHandler` is called after a successful logout by the `LogoutFilter`, to handle e.g. redirection or forwarding to the appropriate destination. Note that the interface is almost the same as the `LogoutHandler` but may raise an exception.

The following implementations are provided:

- [SimpleUrlLogoutSuccessHandler](#)
- `HttpStatusReturningLogoutSuccessHandler`

As mentioned above, you don't need to specify the `SimpleUrlLogoutSuccessHandler` directly. Instead, the fluent API provides a shortcut by setting the `logoutSuccessUrl()`. This will setup the `SimpleUrlLogoutSuccessHandler` under the covers. The provided URL will be redirected to after a logout has occurred. The default is `/login?logout`.

The `HttpStatusReturningLogoutSuccessHandler` can be interesting in REST API type scenarios. Instead of redirecting to a URL upon the successful logout, this `LogoutSuccessHandler` allows you to provide a plain HTTP status code to be returned. If not configured a status code 200 will be returned by default.

Further Logout-Related References

- [Logout Handling](#)
- [Testing Logout](#)
- [HttpServletRequest.logout\(\)](#)
- the section called "Remember-Me Interfaces and Implementations"
- [Logging Out](#) in section CSRF Caveats
- Section [Single Logout](#) (CAS protocol)
- Documentation for the [logout element](#) in the Spring Security XML Namespace section

10.25 Setting a Custom AuthenticationEntryPoint

If you aren't using form login, OpenID or basic authentication through the namespace, you may want to define an authentication filter and entry point using a traditional bean syntax and link them into the namespace, as we've just seen. The corresponding `AuthenticationEntryPoint` can be set using the `entry-point-ref` attribute on the `<http>` element.

The CAS sample application is a good example of the use of custom beans with the namespace, including this syntax. If you aren't familiar with authentication entry points, they are discussed in the [technical overview](#) chapter.

11. Authorization

The advanced authorization capabilities within Spring Security represent one of the most compelling reasons for its popularity. Irrespective of how you choose to authenticate - whether using a Spring Security-provided mechanism and provider, or integrating with a container or other non-Spring Security authentication authority - you will find the authorization services can be used within your application in a consistent and simple way.

In this part we'll explore the different `AbstractSecurityInterceptor` implementations, which were introduced in Part I. We then move on to explore how to fine-tune authorization through use of domain access control lists.

11.1 Authorization Architecture

Authorities

As we saw in the [technical overview](#), all `Authentication` implementations store a list of `GrantedAuthority` objects. These represent the authorities that have been granted to the principal. The `GrantedAuthority` objects are inserted into the `Authentication` object by the `AuthenticationManager` and are later read by `AccessDecisionManager`s when making authorization decisions.

`GrantedAuthority` is an interface with only one method:

```
String getAuthority();
```

This method allows `AccessDecisionManager`s to obtain a precise `String` representation of the `GrantedAuthority`. By returning a representation as a `String`, a `GrantedAuthority` can be easily "read" by most `AccessDecisionManager`s. If a `GrantedAuthority` cannot be precisely represented as a `String`, the `GrantedAuthority` is considered "complex" and `getAuthority()` must return `null`.

An example of a "complex" `GrantedAuthority` would be an implementation that stores a list of operations and authority thresholds that apply to different customer account numbers. Representing this complex `GrantedAuthority` as a `String` would be quite difficult, and as a result the `getAuthority()` method should return `null`. This will indicate to any `AccessDecisionManager` that it will need to specifically support the `GrantedAuthority` implementation in order to understand its contents.

Spring Security includes one concrete `GrantedAuthority` implementation, `SimpleGrantedAuthority`. This allows any user-specified `String` to be converted into a `GrantedAuthority`. All `AuthenticationProvider`s included with the security architecture use `SimpleGrantedAuthority` to populate the `Authentication` object.

Pre-Invocation Handling

As we've also seen in the [Technical Overview](#) chapter, Spring Security provides interceptors which control access to secure objects such as method invocations or web requests. A pre-invocation decision on whether the invocation is allowed to proceed is made by the `AccessDecisionManager`.

The AccessDecisionManager

The `AccessDecisionManager` is called by the `AbstractSecurityInterceptor` and is responsible for making final access control decisions. The `AccessDecisionManager` interface contains three methods:

```
void decide(Authentication authentication, Object secureObject,
            Collection<ConfigAttribute> attrs) throws AccessDeniedException;

boolean supports(ConfigAttribute attribute);

boolean supports(Class clazz);
```

The `AccessDecisionManager`'s `decide` method is passed all the relevant information it needs in order to make an authorization decision. In particular, passing the secure `Object` enables those arguments contained in the actual secure object invocation to be inspected. For example, let's assume the secure object was a `MethodInvocation`. It would be easy to query the `MethodInvocation` for any `Customer` argument, and then implement some sort of security logic in the `AccessDecisionManager` to ensure the principal is permitted to operate on that customer. Implementations are expected to throw an `AccessDeniedException` if access is denied.

The `supports(ConfigAttribute)` method is called by the `AbstractSecurityInterceptor` at startup time to determine if the `AccessDecisionManager` can process the passed `ConfigAttribute`. The `supports(Class)` method is called by a security interceptor implementation to ensure the configured `AccessDecisionManager` supports the type of secure object that the security interceptor will present.

Voting-Based AccessDecisionManager Implementations

Whilst users can implement their own `AccessDecisionManager` to control all aspects of authorization, Spring Security includes several `AccessDecisionManager` implementations that are based on voting. Figure 11.1, "Voting Decision Manager" illustrates the relevant classes.




Figure 11.1. Voting Decision Manager

Using this approach, a series of `AccessDecisionVoter` implementations are polled on an authorization decision. The `AccessDecisionManager` then decides whether or not to throw an `AccessDeniedException` based on its assessment of the votes.

The `AccessDecisionVoter` interface has three methods:

```

int vote(Authentication authentication, Object object, Collection<ConfigAttribute> attrs);

boolean supports(ConfigAttribute attribute);

boolean supports(Class clazz);

```

Concrete implementations return an `int`, with possible values being reflected in the `AccessDecisionVoter` static fields `ACCESS_ABSTAIN`, `ACCESS_DENIED` and `ACCESS_GRANTED`. A voting implementation will return `ACCESS_ABSTAIN` if it has no opinion on an authorization decision. If it does have an opinion, it must return either `ACCESS_DENIED` or `ACCESS_GRANTED`.

There are three concrete `AccessDecisionManager`s provided with Spring Security that tally the votes. The `ConsensusBased` implementation will grant or deny access based on the consensus of non-abstain votes. Properties are provided to control behavior in the event of an equality of votes or if all votes are abstain. The `AffirmativeBased` implementation will grant access if one or more `ACCESS_GRANTED` votes were received (i.e. a deny vote will be ignored, provided there was at least one grant vote). Like the `ConsensusBased` implementation, there is a parameter that controls the behavior if all voters abstain. The `UnanimousBased` provider expects unanimous `ACCESS_GRANTED` votes in order to grant access, ignoring abstains. It will deny access if there is any `ACCESS_DENIED` vote. Like the other implementations, there is a parameter that controls the behaviour if all voters abstain.

It is possible to implement a custom `AccessDecisionManager` that tallies votes differently. For example, votes from a particular `AccessDecisionVoter` might receive additional weighting, whilst a deny vote from a particular voter may have a veto effect.

RoleVoter

The most commonly used `AccessDecisionVoter` provided with Spring Security is the simple `RoleVoter`, which treats configuration attributes as simple role names and votes to grant access if the user has been assigned that role.

It will vote if any `ConfigAttribute` begins with the prefix `ROLE_`. It will vote to grant access if there is a `GrantedAuthority` which returns a `String` representation (via the `getAuthority()` method) exactly equal to one or more `ConfigAttributes` starting with the prefix `ROLE_`. If there is no exact match of any `ConfigAttribute` starting with `ROLE_`, the `RoleVoter` will vote to deny access. If no `ConfigAttribute` begins with `ROLE_`, the voter will abstain.

AuthenticatedVoter

Another voter which we've implicitly seen is the `AuthenticatedVoter`, which can be used to differentiate between anonymous, fully-authenticated and remember-me authenticated users. Many sites allow certain limited access under remember-me authentication, but require a user to confirm their identity by logging in for full access.

When we've used the attribute `IS_AUTHENTICATED_ANONYMOUSLY` to grant anonymous access, this attribute was being processed by the `AuthenticatedVoter`. See the Javadoc for this class for more information.

Custom Voters

Obviously, you can also implement a custom `AccessDecisionVoter` and you can put just about any access-control logic you want in it. It might be specific to your application (business-logic related) or it might implement some security administration logic. For example, you'll find a [blog article](#) on the Spring web site which describes how to use a voter to deny access in real-time to users whose accounts have been suspended.

After Invocation Handling

Whilst the `AccessDecisionManager` is called by the `AbstractSecurityInterceptor` before proceeding with the secure object invocation, some applications need a way of modifying the object actually returned by the secure object invocation. Whilst you could easily implement your own AOP concern to achieve this, Spring Security provides a convenient hook that has several concrete implementations that integrate with its ACL capabilities.

Figure 11.2, "After Invocation Implementation" illustrates Spring Security's `AfterInvocationManager` and its concrete implementations.




Figure 11.2. After Invocation Implementation

Like many other parts of Spring Security, `AfterInvocationManager` has a single concrete implementation, `AfterInvocationProviderManager`, which polls a list of `AfterInvocationProvider`s. Each `AfterInvocationProvider` is allowed to modify the return object or throw an `AccessDeniedException`. Indeed multiple providers can modify the object, as the result of the previous provider is passed to the next in the list.

Please be aware that if you're using `AfterInvocationManager`, you will still need configuration attributes that allow the `MethodSecurityInterceptor`'s `AccessDecisionManager` to allow an operation. If you're using the typical Spring Security included `AccessDecisionManager` implementations, having no configuration attributes defined for a particular secure method invocation will cause each `AccessDecisionVoter` to abstain from voting. In turn, if the `AccessDecisionManager` property "allowIfAllAbstainDecisions" is false, an `AccessDeniedException` will be thrown. You may avoid this potential issue by either (i) setting "allowIfAllAbstainDecisions" to true (although this is generally not recommended) or (ii) simply ensure that there is at least one configuration attribute that an `AccessDecisionVoter` will vote to grant access for. This latter (recommended) approach is usually achieved through a `ROLE_USER` or `ROLE_AUTHENTICATED` configuration attribute.

Hierarchical Roles

It is a common requirement that a particular role in an application should automatically "include" other roles. For example, in an application which has the concept of an "admin" and a "user" role, you may want an admin to be able to do everything a normal user can. To achieve this, you can either make sure that all admin users are also assigned the "user" role. Alternatively, you can modify every access constraint which requires the "user" role to also include the "admin" role. This can get quite complicated if you have a lot of different roles in your application.

The use of a role-hierarchy allows you to configure which roles (or authorities) should include others. An extended version of Spring Security's [RoleVoter](#), `RoleHierarchyVoter`, is configured with a `RoleHierarchy`, from which it obtains all the "reachable authorities" which the user is assigned. A typical configuration might look like this:

```

<bean id="roleVoter" class="org.springframework.security.access.vote.RoleHierarchyVoter">
  <constructor-arg ref="roleHierarchy" />
</bean>
<bean id="roleHierarchy"
  class="org.springframework.security.access.hierarchicalroles.RoleHierarchyImpl">
  <property name="hierarchy">
    <value>
      ROLE_ADMIN > ROLE_STAFF
      ROLE_STAFF > ROLE_USER
      ROLE_USER > ROLE_GUEST
    </value>
  </property>
</bean>

```

Here we have four roles in a hierarchy `ROLE_ADMIN # ROLE_STAFF # ROLE_USER # ROLE_GUEST`. A user who is authenticated with `ROLE_ADMIN`, will behave as if they have all four roles when security constraints are evaluated against an `AccessDecisionManager` configured with the above `RoleHierarchyVoter`. The `>` symbol can be thought of as meaning "includes".

Role hierarchies offer a convenient means of simplifying the access-control configuration data for your application and/or reducing the number of authorities which you need to assign to a user. For more complex requirements you may wish to define a logical mapping between the specific access-rights your application requires and the roles that are assigned to users, translating between the two when loading the user information.

11.2 Secure Object Implementations

AOP Alliance (MethodInvocation) Security Interceptor

Prior to Spring Security 2.0, securing `MethodInvocation`s needed quite a lot of boiler plate configuration. Now the recommended approach for method security is to use [namespace configuration](#). This way the method security infrastructure beans are configured automatically for you so you don't really need to know about the implementation classes. We'll just provide a quick overview of the classes that are involved here.

Method security is enforced using a `MethodSecurityInterceptor`, which secures `MethodInvocation`s. Depending on the configuration approach, an interceptor may be specific to a single bean or shared between multiple beans. The interceptor uses a `MethodSecurityMetadataSource` instance to obtain the configuration attributes that apply to a particular method invocation. `MapBasedMethodSecurityMetadataSource` is used to store configuration attributes keyed by method names (which can be wildcarded) and will be used internally when the attributes are defined in the application context using the `<intercept-methods>` or `<protect-point>` elements. Other implementations will be used to handle annotation-based configuration.

Explicit MethodSecurityInterceptor Configuration

You can of course configure a `MethodSecurityInterceptor` directly in your application context for use with one of Spring AOP's proxying mechanisms:

```

<bean id="bankManagerSecurity" class=
    "org.springframework.security.access.intercept.aopalliance.MethodSecurityInterceptor">
  <property name="authenticationManager" ref="authenticationManager"/>
  <property name="accessDecisionManager" ref="accessDecisionManager"/>
  <property name="afterInvocationManager" ref="afterInvocationManager"/>
  <property name="securityMetadataSource">
    <sec:method-security-metadata-source>
      <sec:protect method="com.mycompany.BankManager.delete*" access="ROLE_SUPERVISOR"/>
      <sec:protect method="com.mycompany.BankManager.getBalance" access="ROLE_TELLER,ROLE_SUPERVISOR"/>
    </sec:method-security-metadata-source>
  </property>
</bean>

```

AspectJ (JoinPoint) Security Interceptor

The AspectJ security interceptor is very similar to the AOP Alliance security interceptor discussed in the previous section. Indeed we will only discuss the differences in this section.

The AspectJ interceptor is named `AspectJSecurityInterceptor`. Unlike the AOP Alliance security interceptor, which relies on the Spring application context to weave in the security interceptor via proxying, the `AspectJSecurityInterceptor` is weaved in via the AspectJ compiler. It would not be uncommon to use both types of security interceptors in the same application, with `AspectJSecurityInterceptor` being used for domain object instance security and the AOP Alliance `MethodSecurityInterceptor` being used for services layer security.

Let's first consider how the `AspectJSecurityInterceptor` is configured in the Spring application context:

```

<bean id="bankManagerSecurity" class=
    "org.springframework.security.access.intercept.aspectj.AspectJMethodSecurityInterceptor">
  <property name="authenticationManager" ref="authenticationManager"/>
  <property name="accessDecisionManager" ref="accessDecisionManager"/>
  <property name="afterInvocationManager" ref="afterInvocationManager"/>
  <property name="securityMetadataSource">
    <sec:method-security-metadata-source>
      <sec:protect method="com.mycompany.BankManager.delete*" access="ROLE_SUPERVISOR"/>
      <sec:protect method="com.mycompany.BankManager.getBalance" access="ROLE_TELLER,ROLE_SUPERVISOR"/>
    </sec:method-security-metadata-source>
  </property>
</bean>

```

As you can see, aside from the class name, the `AspectJSecurityInterceptor` is exactly the same as the AOP Alliance security interceptor. Indeed the two interceptors can share the same `securityMetadataSource`, as the `SecurityMetadataSource` works with `java.lang.reflect.Method`s rather than an AOP library-specific class. Of course, your access decisions have access to the relevant AOP library-specific invocation (ie `MethodInvocation` or `JoinPoint`) and as such can consider a range of addition criteria when making access decisions (such as method arguments).

Next you'll need to define an AspectJ aspect. For example:

```

package org.springframework.security.samples.aspectj;

import org.springframework.security.access.intercept.aspectj.AspectJSecurityInterceptor;
import org.springframework.security.access.intercept.aspectj.AspectJCallback;
import org.springframework.beans.factory.InitializingBean;

public aspect DomainObjectInstanceSecurityAspect implements InitializingBean {

    private AspectJSecurityInterceptor securityInterceptor;

    pointcut domainObjectInstanceExecution(): target(PersistableEntity)
        && execution(public * *(..)) && !within(DomainObjectInstanceSecurityAspect);

    Object around(): domainObjectInstanceExecution() {
        if (this.securityInterceptor == null) {
            return proceed();
        }

        AspectJCallback callback = new AspectJCallback() {
            public Object proceedWithObject() {
                return proceed();
            }
        };

        return this.securityInterceptor.invoke(thisJoinPoint, callback);
    }

    public AspectJSecurityInterceptor getSecurityInterceptor() {
        return securityInterceptor;
    }

    public void setSecurityInterceptor(AspectJSecurityInterceptor securityInterceptor) {
        this.securityInterceptor = securityInterceptor;
    }

    public void afterPropertiesSet() throws Exception {
        if (this.securityInterceptor == null)
            throw new IllegalArgumentException("securityInterceptor required");
    }
}

```

In the above example, the security interceptor will be applied to every instance of `PersistableEntity`, which is an abstract class not shown (you can use any other class or pointcut expression you like). For those curious, `AspectJCallback` is needed because the `proceed();` statement has special meaning only within an `around()` body. The `AspectJSecurityInterceptor` calls this anonymous `AspectJCallback` class when it wants the target object to continue.

You will need to configure Spring to load the aspect and wire it with the `AspectJSecurityInterceptor`. A bean declaration which achieves this is shown below:

```

<bean id="domainObjectInstanceSecurityAspect"
    class="security.samples.aspectj.DomainObjectInstanceSecurityAspect"
    factory-method="aspectOf">
<property name="securityInterceptor" ref="bankManagerSecurity"/>
</bean>

```

That's it! Now you can create your beans from anywhere within your application, using whatever means you think fit (eg `new Person();`) and they will have the security interceptor applied.

11.3 Expression-Based Access Control

Spring Security 3.0 introduced the ability to use Spring EL expressions as an authorization mechanism in addition to the simple use of configuration attributes and access-decision voters which have seen before.

Expression-based access control is built on the same architecture but allows complicated Boolean logic to be encapsulated in a single expression.

Overview

Spring Security uses Spring EL for expression support and you should look at how that works if you are interested in understanding the topic in more depth. Expressions are evaluated with a "root object" as part of the evaluation context. Spring Security uses specific classes for web and method security as the root object, in order to provide built-in expressions and access to values such as the current principal.

Common Built-In Expressions

The base class for expression root objects is `SecurityExpressionRoot`. This provides some common expressions which are available in both web and method security.

Table 11.1. Common built-in expressions

Expression	Description
<code>hasRole([role])</code>	Returns <code>true</code> if the current principal has the specified role. By default if the supplied role does not start with 'ROLE_' it will be added. This can be customized by modifying the <code>defaultRolePrefix</code> on <code>DefaultWebSecurityExpressionHandler</code> .
<code>hasAnyRole([role1,role2])</code>	Returns <code>true</code> if the current principal has any of the supplied roles (given as a comma-separated list of strings). By default if the supplied role does not start with 'ROLE_' it will be added. This can be customized by modifying the <code>defaultRolePrefix</code> on <code>DefaultWebSecurityExpressionHandler</code> .
<code>hasAuthority([authority])</code>	Returns <code>true</code> if the current principal has the specified authority.
<code>hasAnyAuthority([authority1,authority2])</code>	Returns <code>true</code> if the current principal has any of the supplied authorities (given as a comma-separated list of strings)
<code>principal</code>	Allows direct access to the principal object representing the current user
<code>authentication</code>	Allows direct access to the current <code>Authentication</code> object obtained from the <code>SecurityContext</code>
<code>permitAll</code>	Always evaluates to <code>true</code>
<code>denyAll</code>	Always evaluates to <code>false</code>
<code>isAnonymous()</code>	Returns <code>true</code> if the current principal is an anonymous user

Expression	Description
<code>isRememberMe()</code>	Returns <code>true</code> if the current principal is a remember-me user
<code>isAuthenticated()</code>	Returns <code>true</code> if the user is not anonymous
<code>isFullyAuthenticated()</code>	Returns <code>true</code> if the user is not an anonymous or a remember-me user
<code>hasPermission(Object target, Object permission)</code>	Returns <code>true</code> if the user has access to the provided target for the given permission. For example, <code>hasPermission(domainObject, 'read')</code>
<code>hasPermission(Object targetId, String targetType, Object permission)</code>	Returns <code>true</code> if the user has access to the provided target for the given permission. For example, <code>hasPermission(1, 'com.example.domain.Message', 'read')</code>

Web Security Expressions

To use expressions to secure individual URLs, you would first need to set the `use-expressions` attribute in the `<http>` element to `true`. Spring Security will then expect the `access` attributes of the `<intercept-url>` elements to contain Spring EL expressions. The expressions should evaluate to a Boolean, defining whether access should be allowed or not. For example:

```
<http>
  <intercept-url pattern="/admin*"
    access="hasRole('admin') and hasIpAddress('192.168.1.0/24')"/>
  ...
</http>
```

Here we have defined that the "admin" area of an application (defined by the URL pattern) should only be available to users who have the granted authority "admin" and whose IP address matches a local subnet. We've already seen the built-in `hasRole` expression in the previous section. The expression `hasIpAddress` is an additional built-in expression which is specific to web security. It is defined by the `WebSecurityExpressionRoot` class, an instance of which is used as the expression root object when evaluating web-access expressions. This object also directly exposed the `HttpServletRequest` object under the name `request` so you can invoke the request directly in an expression. If expressions are being used, a `WebExpressionVoter` will be added to the `AccessDecisionManager` which is used by the namespace. So if you aren't using the namespace and want to use expressions, you will have to add one of these to your configuration.

Referring to Beans in Web Security Expressions

If you wish to extend the expressions that are available, you can easily refer to any Spring Bean you expose. For example, assuming you have a Bean with the name of `webSecurity` that contains the following method signature:

```
public class WebSecurity {
    public boolean check(Authentication authentication, HttpServletRequest request) {
        ...
    }
}
```

You could refer to the method using:

```
<http>
  <intercept-url pattern="/user/**"
    access="@webSecurity.check(authentication,request)"/>
  ...
</http>
```

or in Java configuration

```
http
    .authorizeRequests()
      .antMatchers("/user/**").access("@webSecurity.check(authentication,request)")
      ...
```

Path Variables in Web Security Expressions

At times it is nice to be able to refer to path variables within a URL. For example, consider a RESTful application that looks up a user by id from the URL path in the format `/user/{userId}`.

You can easily refer to the path variable by placing it in the pattern. For example, if you had a Bean with the name of `webSecurity` that contains the following method signature:

```
public class WebSecurity {
    public boolean checkUserId(Authentication authentication, int id) {
        ...
    }
}
```

You could refer to the method using:

```
<http>
  <intercept-url pattern="/user/{userId}/**"
    access="@webSecurity.checkUserId(authentication,#userId)"/>
  ...
</http>
```

or in Java configuration

```
http
    .authorizeRequests(authorizeRequests ->
      authorizeRequests
        .antMatchers("/user/{userId}/**").access("@webSecurity.checkUserId(authentication,#userId)")
        ...
    );
```

In both configurations URLs that match would pass in the path variable (and convert it) into `checkUserId` method. For example, if the URL were `/user/123/resource`, then the id passed in would be 123.

Method Security Expressions

Method security is a bit more complicated than a simple allow or deny rule. Spring Security 3.0 introduced some new annotations in order to allow comprehensive support for the use of expressions.

@Pre and @Post Annotations

There are four annotations which support expression attributes to allow pre and post-invocation authorization checks and also to support filtering of submitted collection arguments or return values. They are `@PreAuthorize`, `@PreFilter`, `@PostAuthorize` and `@PostFilter`. Their use is enabled through the `global-method-security` namespace element:

```
<global-method-security pre-post-annotations="enabled"/>
```

Access Control using @PreAuthorize and @PostAuthorize

The most obviously useful annotation is @PreAuthorize which decides whether a method can actually be invoked or not. For example (from the "Contacts" sample application)

```
@PreAuthorize("hasRole('USER')")
public void create(Contact contact);
```

which means that access will only be allowed for users with the role "ROLE_USER". Obviously the same thing could easily be achieved using a traditional configuration and a simple configuration attribute for the required role. But what about:

```
@PreAuthorize("hasPermission(#contact, 'admin')")
public void deletePermission(Contact contact, Sid recipient, Permission permission);
```

Here we're actually using a method argument as part of the expression to decide whether the current user has the "admin" permission for the given contact. The built-in hasPermission() expression is linked into the Spring Security ACL module through the application context, as we'll [see below](#). You can access any of the method arguments by name as expression variables.

There are a number of ways in which Spring Security can resolve the method arguments. Spring Security uses DefaultSecurityParameterNameDiscoverer to discover the parameter names. By default, the following options are tried for a method as a whole.

- If Spring Security's @P annotation is present on a single argument to the method, the value will be used. This is useful for interfaces compiled with a JDK prior to JDK 8 which do not contain any information about the parameter names. For example:

```
import org.springframework.security.access.method.P;
...
@PreAuthorize("#c.name == authentication.name")
public void doSomething(@P("c") Contact contact);
```

Behind the scenes this use implemented using AnnotationParameterNameDiscoverer which can be customized to support the value attribute of any specified annotation.

- If Spring Data's @Param annotation is present on at least one parameter for the method, the value will be used. This is useful for interfaces compiled with a JDK prior to JDK 8 which do not contain any information about the parameter names. For example:

```
import org.springframework.data.repository.query.Param;
...
@PreAuthorize("#n == authentication.name")
Contact findContactByName(@Param("n") String name);
```

Behind the scenes this use implemented using AnnotationParameterNameDiscoverer which can be customized to support the value attribute of any specified annotation.

- If JDK 8 was used to compile the source with the -parameters argument and Spring 4+ is being used, then the standard JDK reflection API is used to discover the parameter names. This works on both classes and interfaces.

- Last, if the code was compiled with the debug symbols, the parameter names will be discovered using the debug symbols. This will not work for interfaces since they do not have debug information about the parameter names. For interfaces, annotations or the JDK 8 approach must be used.

Any Spring-EL functionality is available within the expression, so you can also access properties on the arguments. For example, if you wanted a particular method to only allow access to a user whose username matched that of the contact, you could write

```
@PreAuthorize("#contact.name == authentication.name")
public void doSomething(Contact contact);
```

Here we are accessing another built-in expression, `authentication`, which is the `Authentication` stored in the security context. You can also access its "principal" property directly, using the expression `principal`. The value will often be a `UserDetails` instance, so you might use an expression like `principal.username` or `principal.enabled`.

Less commonly, you may wish to perform an access-control check after the method has been invoked. This can be achieved using the `@PostAuthorize` annotation. To access the return value from a method, use the built-in name `returnObject` in the expression.

Filtering using `@PreFilter` and `@PostFilter`

As you may already be aware, Spring Security supports filtering of collections and arrays and this can now be achieved using expressions. This is most commonly performed on the return value of a method. For example:

```
@PreAuthorize("hasRole('USER')")
@PostFilter("hasPermission(filterObject, 'read') or hasPermission(filterObject, 'admin')")
public List<Contact> getAll();
```

When using the `@PostFilter` annotation, Spring Security iterates through the returned collection and removes any elements for which the supplied expression is false. The name `filterObject` refers to the current object in the collection. You can also filter before the method call, using `@PreFilter`, though this is a less common requirement. The syntax is just the same, but if there is more than one argument which is a collection type then you have to select one by name using the `filterTarget` property of this annotation.

Note that filtering is obviously not a substitute for tuning your data retrieval queries. If you are filtering large collections and removing many of the entries then this is likely to be inefficient.

Built-In Expressions

There are some built-in expressions which are specific to method security, which we have already seen in use above. The `filterTarget` and `returnValue` values are simple enough, but the use of the `hasPermission()` expression warrants a closer look.

The `PermissionEvaluator` interface

`hasPermission()` expressions are delegated to an instance of `PermissionEvaluator`. It is intended to bridge between the expression system and Spring Security's ACL system, allowing you to specify authorization constraints on domain objects, based on abstract permissions. It has no explicit dependencies on the ACL module, so you could swap that out for an alternative implementation if required. The interface has two methods:

```

boolean hasPermission(Authentication authentication, Object targetDomainObject,
                      Object permission);

boolean hasPermission(Authentication authentication, Serializable targetId,
                      String targetType, Object permission);

```

which map directly to the available versions of the expression, with the exception that the first argument (the `Authentication` object) is not supplied. The first is used in situations where the domain object, to which access is being controlled, is already loaded. Then expression will return true if the current user has the given permission for that object. The second version is used in cases where the object is not loaded, but its identifier is known. An abstract "type" specifier for the domain object is also required, allowing the correct ACL permissions to be loaded. This has traditionally been the Java class of the object, but does not have to be as long as it is consistent with how the permissions are loaded.

To use `hasPermission()` expressions, you have to explicitly configure a `PermissionEvaluator` in your application context. This would look something like this:

```

<security:global-method-security pre-post-annotations="enabled">
<security:expression-handler ref="expressionHandler"/>
</security:global-method-security>

<bean id="expressionHandler" class=
"org.springframework.security.access.expression.method.DefaultMethodSecurityExpressionHandler">
  <property name="permissionEvaluator" ref="myPermissionEvaluator"/>
</bean>

```

Where `myPermissionEvaluator` is the bean which implements `PermissionEvaluator`. Usually this will be the implementation from the ACL module which is called `AclPermissionEvaluator`. See the "Contacts" sample application configuration for more details.

Method Security Meta Annotations

You can make use of meta annotations for method security to make your code more readable. This is especially convenient if you find that you are repeating the same complex expression throughout your code base. For example, consider the following:

```
@PreAuthorize("#contact.name == authentication.name")
```

Instead of repeating this everywhere, we can create a meta annotation that can be used instead.

```

@Retention(RetentionPolicy.RUNTIME)
@PreAuthorize("#contact.name == authentication.name")
public @interface ContactPermission {}

```

Meta annotations can be used for any of the Spring Security method security annotations. In order to remain compliant with the specification JSR-250 annotations do not support meta annotations.

11.4 Authorize Requests

Our examples have only required users to be authenticated and have done so for every URL in our application. We can specify custom requirements for our URLs by adding multiple children to our `http.authorizeRequests()` method. For example:

```

protected void configure(HttpSecurity http) throws Exception {
    http
        .authorizeRequests(authorizeRequests ->                                ❶
            authorizeRequests
                .antMatchers("/resources/**", "/signup", "/about").permitAll()    ❷
                .antMatchers("/admin/**").hasRole("ADMIN")                       ❸
                .antMatchers("/db/**").access("hasRole('ADMIN') and hasRole('DBA')") ❹
                .anyRequest().authenticated()                                     ❺
            )
        .formLogin(withDefaults());
}

```

- ❶ There are multiple children to the `http.authorizeRequests()` method each matcher is considered in the order they were declared.
- ❷ We specified multiple URL patterns that any user can access. Specifically, any user can access a request if the URL starts with `/resources/`, equals `/signup`, or equals `/about`.
- ❸ Any URL that starts with `/admin/` will be restricted to users who have the role `"ROLE_ADMIN"`. You will notice that since we are invoking the `hasRole` method we do not need to specify the `"ROLE_"` prefix.
- ❹ Any URL that starts with `/db/` requires the user to have both `"ROLE_ADMIN"` and `"ROLE_DBA"`. You will notice that since we are using the `hasRole` expression we do not need to specify the `"ROLE_"` prefix.
- ❺ Any URL that has not already been matched on only requires that the user be authenticated

11.5 Method Security

From version 2.0 onwards Spring Security has improved support substantially for adding security to your service layer methods. It provides support for JSR-250 annotation security as well as the framework's original `@Secured` annotation. From 3.0 you can also make use of new [expression-based annotations](#). You can apply security to a single bean, using the `intercept-methods` element to decorate the bean declaration, or you can secure multiple beans across the entire service layer using the AspectJ style pointcuts.

EnableGlobalMethodSecurity

We can enable annotation-based security using the `@EnableGlobalMethodSecurity` annotation on any `@Configuration` instance. For example, the following would enable Spring Security's `@Secured` annotation.

```

@EnableGlobalMethodSecurity(securedEnabled = true)
public class MethodSecurityConfig {
    // ...
}

```

Adding an annotation to a method (on a class or interface) would then limit the access to that method accordingly. Spring Security's native annotation support defines a set of attributes for the method. These will be passed to the `AccessDecisionManager` for it to make the actual decision:

```
public interface BankService {

    @Secured("IS_AUTHENTICATED_ANONYMOUSLY")
    public Account readAccount(Long id);

    @Secured("IS_AUTHENTICATED_ANONYMOUSLY")
    public Account[] findAccounts();

    @Secured("ROLE_TELLER")
    public Account post(Account account, double amount);
}
```

Support for JSR-250 annotations can be enabled using

```
@EnableGlobalMethodSecurity(jsr250Enabled = true)
public class MethodSecurityConfig {
    // ...
}
```

These are standards-based and allow simple role-based constraints to be applied but do not have the power Spring Security's native annotations. To use the new expression-based syntax, you would use

```
@EnableGlobalMethodSecurity(prePostEnabled = true)
public class MethodSecurityConfig {
    // ...
}
```

and the equivalent Java code would be

```
public interface BankService {

    @PreAuthorize("isAnonymous()")
    public Account readAccount(Long id);

    @PreAuthorize("isAnonymous()")
    public Account[] findAccounts();

    @PreAuthorize("hasAuthority('ROLE_TELLER')")
    public Account post(Account account, double amount);
}
```

GlobalMethodSecurityConfiguration

Sometimes you may need to perform operations that are more complicated than are possible with the `@EnableGlobalMethodSecurity` annotation allow. For these instances, you can extend the `GlobalMethodSecurityConfiguration` ensuring that the `@EnableGlobalMethodSecurity` annotation is present on your subclass. For example, if you wanted to provide a custom `MethodSecurityExpressionHandler`, you could use the following configuration:

```
@EnableGlobalMethodSecurity(prePostEnabled = true)
public class MethodSecurityConfig extends GlobalMethodSecurityConfiguration {
    @Override
    protected MethodSecurityExpressionHandler createExpressionHandler() {
        // ... create and return custom MethodSecurityExpressionHandler ...
        return expressionHandler;
    }
}
```

For additional information about methods that can be overridden, refer to the `GlobalMethodSecurityConfiguration` Javadoc.

The <global-method-security> Element

This element is used to enable annotation-based security in your application (by setting the appropriate attributes on the element), and also to group together security pointcut declarations which will be applied across your entire application context. You should only declare one <global-method-security> element. The following declaration would enable support for Spring Security's @Secured:

```
<global-method-security secured-annotations="enabled" />
```

Adding an annotation to a method (on a class or interface) would then limit the access to that method accordingly. Spring Security's native annotation support defines a set of attributes for the method. These will be passed to the `AccessDecisionManager` for it to make the actual decision:

```
public interface BankService {

    @Secured("IS_AUTHENTICATED_ANONYMOUSLY")
    public Account readAccount(Long id);

    @Secured("IS_AUTHENTICATED_ANONYMOUSLY")
    public Account[] findAccounts();

    @Secured("ROLE_TELLER")
    public Account post(Account account, double amount);
}
```

Support for JSR-250 annotations can be enabled using

```
<global-method-security jsr250-annotations="enabled" />
```

These are standards-based and allow simple role-based constraints to be applied but do not have the power Spring Security's native annotations. To use the new expression-based syntax, you would use

```
<global-method-security pre-post-annotations="enabled" />
```

and the equivalent Java code would be

```
public interface BankService {

    @PreAuthorize("isAnonymous()")
    public Account readAccount(Long id);

    @PreAuthorize("isAnonymous()")
    public Account[] findAccounts();

    @PreAuthorize("hasAuthority('ROLE_TELLER')")
    public Account post(Account account, double amount);
}
```

Expression-based annotations are a good choice if you need to define simple rules that go beyond checking the role names against the user's list of authorities.

Note

=== The annotated methods will only be secured for instances which are defined as Spring beans (in the same application context in which method-security is enabled). If you want to secure instances which are not created by Spring (using the `new` operator, for example) then you need to use AspectJ. ===

Note

=== You can enable more than one type of annotation in the same application, but only one type should be used for any interface or class as the behaviour will not be well-defined otherwise. If two annotations are found which apply to a particular method, then only one of them will be applied. ===

Adding Security Pointcuts using protect-pointcut

The use of `protect-pointcut` is particularly powerful, as it allows you to apply security to many beans with only a simple declaration. Consider the following example:

```
<global-method-security>
<protect-pointcut expression="execution(* com.mycompany.*Service.*(..)"
  access="ROLE_USER"/>
</global-method-security>
```

This will protect all methods on beans declared in the application context whose classes are in the `com.mycompany` package and whose class names end in "Service". Only users with the `ROLE_USER` role will be able to invoke these methods. As with URL matching, the most specific matches must come first in the list of pointcuts, as the first matching expression will be used. Security annotations take precedence over pointcuts.

11.6 Domain Object Security (ACLs)**Overview**

Complex applications often will find the need to define access permissions not simply at a web request or method invocation level. Instead, security decisions need to comprise both who (`Authentication`), where (`MethodInvocation`) and what (`SomeDomainObject`). In other words, authorization decisions also need to consider the actual domain object instance subject of a method invocation.

Imagine you're designing an application for a pet clinic. There will be two main groups of users of your Spring-based application: staff of the pet clinic, as well as the pet clinic's customers. The staff will have access to all of the data, whilst your customers will only be able to see their own customer records. To make it a little more interesting, your customers can allow other users to see their customer records, such as their "puppy preschool" mentor or president of their local "Pony Club". Using Spring Security as the foundation, you have several approaches that can be used:

- Write your business methods to enforce the security. You could consult a collection within the `Customer` domain object instance to determine which users have access. By using the `SecurityContextHolder.getContext().getAuthentication()`, you'll be able to access the `Authentication` object.
- Write an `AccessDecisionVoter` to enforce the security from the `GrantedAuthority[]` s stored in the `Authentication` object. This would mean your `AuthenticationManager` would need to populate the `Authentication` with custom `GrantedAuthority[]` s representing each of the `Customer` domain object instances the principal has access to.
- Write an `AccessDecisionVoter` to enforce the security and open the target `Customer` domain object directly. This would mean your voter needs access to a DAO that allows it to retrieve the `Customer` object. It would then access the `Customer` object's collection of approved users and make the appropriate decision.

Each one of these approaches is perfectly legitimate. However, the first couples your authorization checking to your business code. The main problems with this include the enhanced difficulty of unit testing and the fact it would be more difficult to reuse the `Customer` authorization logic elsewhere. Obtaining the `GrantedAuthority[]`s from the `Authentication` object is also fine, but will not scale to large numbers of `Customer`s. If a user might be able to access 5,000 `Customer`s (unlikely in this case, but imagine if it were a popular vet for a large Pony Club!) the amount of memory consumed and time required to construct the `Authentication` object would be undesirable. The final method, opening the `Customer` directly from external code, is probably the best of the three. It achieves separation of concerns, and doesn't misuse memory or CPU cycles, but it is still inefficient in that both the `AccessDecisionVoter` and the eventual business method itself will perform a call to the DAO responsible for retrieving the `Customer` object. Two accesses per method invocation is clearly undesirable. In addition, with every approach listed you'll need to write your own access control list (ACL) persistence and business logic from scratch.

Fortunately, there is another alternative, which we'll talk about below.

Key Concepts

Spring Security's ACL services are shipped in the `spring-security-acl-xxx.jar`. You will need to add this JAR to your classpath to use Spring Security's domain object instance security capabilities.

Spring Security's domain object instance security capabilities centre on the concept of an access control list (ACL). Every domain object instance in your system has its own ACL, and the ACL records details of who can and can't work with that domain object. With this in mind, Spring Security delivers three main ACL-related capabilities to your application:

- A way of efficiently retrieving ACL entries for all of your domain objects (and modifying those ACLs)
- A way of ensuring a given principal is permitted to work with your objects, before methods are called
- A way of ensuring a given principal is permitted to work with your objects (or something they return), after methods are called

As indicated by the first bullet point, one of the main capabilities of the Spring Security ACL module is providing a high-performance way of retrieving ACLs. This ACL repository capability is extremely important, because every domain object instance in your system might have several access control entries, and each ACL might inherit from other ACLs in a tree-like structure (this is supported out-of-the-box by Spring Security, and is very commonly used). Spring Security's ACL capability has been carefully designed to provide high performance retrieval of ACLs, together with pluggable caching, deadlock-minimizing database updates, independence from ORM frameworks (we use JDBC directly), proper encapsulation, and transparent database updating.

Given databases are central to the operation of the ACL module, let's explore the four main tables used by default in the implementation. The tables are presented below in order of size in a typical Spring Security ACL deployment, with the table with the most rows listed last:

- `ACL_SID` allows us to uniquely identify any principal or authority in the system ("SID" stands for "security identity"). The only columns are the ID, a textual representation of the SID, and a flag to indicate whether the textual representation refers to a principal name or a `GrantedAuthority`. Thus, there is a single row for each unique principal or `GrantedAuthority`. When used in the context of receiving a permission, a SID is generally called a "recipient".

- `ACL_CLASS` allows us to uniquely identify any domain object class in the system. The only columns are the ID and the Java class name. Thus, there is a single row for each unique Class we wish to store ACL permissions for.
- `ACL_OBJECT_IDENTITY` stores information for each unique domain object instance in the system. Columns include the ID, a foreign key to the `ACL_CLASS` table, a unique identifier so we know which `ACL_CLASS` instance we're providing information for, the parent, a foreign key to the `ACL_SID` table to represent the owner of the domain object instance, and whether we allow ACL entries to inherit from any parent ACL. We have a single row for every domain object instance we're storing ACL permissions for.
- Finally, `ACL_ENTRY` stores the individual permissions assigned to each recipient. Columns include a foreign key to the `ACL_OBJECT_IDENTITY`, the recipient (ie a foreign key to `ACL_SID`), whether we'll be auditing or not, and the integer bit mask that represents the actual permission being granted or denied. We have a single row for every recipient that receives a permission to work with a domain object.

As mentioned in the last paragraph, the ACL system uses integer bit masking. Don't worry, you need not be aware of the finer points of bit shifting to use the ACL system, but suffice to say that we have 32 bits we can switch on or off. Each of these bits represents a permission, and by default the permissions are read (bit 0), write (bit 1), create (bit 2), delete (bit 3) and administer (bit 4). It's easy to implement your own `Permission` instance if you wish to use other permissions, and the remainder of the ACL framework will operate without knowledge of your extensions.

It is important to understand that the number of domain objects in your system has absolutely no bearing on the fact we've chosen to use integer bit masking. Whilst you have 32 bits available for permissions, you could have billions of domain object instances (which will mean billions of rows in `ACL_OBJECT_IDENTITY` and quite probably `ACL_ENTRY`). We make this point because we've found sometimes people mistakenly believe they need a bit for each potential domain object, which is not the case.

Now that we've provided a basic overview of what the ACL system does, and what it looks like at a table structure, let's explore the key interfaces. The key interfaces are:

- `Acl`: Every domain object has one and only one `Acl` object, which internally holds the `AccessControlEntry` s as well as knows the owner of the `Acl`. An `Acl` does not refer directly to the domain object, but instead to an `ObjectIdentity`. The `Acl` is stored in the `ACL_OBJECT_IDENTITY` table.
- `AccessControlEntry`: An `Acl` holds multiple `AccessControlEntry` s, which are often abbreviated as ACEs in the framework. Each ACE refers to a specific tuple of `Permission`, `Sid` and `Acl`. An ACE can also be granting or non-granting and contain audit settings. The ACE is stored in the `ACL_ENTRY` table.
- `Permission`: A permission represents a particular immutable bit mask, and offers convenience functions for bit masking and outputting information. The basic permissions presented above (bits 0 through 4) are contained in the `BasePermission` class.
- `Sid`: The ACL module needs to refer to principals and `GrantedAuthority[]` s. A level of indirection is provided by the `Sid` interface, which is an abbreviation of "security identity". Common classes include `PrincipalSid` (to represent the principal inside an `Authentication` object) and `GrantedAuthoritySid`. The security identity information is stored in the `ACL_SID` table.

- `ObjectIdentity`: Each domain object is represented internally within the ACL module by an `ObjectIdentity`. The default implementation is called `ObjectIdentityImpl`.
- `AclService`: Retrieves the `Acl` applicable for a given `ObjectIdentity`. In the included implementation (`JdbcAclService`), retrieval operations are delegated to a `LookupStrategy`. The `LookupStrategy` provides a highly optimized strategy for retrieving ACL information, using batched retrievals (`BasicLookupStrategy`) and supporting custom implementations that leverage materialized views, hierarchical queries and similar performance-centric, non-ANSI SQL capabilities.
- `MutableAclService`: Allows a modified `Acl` to be presented for persistence. It is not essential to use this interface if you do not wish.

Please note that our out-of-the-box `AclService` and related database classes all use ANSI SQL. This should therefore work with all major databases. At the time of writing, the system had been successfully tested using Hypersonic SQL, PostgreSQL, Microsoft SQL Server and Oracle.

Two samples ship with Spring Security that demonstrate the ACL module. The first is the Contacts Sample, and the other is the Document Management System (DMS) Sample. We suggest taking a look over these for examples.

Getting Started

To get starting using Spring Security's ACL capability, you will need to store your ACL information somewhere. This necessitates the instantiation of a `DataSource` using Spring. The `DataSource` is then injected into a `JdbcMutableAclService` and `BasicLookupStrategy` instance. The latter provides high-performance ACL retrieval capabilities, and the former provides mutator capabilities. Refer to one of the samples that ship with Spring Security for an example configuration. You'll also need to populate the database with the four ACL-specific tables listed in the last section (refer to the ACL samples for the appropriate SQL statements).

Once you've created the required schema and instantiated `JdbcMutableAclService`, you'll next need to ensure your domain model supports interoperability with the Spring Security ACL package. Hopefully `ObjectIdentityImpl` will prove sufficient, as it provides a large number of ways in which it can be used. Most people will have domain objects that contain a public `Serializable getId()` method. If the return type is `long`, or compatible with `long` (eg an `int`), you will find you need not give further consideration to `ObjectIdentity` issues. Many parts of the ACL module rely on long identifiers. If you're not using `long` (or an `int`, `byte` etc), there is a very good chance you'll need to reimplement a number of classes. We do not intend to support non-long identifiers in Spring Security's ACL module, as longs are already compatible with all database sequences, the most common identifier data type, and are of sufficient length to accommodate all common usage scenarios.

The following fragment of code shows how to create an `Acl`, or modify an existing `Acl`:

```
// Prepare the information we'd like in our access control entry (ACE)
ObjectIdentity oi = new ObjectIdentityImpl(Foo.class, new Long(44));
Sid sid = new PrincipalSid("Samantha");
Permission p = BasePermission.ADMINISTRATION;

// Create or update the relevant ACL
MutableAcl acl = null;
try {
    acl = (MutableAcl) aclService.readAclById(oi);
} catch (NotFoundException nfe) {
    acl = aclService.createAcl(oi);
}

// Now grant some permissions via an access control entry (ACE)
acl.insertAce(acl.getEntries().length, p, sid, true);
aclService.updateAcl(acl);
```

In the example above, we're retrieving the ACL associated with the "Foo" domain object with identifier number 44. We're then adding an ACE so that a principal named "Samantha" can "administer" the object. The code fragment is relatively self-explanatory, except the `insertAce` method. The first argument to the `insertAce` method is determining at what position in the `Acl` the new entry will be inserted. In the example above, we're just putting the new ACE at the end of the existing ACEs. The final argument is a `Boolean` indicating whether the ACE is granting or denying. Most of the time it will be granting (`true`), but if it is denying (`false`), the permissions are effectively being blocked.

Spring Security does not provide any special integration to automatically create, update or delete ACLs as part of your DAO or repository operations. Instead, you will need to write code like shown above for your individual domain objects. It's worth considering using AOP on your services layer to automatically integrate the ACL information with your services layer operations. We've found this quite an effective approach in the past.

Once you've used the above techniques to store some ACL information in the database, the next step is to actually use the ACL information as part of authorization decision logic. You have a number of choices here. You could write your own `AccessDecisionVoter` or `AfterInvocationProvider` that respectively fires before or after a method invocation. Such classes would use `AclService` to retrieve the relevant ACL and then call `Acl.isGranted(Permission[] permission, Sid[] sids, boolean administrativeMode)` to decide whether permission is granted or denied. Alternately, you could use our `AclEntryVoter`, `AclEntryAfterInvocationProvider` or `AclEntryAfterInvocationCollectionFilteringProvider` classes. All of these classes provide a declarative-based approach to evaluating ACL information at runtime, freeing you from needing to write any code. Please refer to the sample applications to learn how to use these classes.

12. OAuth2

12.1 OAuth 2.0 Login

The OAuth 2.0 Login feature provides an application with the capability to have users log in to the application by using their existing account at an OAuth 2.0 Provider (e.g. GitHub) or OpenID Connect 1.0 Provider (such as Google). OAuth 2.0 Login implements the use cases: "Login with Google" or "Login with GitHub".

Note

OAuth 2.0 Login is implemented by using the **Authorization Code Grant**, as specified in the [OAuth 2.0 Authorization Framework](#) and [OpenID Connect Core 1.0](#).

Spring Boot 2.x Sample

Spring Boot 2.x brings full auto-configuration capabilities for OAuth 2.0 Login.

This section shows how to configure the [OAuth 2.0 Login sample](#) using *Google* as the *Authentication Provider* and covers the following topics:

- [Initial setup](#)
- [Setting the redirect URI](#)
- [Configure application.yml](#)
- [Boot up the application](#)

Initial setup

To use Google's OAuth 2.0 authentication system for login, you must set up a project in the Google API Console to obtain OAuth 2.0 credentials.

Note

[Google's OAuth 2.0 implementation](#) for authentication conforms to the [OpenID Connect 1.0](#) specification and is [OpenID Certified](#).

Follow the instructions on the [OpenID Connect](#) page, starting in the section, "Setting up OAuth 2.0".

After completing the "Obtain OAuth 2.0 credentials" instructions, you should have a new OAuth Client with credentials consisting of a Client ID and a Client Secret.

Setting the redirect URI

The redirect URI is the path in the application that the end-user's user-agent is redirected back to after they have authenticated with Google and have granted access to the OAuth Client ([created in the previous step](#)) on the Consent page.

In the "Set a redirect URI" sub-section, ensure that the **Authorized redirect URIs** field is set to <http://localhost:8080/login/oauth2/code/google>.

Tip

The default redirect URI template is `{baseUrl}/login/oauth2/code/{registrationId}`. The `registrationId` is a unique identifier for the [ClientRegistration](#).

Important

If the OAuth Client is running behind a proxy server, it is recommended to check [Proxy Server Configuration](#) to ensure the application is correctly configured. Also, see the supported [URI template variables](#) for `redirect-uri`.

Configure application.yml

Now that you have a new OAuth Client with Google, you need to configure the application to use the OAuth Client for the *authentication flow*. To do so:

1. Go to `application.yml` and set the following configuration:

```
spring:
  security:
    oauth2:
      client:
        registration: ❶
        google: ❷
          client-id: google-client-id
          client-secret: google-client-secret
```

- ❶ `spring.security.oauth2.client.registration` is the base property prefix for OAuth Client properties.
- ❷ Following the base property prefix is the ID for the [ClientRegistration](#), such as `google`.

Example 12.1 OAuth Client properties

2. Replace the values in the `client-id` and `client-secret` property with the OAuth 2.0 credentials you created earlier.

Boot up the application

Launch the Spring Boot 2.x sample and go to <http://localhost:8080>. You are then redirected to the default *auto-generated* login page, which displays a link for Google.

Click on the Google link, and you are then redirected to Google for authentication.

After authenticating with your Google account credentials, the next page presented to you is the Consent screen. The Consent screen asks you to either allow or deny access to the OAuth Client you created earlier. Click **Allow** to authorize the OAuth Client to access your email address and basic profile information.

At this point, the OAuth Client retrieves your email address and basic profile information from the [UserInfo Endpoint](#) and establishes an authenticated session.

Spring Boot 2.x Property Mappings

The following table outlines the mapping of the Spring Boot 2.x OAuth Client properties to the [ClientRegistration](#) properties.

Spring Boot 2.x	ClientRegistration
spring.security.oauth2.client.registration.registrationId	registrationId
spring.security.oauth2.client.registration.registrationId.client-id	clientId
spring.security.oauth2.client.registration.registrationId.client-secret	secret
spring.security.oauth2.client.registration.registrationId.authentication-method	authenticationMethod
spring.security.oauth2.client.registration.registrationId.authorization-grant-type	authorizationGrantType
spring.security.oauth2.client.registration.registrationId.redirect-uri	redirectUri
spring.security.oauth2.client.registration.registrationId.scope	scope
spring.security.oauth2.client.registration.registrationId.client-name	clientId
spring.security.oauth2.client.provider.providerId.authorization-uri	authorizationUri
spring.security.oauth2.client.provider.providerId.token-uri	tokenUri
spring.security.oauth2.client.provider.providerId.jwt-set-uri	jwtSetUri
spring.security.oauth2.client.provider.providerId.user-info-endpoint-uri	userInfoEndpointUri
spring.security.oauth2.client.provider.providerId.user-info-endpoint-authentication-method	userInfoEndpointAuthenticationMethod
spring.security.oauth2.client.provider.providerId.user-name-endpoint-uri	userNameEndpointUri

Tip

A `ClientRegistration` can be initially configured using discovery of an OpenID Connect Provider's [Configuration endpoint](#) or an Authorization Server's [Metadata endpoint](#), by specifying the `spring.security.oauth2.client.provider.[providerId].issuer-uri` property.

CommonOAuth2Provider

`CommonOAuth2Provider` pre-defines a set of default client properties for a number of well known providers: Google, GitHub, Facebook, and Okta.

For example, the `authorization-uri`, `token-uri`, and `user-info-uri` do not change often for a Provider. Therefore, it makes sense to provide default values in order to reduce the required configuration.

As demonstrated previously, when we [configured a Google client](#), only the `client-id` and `client-secret` properties are required.

The following listing shows an example:

```
spring:
  security:
    oauth2:
      client:
        registration:
          google:
            client-id: google-client-id
            client-secret: google-client-secret
```

Tip

The auto-defaulting of client properties works seamlessly here because the `registrationId` (`google`) matches the `GOOGLE` enum (case-insensitive) in `CommonOAuth2Provider`.

For cases where you may want to specify a different `registrationId`, such as `google-login`, you can still leverage auto-defaulting of client properties by configuring the `provider` property.

The following listing shows an example:

```
spring:
  security:
    oauth2:
      client:
        registration:
          google-login: ❶
            provider: google ❷
            client-id: google-client-id
            client-secret: google-client-secret
```

- ❶ The `registrationId` is set to `google-login`.
- ❷ The `provider` property is set to `google`, which will leverage the auto-defaulting of client properties set in `CommonOAuth2Provider.GOOGLE.getBuilder()`.

Configuring Custom Provider Properties

There are some OAuth 2.0 Providers that support multi-tenancy, which results in different protocol endpoints for each tenant (or sub-domain).

For example, an OAuth Client registered with Okta is assigned to a specific sub-domain and have their own protocol endpoints.

For these cases, Spring Boot 2.x provides the following base property for configuring custom provider properties: `spring.security.oauth2.client.provider.[providerId]`.

The following listing shows an example:

```
spring:
  security:
    oauth2:
      client:
        registration:
          okta:
            client-id: okta-client-id
            client-secret: okta-client-secret
        provider:
          okta: ❶
            authorization-uri: https://your-subdomain.oktapreview.com/oauth2/v1/authorize
            token-uri: https://your-subdomain.oktapreview.com/oauth2/v1/token
            user-info-uri: https://your-subdomain.oktapreview.com/oauth2/v1/userinfo
            user-name-attribute: sub
            jwk-set-uri: https://your-subdomain.oktapreview.com/oauth2/v1/keys
```

- ❶ The base property (`spring.security.oauth2.client.provider.okta`) allows for custom configuration of protocol endpoint locations.

Overriding Spring Boot 2.x Auto-configuration

The Spring Boot 2.x auto-configuration class for OAuth Client support is `OAuth2ClientAutoConfiguration`.

It performs the following tasks:

- Registers a `ClientRegistrationRepository @Bean` composed of `ClientRegistration(s)` from the configured OAuth Client properties.
- Provides a `WebSecurityConfigurerAdapter @Configuration` and enables OAuth 2.0 Login through `httpSecurity.oauth2Login()`.

If you need to override the auto-configuration based on your specific requirements, you may do so in the following ways:

- [Register a ClientRegistrationRepository @Bean](#)
- [Provide a WebSecurityConfigurerAdapter](#)
- [Completely Override the Auto-configuration](#)

Register a ClientRegistrationRepository @Bean

The following example shows how to register a `ClientRegistrationRepository @Bean`:

```

@Configuration
public class OAuth2LoginConfig {

    @Bean
    public ClientRegistrationRepository clientRegistrationRepository() {
        return new InMemoryClientRegistrationRepository(this.googleClientRegistration());
    }

    private ClientRegistration googleClientRegistration() {
        return ClientRegistration.withRegistrationId("google")
            .clientId("google-client-id")
            .clientSecret("google-client-secret")
            .clientAuthenticationMethod(ClientAuthenticationMethod.BASIC)
            .authorizationGrantType(AuthorizationGrantType.AUTHORIZATION_CODE)
            .redirectUriTemplate("{baseUrl}/login/oauth2/code/{registrationId}")
            .scope("openid", "profile", "email", "address", "phone")
            .authorizationUri("https://accounts.google.com/o/oauth2/v2/auth")
            .tokenUri("https://www.googleapis.com/oauth2/v4/token")
            .userInfoUri("https://www.googleapis.com/oauth2/v3/userinfo")
            .userNameAttributeName(IdTokenClaimNames.SUB)
            .jwkSetUri("https://www.googleapis.com/oauth2/v3/certs")
            .clientName("Google")
            .build();
    }
}

```

Provide a WebSecurityConfigurerAdapter

The following example shows how to provide a `WebSecurityConfigurerAdapter` with `@EnableWebSecurity` and enable OAuth 2.0 login through `httpSecurity.oauth2Login()`:

```

@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .authorizeRequests(authorizeRequests ->
                authorizeRequests
                    .anyRequest().authenticated()
            )
            .oauth2Login(withDefaults());
    }
}

```

Completely Override the Auto-configuration

The following example shows how to completely override the auto-configuration by registering a `ClientRegistrationRepository` `@Bean` and providing a `WebSecurityConfigurerAdapter`.

```

@Configuration
public class OAuth2LoginConfig {

    @EnableWebSecurity
    public static class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

        @Override
        protected void configure(HttpSecurity http) throws Exception {
            http
                .authorizeRequests(authorizeRequests ->
                    authorizeRequests
                        .anyRequest().authenticated()
                )
                .oauth2Login(withDefaults());
        }
    }

    @Bean
    public ClientRegistrationRepository clientRegistrationRepository() {
        return new InMemoryClientRegistrationRepository(this.googleClientRegistration());
    }

    private ClientRegistration googleClientRegistration() {
        return ClientRegistration.withRegistrationId("google")
            .clientId("google-client-id")
            .clientSecret("google-client-secret")
            .clientAuthenticationMethod(ClientAuthenticationMethod.BASIC)
            .authorizationGrantType(AuthorizationGrantType.AUTHORIZATION_CODE)
            .redirectUriTemplate("{baseUrl}/login/oauth2/code/{registrationId}")
            .scope("openid", "profile", "email", "address", "phone")
            .authorizationUri("https://accounts.google.com/o/oauth2/v2/auth")
            .tokenUri("https://www.googleapis.com/oauth2/v4/token")
            .userInfoUri("https://www.googleapis.com/oauth2/v3/userinfo")
            .userNameAttributeName(IdTokenClaimNames.SUB)
            .jwkSetUri("https://www.googleapis.com/oauth2/v3/certs")
            .clientName("Google")
            .build();
    }
}

```

Java Configuration without Spring Boot 2.x

If you are not able to use Spring Boot 2.x and would like to configure one of the pre-defined providers in `CommonOAuth2Provider` (for example, Google), apply the following configuration:

```

@Configuration
public class OAuth2LoginConfig {

    @EnableWebSecurity
    public static class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

        @Override
        protected void configure(HttpSecurity http) throws Exception {
            http
                .authorizeRequests(authorizeRequests ->
                    authorizeRequests
                        .anyRequest().authenticated()
                )
                .oauth2Login(withDefaults());
        }
    }

    @Bean
    public ClientRegistrationRepository clientRegistrationRepository() {
        return new InMemoryClientRegistrationRepository(this.googleClientRegistration());
    }

    @Bean
    public OAuth2AuthorizedClientService authorizedClientService(
        ClientRegistrationRepository clientRegistrationRepository) {
        return new InMemoryOAuth2AuthorizedClientService(clientRegistrationRepository);
    }

    @Bean
    public OAuth2AuthorizedClientRepository authorizedClientRepository(
        OAuth2AuthorizedClientService authorizedClientService) {
        return new AuthenticatedPrincipalOAuth2AuthorizedClientRepository(authorizedClientService);
    }

    private ClientRegistration googleClientRegistration() {
        return CommonOAuth2Provider.GOOGLE.getBuilder("google")
            .clientId("google-client-id")
            .clientSecret("google-client-secret")
            .build();
    }
}

```

Advanced Configuration

`HttpSecurity.oauth2Login()` provides a number of configuration options for customizing OAuth 2.0 Login. The main configuration options are grouped into their protocol endpoint counterparts.

For example, `oauth2Login().authorizationEndpoint()` allows configuring the *Authorization Endpoint*, whereas `oauth2Login().tokenEndpoint()` allows configuring the *Token Endpoint*.

The following code shows an example:

```

@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .authorizationEndpoint(authorizationEndpoint ->
                        authorizationEndpoint
                            ...
                    )
                    .redirectionEndpoint(redirectionEndpoint ->
                        redirectionEndpoint
                            ...
                    )
                    .tokenEndpoint(tokenEndpoint ->
                        tokenEndpoint
                            ...
                    )
                    .userInfoEndpoint(userInfoEndpoint ->
                        userInfoEndpoint
                            ...
                    )
            );
    }
}

```

The main goal of the `oauth2Login()` DSL was to closely align with the naming, as defined in the specifications.

The OAuth 2.0 Authorization Framework defines the [Protocol Endpoints](#) as follows:

The authorization process utilizes two authorization server endpoints (HTTP resources):

- Authorization Endpoint: Used by the client to obtain authorization from the resource owner via user-agent redirection.
- Token Endpoint: Used by the client to exchange an authorization grant for an access token, typically with client authentication.

As well as one client endpoint:

- Redirection Endpoint: Used by the authorization server to return responses containing authorization credentials to the client via the resource owner user-agent.

The OpenID Connect Core 1.0 specification defines the [UserInfo Endpoint](#) as follows:

The UserInfo Endpoint is an OAuth 2.0 Protected Resource that returns claims about the authenticated end-user. To obtain the requested claims about the end-user, the client makes a request to the UserInfo Endpoint by using an access token obtained through OpenID Connect Authentication. These claims are normally represented by a JSON object that contains a collection of name-value pairs for the claims.

The following code shows the complete configuration options available for the `oauth2Login()` DSL:

```

@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .clientRegistrationRepository(this.clientRegistrationRepository())
                    .authorizedClientRepository(this.authorizedClientRepository())
                    .authorizedClientService(this.authorizedClientService())
                    .loginPage("/login")
                    .authorizationEndpoint(authorizationEndpoint ->
                        authorizationEndpoint
                            .baseUri(this.authorizationRequestBaseUri())
                            .authorizationRequestRepository(this.authorizationRequestRepository())
                            .authorizationRequestResolver(this.authorizationRequestResolver())
                        )
                    .redirectionEndpoint(redirectionEndpoint ->
                        redirectionEndpoint
                            .baseUri(this.authorizationResponseBaseUri())
                        )
                    .tokenEndpoint(tokenEndpoint ->
                        tokenEndpoint
                            .accessTokenResponseClient(this.accessTokenResponseClient())
                        )
                    .userInfoEndpoint(userInfoEndpoint ->
                        userInfoEndpoint
                            .userAuthoritiesMapper(this.userAuthoritiesMapper())
                            .userService(this.oauth2UserService())
                            .oidcUserService(this.oidcUserService())
                            .customUserType(GitHubOAuth2User.class, "github")
                        )
                )
            );
    }
}

```

The following sections go into more detail on each of the configuration options available:

- [OAuth 2.0 Login Page](#)
- [Redirection Endpoint](#)
- [UserInfo Endpoint](#)

OAuth 2.0 Login Page

By default, the OAuth 2.0 Login Page is auto-generated by the `DefaultLoginPageGeneratingFilter`. The default login page shows each configured OAuth Client with its `ClientRegistration.clientName` as a link, which is capable of initiating the Authorization Request (or OAuth 2.0 Login).

Note

In order for `DefaultLoginPageGeneratingFilter` to show links for configured OAuth Clients, the registered `ClientRegistrationRepository` needs to also implement `Iterable<ClientRegistration>`. See `InMemoryClientRegistrationRepository` for reference.

The link's destination for each OAuth Client defaults to the following:

```

OAuth2AuthorizationRequestRedirectFilter.DEFAULT_AUTHORIZATION_REQUEST_BASE_URI
+ "{registrationId}"

```


The following line shows an example:

```
<a href="/oauth2/authorization/google">Google</a>
```

To override the default login page, configure `oauth2Login().loginPage()` and (optionally) `oauth2Login().authorizationEndpoint().baseUri()`.

The following listing shows an example:

```
@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .loginPage("/login/oauth2")
                    ...
                .authorizationEndpoint(authorizationEndpoint ->
                    authorizationEndpoint
                        .baseUri("/login/oauth2/authorization")
                    ...
                )
            );
    }
}
```

Important

You need to provide a `@Controller` with a `@RequestMapping("/login/oauth2")` that is capable of rendering the custom login page.

Tip

As noted earlier, configuring `oauth2Login().authorizationEndpoint().baseUri()` is optional. However, if you choose to customize it, ensure the link to each OAuth Client matches the `authorizationEndpoint().baseUri()`.

The following line shows an example:

```
<a href="/login/oauth2/authorization/google">Google</a>
```

Redirection Endpoint

The Redirection Endpoint is used by the Authorization Server for returning the Authorization Response (which contains the authorization credentials) to the client via the Resource Owner user-agent.

Tip

OAuth 2.0 Login leverages the Authorization Code Grant. Therefore, the authorization credential is the authorization code.

The default Authorization Response `baseUri` (redirection endpoint) is `/login/oauth2/code/*`, which is defined in `OAuth2LoginAuthenticationFilter.DEFAULT_FILTER_PROCESSES_URI`.

If you would like to customize the Authorization Response `baseUri`, configure it as shown in the following example:

```

@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .redirectEndpoint(redirectEndpoint ->
                        redirectEndpoint
                            .baseUrl("/login/oauth2/callback/*")
                            ...
                    )
            );
    }
}

```

Important

You also need to ensure the `ClientRegistration.redirectUriTemplate` matches the custom Authorization Response `baseUrl`.

The following listing shows an example:

```

return CommonOAuth2Provider.GOOGLE.getBuilder("google")
    .clientId("google-client-id")
    .clientSecret("google-client-secret")
    .redirectUriTemplate("{baseUrl}/login/oauth2/callback/{registrationId}")
    .build();

```

User Info Endpoint

The User Info Endpoint includes a number of configuration options, as described in the following subsections:

- [Mapping User Authorities](#)
- [Configuring a Custom OAuth2User](#)
- [OAuth 2.0 UserService](#)
- [OpenID Connect 1.0 UserService](#)

Mapping User Authorities

After the user successfully authenticates with the OAuth 2.0 Provider, the `OAuth2User.getAuthorities()` (or `OidcUser.getAuthorities()`) may be mapped to a new set of `GrantedAuthority` instances, which will be supplied to `OAuth2AuthenticationToken` when completing the authentication.

Tip

`OAuth2AuthenticationToken.getAuthorities()` is used for authorizing requests, such as in `hasRole('USER')` or `hasRole('ADMIN')`.

There are a couple of options to choose from when mapping user authorities:

- [Using a GrantedAuthoritiesMapper](#)
- [Delegation-based strategy with OAuth2UserService](#)

Using a GrantedAuthoritiesMapper

Provide an implementation of `GrantedAuthoritiesMapper` and configure it as shown in the following example:

```
@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .userInfoEndpoint(userInfoEndpoint ->
                        userInfoEndpoint
                            .userAuthoritiesMapper(this.userAuthoritiesMapper())
                            ...
                    )
            );
    }

    private GrantedAuthoritiesMapper userAuthoritiesMapper() {
        return (authorities) -> {
            Set<GrantedAuthority> mappedAuthorities = new HashSet<>();

            authorities.forEach(authority -> {
                if (OidcUserAuthority.class.isInstance(authority)) {
                    OidcUserAuthority oidcUserAuthority = (OidcUserAuthority)authority;

                    OidcIdToken idToken = oidcUserAuthority.getIdToken();
                    OidcUserInfo userInfo = oidcUserAuthority.getUserInfo();

                    // Map the claims found in idToken and/or userInfo
                    // to one or more GrantedAuthority's and add it to mappedAuthorities
                } else if (OAuth2UserAuthority.class.isInstance(authority)) {
                    OAuth2UserAuthority oauth2UserAuthority = (OAuth2UserAuthority)authority;

                    Map<String, Object> userAttributes = oauth2UserAuthority.getAttributes();

                    // Map the attributes found in userAttributes
                    // to one or more GrantedAuthority's and add it to mappedAuthorities
                }
            });

            return mappedAuthorities;
        };
    }
}
```

Alternatively, you may register a `GrantedAuthoritiesMapper` `@Bean` to have it automatically applied to the configuration, as shown in the following example:

```
@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(withDefaults());
    }

    @Bean
    public GrantedAuthoritiesMapper userAuthoritiesMapper() {
        ...
    }
}
```

Delegation-based strategy with OAuth2UserService

This strategy is advanced compared to using a `GrantedAuthoritiesMapper`, however, it's also more flexible as it gives you access to the `OAuth2UserRequest` and `OAuth2User` (when using an OAuth 2.0 `UserService`) or `OidcUserRequest` and `OidcUser` (when using an OpenID Connect 1.0 `UserService`).

The `OAuth2UserRequest` (and `OidcUserRequest`) provides you access to the associated `OAuth2AccessToken`, which is very useful in the cases where the *delegator* needs to fetch authority information from a protected resource before it can map the custom authorities for the user.

The following example shows how to implement and configure a delegation-based strategy using an OpenID Connect 1.0 `UserService`:

```
@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .userInfoEndpoint(userInfoEndpoint ->
                        userInfoEndpoint
                            .oidcUserService(this.oidcUserService())
                            ...
                    )
            );
    }

    private OAuth2UserService<OidcUserRequest, OidcUser> oidcUserService() {
        final OidcUserService delegate = new OidcUserService();

        return (userRequest) -> {
            // Delegate to the default implementation for loading a user
            OidcUser oidcUser = delegate.loadUser(userRequest);

            OAuth2AccessToken accessToken = userRequest.getAccessToken();
            Set<GrantedAuthority> mappedAuthorities = new HashSet<>();

            // TODO
            // 1) Fetch the authority information from the protected resource using accessToken
            // 2) Map the authority information to one or more GrantedAuthority's and add it to
            mappedAuthorities

            // 3) Create a copy of oidcUser but use the mappedAuthorities instead
            oidcUser = new DefaultOidcUser(mappedAuthorities, oidcUser.getIdToken(),
            oidcUser.getUserInfo());

            return oidcUser;
        };
    }
}
```

Configuring a Custom OAuth2User

`CustomUserTypesOAuth2UserService` is an implementation of an `OAuth2UserService` that provides support for custom `OAuth2User` types.

If the default implementation (`DefaultOAuth2User`) does not suit your needs, you can define your own implementation of `OAuth2User`.

The following code demonstrates how you would register a custom `OAuth2User` type for GitHub:

```
@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .userInfoEndpoint(userInfoEndpoint ->
                        userInfoEndpoint
                            .customUserType(GitHubOAuth2User.class, "github")
                            ...
                    )
            );
    }
}
```

The following code shows an example of a custom `OAuth2User` type for GitHub:

```
public class GitHubOAuth2User implements OAuth2User {
    private List<GrantedAuthority> authorities =
        AuthorityUtils.createAuthorityList("ROLE_USER");
    private Map<String, Object> attributes;
    private String id;
    private String name;
    private String login;
    private String email;

    @Override
    public Collection<? extends GrantedAuthority> getAuthorities() {
        return this.authorities;
    }

    @Override
    public Map<String, Object> getAttributes() {
        if (this.attributes == null) {
            this.attributes = new HashMap<>();
            this.attributes.put("id", this.getId());
            this.attributes.put("name", this.getName());
            this.attributes.put("login", this.getLogin());
            this.attributes.put("email", this.getEmail());
        }
        return attributes;
    }

    public String getId() {
        return this.id;
    }

    public void setId(String id) {
        this.id = id;
    }

    @Override
    public String getName() {
        return this.name;
    }

    public void setName(String name) {
        this.name = name;
    }

    public String getLogin() {
        return this.login;
    }

    public void setLogin(String login) {
        this.login = login;
    }

    public String getEmail() {
        return this.email;
    }

    public void setEmail(String email) {
        this.email = email;
    }
}
```

Tip

id, name, login, and email are attributes returned in GitHub's UserInfo Response. For detailed information returned from the UserInfo Endpoint, see the API documentation for ["Get the authenticated user"](#).

OAuth 2.0 UserService

`DefaultOAuth2UserService` is an implementation of an `OAuth2UserService` that supports standard OAuth 2.0 Provider's.

Note

`OAuth2UserService` obtains the user attributes of the end-user (the resource owner) from the `Userinfo` Endpoint (by using the access token granted to the client during the authorization flow) and returns an `AuthenticatedPrincipal` in the form of an `OAuth2User`.

`DefaultOAuth2UserService` uses a `RestOperations` when requesting the user attributes at the `Userinfo` Endpoint.

If you need to customize the pre-processing of the `Userinfo` Request, you can provide `DefaultOAuth2UserService.setRequestEntityConverter()` with a custom `Converter<OAuth2UserRequest, RequestEntity<?>>`. The default implementation `OAuth2UserRequestEntityConverter` builds a `RequestEntity` representation of a `Userinfo` Request that sets the `OAuth2AccessToken` in the `Authorization` header by default.

On the other end, if you need to customize the post-handling of the `Userinfo` Response, you will need to provide `DefaultOAuth2UserService.setRestOperations()` with a custom configured `RestOperations`. The default `RestOperations` is configured as follows:

```
RestTemplate restTemplate = new RestTemplate();
restTemplate.setErrorHandler(new OAuth2ErrorResponseErrorHandler());
```

`OAuth2ErrorResponseErrorHandler` is a `ResponseErrorHandler` that can handle an OAuth 2.0 Error (400 Bad Request). It uses an `OAuth2ErrorHttpMessageConverter` for converting the OAuth 2.0 Error parameters to an `OAuth2Error`.

Whether you customize `DefaultOAuth2UserService` or provide your own implementation of `OAuth2UserService`, you'll need to configure it as shown in the following example:

```
@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .userInfoEndpoint(userInfoEndpoint ->
                        userInfoEndpoint
                            .userService(this.oauth2UserService())
                            ...
                    )
            );
    }

    private OAuth2UserService<OAuth2UserRequest, OAuth2User> oauth2UserService() {
        ...
    }
}
```

OpenID Connect 1.0 UserService

`OidcUserService` is an implementation of an `OAuth2UserService` that supports OpenID Connect 1.0 Provider's.

The `OidcUserService` leverages the `DefaultOAuth2UserService` when requesting the user attributes at the `UserInfo` Endpoint.

If you need to customize the pre-processing of the `UserInfo` Request and/or the post-handling of the `UserInfo` Response, you will need to provide `OidcUserService.setOAuth2UserService()` with a custom configured `DefaultOAuth2UserService`.

Whether you customize `OidcUserService` or provide your own implementation of `OAuth2UserService` for OpenID Connect 1.0 Provider's, you'll need to configure it as shown in the following example:

```
@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .userInfoEndpoint(userInfoEndpoint ->
                        userInfoEndpoint
                            .oidcUserService(this.oidcUserService())
                            ...
                    )
            );
    }

    private OAuth2UserService<OidcUserRequest, OidcUser> oidcUserService() {
        ...
    }
}
```

ID Token Signature Verification

OpenID Connect 1.0 Authentication introduces the [ID Token](#), which is a security token that contains Claims about the Authentication of an End-User by an Authorization Server when used by a Client.

The ID Token is represented as a [JSON Web Token](#) (JWT) and MUST be signed using [JSON Web Signature](#) (JWS).

The `OidcIdTokenDecoderFactory` provides a `JwtDecoder` used for `OidcIdToken` signature verification. The default algorithm is `RS256` but may be different when assigned during client registration. For these cases, a resolver may be configured to return the expected JWS algorithm assigned for a specific client.

The JWS algorithm resolver is a Function that accepts a `ClientRegistration` and returns the expected `JwsAlgorithm` for the client, eg. `SignatureAlgorithm.RS256` or `MacAlgorithm.HS256`

The following code shows how to configure the `OidcIdTokenDecoderFactory` @Bean to default to `MacAlgorithm.HS256` for all `ClientRegistration`:

```
@Bean
public JwtDecoderFactory<ClientRegistration> idTokenDecoderFactory() {
    OidcIdTokenDecoderFactory idTokenDecoderFactory = new OidcIdTokenDecoderFactory();
    idTokenDecoderFactory.setJwsAlgorithmResolver(clientRegistration -> MacAlgorithm.HS256);
    return idTokenDecoderFactory;
}
```


Note

For MAC based algorithms such as HS256, HS384 or HS512, the `client-secret` corresponding to the `client-id` is used as the symmetric key for signature verification.

Tip

If more than one `ClientRegistration` is configured for OpenID Connect 1.0 Authentication, the JWS algorithm resolver may evaluate the provided `ClientRegistration` to determine which algorithm to return.

OpenID Connect 1.0 Logout

OpenID Connect Session Management 1.0 allows the ability to log out the End-User at the Provider using the Client. One of the strategies available is [RP-Initiated Logout](#).

If the OpenID Provider supports both Session Management and [Discovery](#), the client may obtain the `end_session_endpoint` URL from the OpenID Provider's [Discovery Metadata](#). This can be achieved by configuring the `ClientRegistration` with the `issuer-uri`, as in the following example:

```
spring:
  security:
    oauth2:
      client:
        registration:
          okta:
            client-id: okta-client-id
            client-secret: okta-client-secret
            ...
        provider:
          okta:
            issuer-uri: https://dev-1234.oktapreview.com
```

...and the `OidcClientInitiatedLogoutSuccessHandler`, which implements RP-Initiated Logout, may be configured as follows:

```

@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Autowired
    private ClientRegistrationRepository clientRegistrationRepository;

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .authorizeRequests(authorizeRequests ->
                authorizeRequests
                    .anyRequest().authenticated()
                )
            .oauth2Login(withDefaults())
            .logout(logout ->
                logout
                    .logoutSuccessHandler(oidcLogoutSuccessHandler())
                );
    }

    private LogoutSuccessHandler oidcLogoutSuccessHandler() {
        OidcClientInitiatedLogoutSuccessHandler oidcLogoutSuccessHandler =
            new OidcClientInitiatedLogoutSuccessHandler(this.clientRegistrationRepository);

        // Sets the `URI` that the End-User's User Agent will be redirected to
        // after the logout has been performed at the Provider
        oidcLogoutSuccessHandler.setPostLogoutRedirectUri(URI.create("https://localhost:8080"));

        return oidcLogoutSuccessHandler;
    }
}

```

12.2 OAuth 2.0 Client

The OAuth 2.0 Client features provide support for the Client role as defined in the [OAuth 2.0 Authorization Framework](#).

At a high-level, the core features available are:

Authorization Grant support

- [Authorization Code](#)
- [Refresh Token](#)
- [Client Credentials](#)
- [Resource Owner Password Credentials](#)

HTTP Client support

- [webClient integration for Servlet Environments](#) (for requesting protected resources)

The `HttpSecurity.oauth2Client()` DSL provides a number of configuration options for customizing the core components used by OAuth 2.0 Client. In addition, `HttpSecurity.oauth2Client().authorizationCodeGrant()` enables the customization of the Authorization Code grant.

The following code shows the complete configuration options provided by the `HttpSecurity.oauth2Client()` DSL:

```

@EnableWebSecurity
public class OAuth2ClientSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Client(oauth2Client ->
                oauth2Client
                    .clientRegistrationRepository(this.clientRegistrationRepository())
                    .authorizedClientRepository(this.authorizedClientRepository())
                    .authorizedClientService(this.authorizedClientService())
                    .authorizationCodeGrant(authorizationCodeGrant ->
                        authorizationCodeGrant
                            .authorizationRequestRepository(this.authorizationRequestRepository())
                            .authorizationRequestResolver(this.authorizationRequestResolver())
                            .accessTokenResponseClient(this.accessTokenResponseClient())
                    )
                )
            );
    }
}

```

The `OAuth2AuthorizedClientManager` is responsible for managing the authorization (or re-authorization) of an OAuth 2.0 Client, in collaboration with one or more `OAuth2AuthorizedClientProvider(s)`.

The following code shows an example of how to register an `OAuth2AuthorizedClientManager` `@Bean` and associate it with an `OAuth2AuthorizedClientProvider` composite that provides support for the `authorization_code`, `refresh_token`, `client_credentials` and `password` authorization grant types:

```

@Bean
public OAuth2AuthorizedClientManager authorizedClientManager(
    ClientRegistrationRepository clientRegistrationRepository,
    OAuth2AuthorizedClientRepository authorizedClientRepository) {

    OAuth2AuthorizedClientProvider authorizedClientProvider =
        OAuth2AuthorizedClientProviderBuilder.builder()
            .authorizationCode()
            .refreshToken()
            .clientCredentials()
            .password()
            .build();

    DefaultOAuth2AuthorizedClientManager authorizedClientManager =
        new DefaultOAuth2AuthorizedClientManager(
            clientRegistrationRepository, authorizedClientRepository);
    authorizedClientManager.setAuthorizedClientProvider(authorizedClientProvider);

    return authorizedClientManager;
}

```

The following sections will go into more detail on the core components used by OAuth 2.0 Client and the configuration options available:

- the section called “Core Interfaces / Classes”
 - [ClientRegistration](#)
 - [ClientRegistrationRepository](#)
 - [OAuth2AuthorizedClient](#)
 - [OAuth2AuthorizedClientRepository / OAuth2AuthorizedClientService](#)

- [OAuth2AuthorizedClientManager / OAuth2AuthorizedClientProvider](#)
- the section called “Authorization Grant Support”
 - [Authorization Code](#)
 - [Refresh Token](#)
 - [Client Credentials](#)
 - [Resource Owner Password Credentials](#)
- the section called “Additional Features”
 - [Resolving an Authorized Client](#)
- the section called “WebClient integration for Servlet Environments”

Core Interfaces / Classes

ClientRegistration

`ClientRegistration` is a representation of a client registered with an OAuth 2.0 or OpenID Connect 1.0 Provider.

A client registration holds information, such as client id, client secret, authorization grant type, redirect URI, scope(s), authorization URI, token URI, and other details.

`ClientRegistration` and its properties are defined as follows:

```
public final class ClientRegistration {
    private String registrationId; ❶
    private String clientId; ❷
    private String clientSecret; ❸
    private ClientAuthenticationMethod clientAuthenticationMethod; ❹
    private AuthorizationGrantType authorizationGrantType; ❺
    private String redirectUriTemplate; ❻
    private Set<String> scopes; ❼
    private ProviderDetails providerDetails;
    private String clientName; ❽

    public class ProviderDetails {
        private String authorizationUri; ❾
        private String tokenUri; ❿
        private UserInfoEndpoint userInfoEndpoint;
        private String jwkSetUri; 11
        private Map<String, Object> configurationMetadata; 12

        public class UserInfoEndpoint {
            private String uri; 13
            private AuthenticationMethod authenticationMethod; 14
            private String userNameAttributeName; 15
        }
    }
}
```

- ❶ `registrationId`: The ID that uniquely identifies the `ClientRegistration`.
- ❷ `clientId`: The client identifier.

- ③ `clientSecret`: The client secret.
- ④ `clientAuthenticationMethod`: The method used to authenticate the Client with the Provider. The supported values are **basic**, **post** and **none** ([public clients](#)).
- ⑤ `authorizationGrantType`: The OAuth 2.0 Authorization Framework defines four [Authorization Grant](#) types. The supported values are `authorization_code`, `client_credentials`, `password` and `implicit`.
- ⑥ `redirectUriTemplate`: The client's registered redirect URI that the *Authorization Server* redirects the end-user's user-agent to after the end-user has authenticated and authorized access to the client.
- ⑦ `scopes`: The scope(s) requested by the client during the Authorization Request flow, such as `openid`, `email`, or `profile`.
- ⑧ `clientName`: A descriptive name used for the client. The name may be used in certain scenarios, such as when displaying the name of the client in the auto-generated login page.
- ⑨ `authorizationUri`: The Authorization Endpoint URI for the Authorization Server.
- ⑩ `tokenUri`: The Token Endpoint URI for the Authorization Server.
- ⑪ `jwkSetUri`: The URI used to retrieve the [JSON Web Key \(JWK\)](#) Set from the Authorization Server, which contains the cryptographic key(s) used to verify the [JSON Web Signature \(JWS\)](#) of the ID Token and optionally the UserInfo Response.
- ⑫ `configurationMetadata`: The [OpenID Provider Configuration Information](#). This information will only be available if the Spring Boot 2.x property `spring.security.oauth2.client.provider.[providerId].issuerUri` is configured.
- ⑬ `(userInfoEndpoint)uri`: The UserInfo Endpoint URI used to access the claims/attributes of the authenticated end-user.
- ⑭ `(userInfoEndpoint)authenticationMethod`: The authentication method used when sending the access token to the UserInfo Endpoint. The supported values are **header**, **form** and **query**.
- ⑮ `userNameAttributeName`: The name of the attribute returned in the UserInfo Response that references the Name or Identifier of the end-user.

A `ClientRegistration` can be initially configured using discovery of an OpenID Connect Provider's [Configuration endpoint](#) or an Authorization Server's [Metadata endpoint](#).

`ClientRegistrations` provides convenience methods for configuring a `ClientRegistration` in this way, as can be seen in the following example:

```
ClientRegistration clientRegistration =
    ClientRegistrations.fromIssuerLocation("https://idp.example.com/issuer").build();
```

The above code will query in series <https://idp.example.com/issuer/.well-known/openid-configuration>, and then <https://idp.example.com/.well-known/openid-configuration/issuer>, and finally <https://idp.example.com/.well-known/oauth-authorization-server/issuer>, stopping at the first to return a 200 response.

As an alternative, you can use `ClientRegistrations.fromOidcIssuerLocation()` to only query the OpenID Connect Provider's Configuration endpoint.

ClientRegistrationRepository

The `ClientRegistrationRepository` serves as a repository for OAuth 2.0 / OpenID Connect 1.0 `ClientRegistration(s)`.

Note

Client registration information is ultimately stored and owned by the associated Authorization Server. This repository provides the ability to retrieve a sub-set of the primary client registration information, which is stored with the Authorization Server.

Spring Boot 2.x auto-configuration binds each of the properties under `spring.security.oauth2.client.registration.[registrationId]` to an instance of `ClientRegistration` and then composes each of the `ClientRegistration` instance(s) within a `ClientRegistrationRepository`.

Note

The default implementation of `ClientRegistrationRepository` is `InMemoryClientRegistrationRepository`.

The auto-configuration also registers the `ClientRegistrationRepository` as a `@Bean` in the `ApplicationContext` so that it is available for dependency-injection, if needed by the application.

The following listing shows an example:

```
@Controller
public class OAuth2ClientController {

    @Autowired
    private ClientRegistrationRepository clientRegistrationRepository;

    @GetMapping("/")
    public String index() {
        ClientRegistration oktaRegistration =
            this.clientRegistrationRepository.findByRegistrationId("okta");

        ...

        return "index";
    }
}
```

OAuth2AuthorizedClient

`OAuth2AuthorizedClient` is a representation of an Authorized Client. A client is considered to be authorized when the end-user (Resource Owner) has granted authorization to the client to access its protected resources.

`OAuth2AuthorizedClient` serves the purpose of associating an `OAuth2AccessToken` (and optional `OAuth2RefreshToken`) to a `ClientRegistration` (client) and resource owner, who is the Principal end-user that granted the authorization.

OAuth2AuthorizedClientRepository / OAuth2AuthorizedClientService

`OAuth2AuthorizedClientRepository` is responsible for persisting `OAuth2AuthorizedClient(s)` between web requests. Whereas, the primary role of `OAuth2AuthorizedClientService` is to manage `OAuth2AuthorizedClient(s)` at the application-level.

From a developer perspective, the `OAuth2AuthorizedClientRepository` or `OAuth2AuthorizedClientService` provides the capability to lookup an `OAuth2AccessToken` associated with a client so that it may be used to initiate a protected resource request.

The following listing shows an example:

```
@Controller
public class OAuth2ClientController {

    @Autowired
    private OAuth2AuthorizedClientService authorizedClientService;

    @GetMapping("/")
    public String index(Authentication authentication) {
        OAuth2AuthorizedClient authorizedClient =
            this.authorizedClientService.loadAuthorizedClient("okta", authentication.getName());

        OAuth2AccessToken accessToken = authorizedClient.getAccessToken();

        ...

        return "index";
    }
}
```

Note

Spring Boot 2.x auto-configuration registers an `OAuth2AuthorizedClientRepository` and/or `OAuth2AuthorizedClientService` @Bean in the `ApplicationContext`. However, the application may choose to override and register a custom `OAuth2AuthorizedClientRepository` or `OAuth2AuthorizedClientService` @Bean.

OAuth2AuthorizedClientManager / OAuth2AuthorizedClientProvider

The `OAuth2AuthorizedClientManager` is responsible for the overall management of `OAuth2AuthorizedClient(s)`.

The primary responsibilities include:

- Authorizing (or re-authorizing) an OAuth 2.0 Client, using an `OAuth2AuthorizedClientProvider`.
- Delegating the persistence of an `OAuth2AuthorizedClient`, typically using an `OAuth2AuthorizedClientService` or `OAuth2AuthorizedClientRepository`.

An `OAuth2AuthorizedClientProvider` implements a strategy for authorizing (or re-authorizing) an OAuth 2.0 Client. Implementations will typically implement an authorization grant type, eg. `authorization_code`, `client_credentials`, etc.

The default implementation of `OAuth2AuthorizedClientManager` is `DefaultOAuth2AuthorizedClientManager`, which is associated with an `OAuth2AuthorizedClientProvider` that may support multiple authorization grant types using a delegation-based composite. The `OAuth2AuthorizedClientProviderBuilder` may be used to configure and build the delegation-based composite.

The following code shows an example of how to configure and build an `OAuth2AuthorizedClientProvider` composite that provides support for the `authorization_code`, `refresh_token`, `client_credentials` and `password` authorization grant types:

```
@Bean
public OAuth2AuthorizedClientManager authorizedClientManager(
    ClientRegistrationRepository clientRegistrationRepository,
    OAuth2AuthorizedClientRepository authorizedClientRepository) {

    OAuth2AuthorizedClientProvider authorizedClientProvider =
        OAuth2AuthorizedClientProviderBuilder.builder()
            .authorizationCode()
            .refreshToken()
            .clientCredentials()
            .password()
            .build();

    DefaultOAuth2AuthorizedClientManager authorizedClientManager =
        new DefaultOAuth2AuthorizedClientManager(
            clientRegistrationRepository, authorizedClientRepository);
    authorizedClientManager.setAuthorizedClientProvider(authorizedClientProvider);

    return authorizedClientManager;
}
```

The `DefaultOAuth2AuthorizedClientManager` is also associated with a `contextAttributesMapper` of type `Function<OAuth2AuthorizeRequest, Map<String, Object>>`, which is responsible for mapping attribute(s) from the `OAuth2AuthorizeRequest` to a `Map` of attributes to be associated to the `OAuth2AuthorizationContext`. This can be useful when you need to supply an `OAuth2AuthorizedClientProvider` with required (supported) attribute(s), eg. the `PasswordOAuth2AuthorizedClientProvider` requires the resource owner's username and password to be available in `OAuth2AuthorizationContext.getAttributes()`.

The following code shows an example of the `contextAttributesMapper`:


```

@Bean
public OAuth2AuthorizedClientManager authorizedClientManager(
    ClientRegistrationRepository clientRegistrationRepository,
    OAuth2AuthorizedClientRepository authorizedClientRepository) {

    OAuth2AuthorizedClientProvider authorizedClientProvider =
        OAuth2AuthorizedClientProviderBuilder.builder()
            .password()
            .refreshToken()
            .build();

    DefaultOAuth2AuthorizedClientManager authorizedClientManager =
        new DefaultOAuth2AuthorizedClientManager(
            clientRegistrationRepository, authorizedClientRepository);
    authorizedClientManager.setAuthorizedClientProvider(authorizedClientProvider);

    // Assuming the `username` and `password` are supplied as `HttpServletRequest` parameters,
    // map the `HttpServletRequest` parameters to `OAuth2AuthorizationContext.getAttributes()`
    authorizedClientManager.setContextAttributesMapper(contextAttributesMapper());

    return authorizedClientManager;
}

private Function<OAuth2AuthorizeRequest, Map<String, Object>> contextAttributesMapper() {
    return authorizeRequest -> {
        Map<String, Object> contextAttributes = Collections.emptyMap();
        HttpServletRequest servletRequest =
            authorizeRequest.getAttribute(HttpServletRequest.class.getName());
        String username = servletRequest.getParameter(OAuth2ParameterNames.USERNAME);
        String password = servletRequest.getParameter(OAuth2ParameterNames.PASSWORD);
        if (StringUtils.hasText(username) && StringUtils.hasText(password)) {
            contextAttributes = new HashMap<>();

            // `PasswordOAuth2AuthorizedClientProvider` requires both attributes
            contextAttributes.put(OAuth2AuthorizationContext.USERNAME_ATTRIBUTE_NAME, username);
            contextAttributes.put(OAuth2AuthorizationContext.PASSWORD_ATTRIBUTE_NAME, password);
        }
        return contextAttributes;
    };
}

```

Authorization Grant Support

Authorization Code

Note

Please refer to the OAuth 2.0 Authorization Framework for further details on the [Authorization Code](#) grant.

Obtaining Authorization

Note

Please refer to the [Authorization Request/Response](#) protocol flow for the Authorization Code grant.

Initiating the Authorization Request

The `OAuth2AuthorizationRequestRedirectFilter` uses an `OAuth2AuthorizationRequestResolver` to resolve an `OAuth2AuthorizationRequest` and initiate the Authorization Code grant flow by redirecting the end-user's user-agent to the Authorization Server's Authorization Endpoint.

The primary role of the `OAuth2AuthorizationRequestResolver` is to resolve an `OAuth2AuthorizationRequest` from the provided web request. The default implementation `DefaultOAuth2AuthorizationRequestResolver` matches on the (default) path `/oauth2/authorization/{registrationId}` extracting the `registrationId` and using it to build the `OAuth2AuthorizationRequest` for the associated `ClientRegistration`.

Given the following Spring Boot 2.x properties for an OAuth 2.0 Client registration:

```
spring:
  security:
    oauth2:
      client:
        registration:
          okta:
            client-id: okta-client-id
            client-secret: okta-client-secret
            authorization-grant-type: authorization_code
            redirect-uri: "{baseUrl}/authorized/okta"
            scope: read, write
        provider:
          okta:
            authorization-uri: https://dev-1234.oktapreview.com/oauth2/v1/authorize
            token-uri: https://dev-1234.oktapreview.com/oauth2/v1/token
```

A request with the base path `/oauth2/authorization/okta` will initiate the Authorization Request redirect by the `OAuth2AuthorizationRequestRedirectFilter` and ultimately start the Authorization Code grant flow.

Note

The `AuthorizationCodeOAuth2AuthorizedClientProvider` is an implementation of `OAuth2AuthorizedClientProvider` for the Authorization Code grant, which also initiates the Authorization Request redirect by the `OAuth2AuthorizationRequestRedirectFilter`.

If the OAuth 2.0 Client is a [Public Client](#), then configure the OAuth 2.0 Client registration as follows:

```
spring:
  security:
    oauth2:
      client:
        registration:
          okta:
            client-id: okta-client-id
            client-authentication-method: none
            authorization-grant-type: authorization_code
            redirect-uri: "{baseUrl}/authorized/okta"
            ...
```

Public Clients are supported using [Proof Key for Code Exchange](#) (PKCE). If the client is running in an untrusted environment (eg. native application or web browser-based application) and therefore incapable of maintaining the confidentiality of its credentials, PKCE will automatically be used when the following conditions are true:

1. `client-secret` is omitted (or empty)
2. `client-authentication-method` is set to "none" (`ClientAuthenticationMethod.NONE`)

The `DefaultOAuth2AuthorizationRequestResolver` also supports URI template variables for the `redirect-uri` using `UriComponentsBuilder`.

The following configuration uses all the supported URI template variables:

```
spring:
  security:
    oauth2:
      client:
        registration:
          okta:
            ...
            redirect-uri: "{baseScheme}://{baseHost}{basePort}{basePath}/authorized/{registrationId}"
            ...
```

Note

{baseUrl} resolves to {baseScheme}://{baseHost}{basePort}{basePath}

Configuring the `redirect-uri` with URI template variables is especially useful when the OAuth 2.0 Client is running behind a [Proxy Server](#). This ensures that the `X-Forwarded-*` headers are used when expanding the `redirect-uri`.

Customizing the Authorization Request

One of the primary use cases an `OAuth2AuthorizationRequestResolver` can realize is the ability to customize the Authorization Request with additional parameters above the standard parameters defined in the OAuth 2.0 Authorization Framework.

For example, OpenID Connect defines additional OAuth 2.0 request parameters for the [Authorization Code Flow](#) extending from the standard parameters defined in the [OAuth 2.0 Authorization Framework](#). One of those extended parameters is the `prompt` parameter.

Note

OPTIONAL. Space delimited, case sensitive list of ASCII string values that specifies whether the Authorization Server prompts the End-User for reauthentication and consent. The defined values are: `none`, `login`, `consent`, `select_account`

The following example shows how to implement an `OAuth2AuthorizationRequestResolver` that customizes the Authorization Request for `oauth2Login()`, by including the request parameter `prompt=consent`.

```

@EnableWebSecurity
public class OAuth2LoginSecurityConfig extends WebSecurityConfigurerAdapter {

    @Autowired
    private ClientRegistrationRepository clientRegistrationRepository;

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .authorizeRequests(authorizeRequests ->
                authorizeRequests
                    .anyRequest().authenticated()
            )
            .oauth2Login(oauth2Login ->
                oauth2Login
                    .authorizationEndpoint(authorizationEndpoint ->
                        authorizationEndpoint
                            .authorizationRequestResolver(
                                new CustomAuthorizationRequestResolver(
                                    this.clientRegistrationRepository) ❶
                            )
                    )
            );
    }
}

public class CustomAuthorizationRequestResolver implements OAuth2AuthorizationRequestResolver {
    private final OAuth2AuthorizationRequestResolver defaultAuthorizationRequestResolver;

    public CustomAuthorizationRequestResolver(
        ClientRegistrationRepository clientRegistrationRepository) {

        this.defaultAuthorizationRequestResolver =
            new DefaultOAuth2AuthorizationRequestResolver(
                clientRegistrationRepository, "/oauth2/authorization");
    }

    @Override
    public OAuth2AuthorizationRequest resolve(HttpServletRequest request) {
        OAuth2AuthorizationRequest authorizationRequest =
            this.defaultAuthorizationRequestResolver.resolve(request); ❷

        return authorizationRequest != null ? ❸
            customAuthorizationRequest(authorizationRequest) :
            null;
    }

    @Override
    public OAuth2AuthorizationRequest resolve(
        HttpServletRequest request, String clientRegistrationId) {

        OAuth2AuthorizationRequest authorizationRequest =
            this.defaultAuthorizationRequestResolver.resolve(
                request, clientRegistrationId); ❹

        return authorizationRequest != null ? ❺
            customAuthorizationRequest(authorizationRequest) :
            null;
    }

    private OAuth2AuthorizationRequest customAuthorizationRequest(
        OAuth2AuthorizationRequest authorizationRequest) {

        Map<String, Object> additionalParameters =
            new LinkedHashMap<>(authorizationRequest.getAdditionalParameters());
        additionalParameters.put("prompt", "consent"); ❻

        return OAuth2AuthorizationRequest.from(authorizationRequest) ❼
            .additionalParameters(additionalParameters) ❽
            .build();
    }
}

```

- ❶ Configure the custom `OAuth2AuthorizationRequestResolver`
- ❷ Attempt to resolve the `OAuth2AuthorizationRequest` using the `DefaultOAuth2AuthorizationRequestResolver`
- ❸ If an `OAuth2AuthorizationRequest` was resolved then return a customized version else return `null`
- ❹ Add custom parameters to the existing `OAuth2AuthorizationRequest.additionalParameters`
- ❺ Create a copy of the default `OAuth2AuthorizationRequest` which returns an `OAuth2AuthorizationRequest.Builder` for further modifications
- ❻ Override the default `additionalParameters`

Tip

`OAuth2AuthorizationRequest.Builder.build()` constructs the `OAuth2AuthorizationRequest.authorizationRequestUri`, which represents the complete Authorization Request URI including all query parameters using the `application/x-www-form-urlencoded` format.

For the simple use case, where the additional request parameter is always the same for a specific provider, it can be added directly in the `authorization-uri`.

For example, if the value for the request parameter `prompt` is always `consent` for the provider `okta`, then simply configure as follows:

```
spring:
  security:
    oauth2:
      client:
        provider:
          okta:
            authorization-uri: https://dev-1234.oktapreview.com/oauth2/v1/authorize?prompt=consent
```

The preceding example shows the common use case of adding a custom parameter on top of the standard parameters. Alternatively, if your requirements are more advanced, then you can take full control in building the Authorization Request URI by simply overriding the `OAuth2AuthorizationRequest.authorizationRequestUri` property.

The following example shows a variation of the `customAuthorizationRequest()` method from the preceding example, and instead overrides the `OAuth2AuthorizationRequest.authorizationRequestUri` property.

```
private OAuth2AuthorizationRequest customAuthorizationRequest(
    OAuth2AuthorizationRequest authorizationRequest) {

    String customAuthorizationRequestUri = UriComponentsBuilder
        .fromUriString(authorizationRequest.getAuthorizationRequestUri())
        .queryParams("prompt", "consent")
        .build(true)
        .toUriString();

    return OAuth2AuthorizationRequest.from(authorizationRequest)
        .authorizationRequestUri(customAuthorizationRequestUri)
        .build();
}
```

Storing the Authorization Request

The `AuthorizationRequestRepository` is responsible for the persistence of the `OAuth2AuthorizationRequest` from the time the Authorization Request is initiated to the time the Authorization Response is received (the callback).

Tip

The `OAuth2AuthorizationRequest` is used to correlate and validate the Authorization Response.

The default implementation of `AuthorizationRequestRepository` is `HttpSessionOAuth2AuthorizationRequestRepository`, which stores the `OAuth2AuthorizationRequest` in the `HttpSession`.

If you have a custom implementation of `AuthorizationRequestRepository`, you may configure it as shown in the following example:

```
@EnableWebSecurity
public class OAuth2ClientSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Client(oauth2Client ->
                oauth2Client
                    .authorizationCodeGrant(authorizationCodeGrant ->
                        authorizationCodeGrant
                            .authorizationRequestRepository(this.authorizationRequestRepository())
                            ...
                    )
            );
    }
}
```

Requesting an Access Token

Note

Please refer to the [Access Token Request/Response](#) protocol flow for the Authorization Code grant.

The default implementation of `OAuth2AccessTokenResponseClient` for the Authorization Code grant is `DefaultAuthorizationCodeTokenResponseClient`, which uses a `RestOperations` for exchanging an authorization code for an access token at the Authorization Server's Token Endpoint.

The `DefaultAuthorizationCodeTokenResponseClient` is quite flexible as it allows you to customize the pre-processing of the Token Request and/or post-handling of the Token Response.

Customizing the Access Token Request

If you need to customize the pre-processing of the Token Request, you can provide `DefaultAuthorizationCodeTokenResponseClient.setRequestEntityConverter()` with a custom `Converter<OAuth2AuthorizationCodeGrantRequest, RequestEntity<?>>`. The default implementation `OAuth2AuthorizationCodeGrantRequestEntityConverter` builds a `RequestEntity` representation of a standard [OAuth 2.0 Access Token Request](#). However, providing a custom `Converter`, would allow you to extend the standard Token Request and add custom parameter(s).

Important

The custom Converter must return a valid RequestEntity representation of an OAuth 2.0 Access Token Request that is understood by the intended OAuth 2.0 Provider.

Customizing the Access Token Response

On the other end, if you need to customize the post-handling of the Token Response, you will need to provide DefaultAuthorizationCodeTokenResponseClient.setRestOperations() with a custom configured RestOperations. The default RestOperations is configured as follows:

```
RestTemplate restTemplate = new RestTemplate(Arrays.asList(
    new FormHttpMessageConverter(),
    new OAuth2AccessTokenResponseHttpMessageConverter());

restTemplate.setErrorHandler(new OAuth2ErrorResponseErrorHandler());
```

Tip

Spring MVC FormHttpMessageConverter is required as it's used when sending the OAuth 2.0 Access Token Request.

OAuth2AccessTokenResponseHttpMessageConverter is a HttpMessageConverter for an OAuth 2.0 Access Token Response. You can provide OAuth2AccessTokenResponseHttpMessageConverter.setTokenResponseConverter() with a custom Converter<Map<String, String>, OAuth2AccessTokenResponse> that is used for converting the OAuth 2.0 Access Token Response parameters to an OAuth2AccessTokenResponse.

OAuth2ErrorResponseErrorHandler is a ResponseErrorHandler that can handle an OAuth 2.0 Error, eg. 400 Bad Request. It uses an OAuth2ErrorHttpMessageConverter for converting the OAuth 2.0 Error parameters to an OAuth2Error.

Whether you customize DefaultAuthorizationCodeTokenResponseClient or provide your own implementation of OAuth2AccessTokenResponseClient, you'll need to configure it as shown in the following example:

```
@EnableWebSecurity
public class OAuth2ClientSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .oauth2Client(oauth2Client ->
                oauth2Client
                    .authorizationCodeGrant(authorizationCodeGrant ->
                        authorizationCodeGrant
                            .accessTokenResponseClient(this.accessTokenResponseClient())
                            ...
                    )
            );
    }
}
```

Refresh Token**Note**

Please refer to the OAuth 2.0 Authorization Framework for further details on the [Refresh Token](#).

Refreshing an Access Token

Note

Please refer to the [Access Token Request/Response](#) protocol flow for the Refresh Token grant.

The default implementation of `OAuth2AccessTokenResponseClient` for the Refresh Token grant is `DefaultRefreshTokenTokenResponseClient`, which uses a `RestOperations` when refreshing an access token at the Authorization Server's Token Endpoint.

The `DefaultRefreshTokenTokenResponseClient` is quite flexible as it allows you to customize the pre-processing of the Token Request and/or post-handling of the Token Response.

Customizing the Access Token Request

If you need to customize the pre-processing of the Token Request, you can provide `DefaultRefreshTokenTokenResponseClient.setRequestEntityConverter()` with a custom `Converter<OAuth2RefreshTokenGrantRequest, RequestEntity<?>>`. The default implementation `OAuth2RefreshTokenGrantRequestEntityConverter` builds a `RequestEntity` representation of a standard [OAuth 2.0 Access Token Request](#). However, providing a custom `Converter`, would allow you to extend the standard Token Request and add custom parameter(s).

Important

The custom `Converter` must return a valid `RequestEntity` representation of an OAuth 2.0 Access Token Request that is understood by the intended OAuth 2.0 Provider.

Customizing the Access Token Response

On the other end, if you need to customize the post-handling of the Token Response, you will need to provide `DefaultRefreshTokenTokenResponseClient.setRestOperations()` with a custom configured `RestOperations`. The default `RestOperations` is configured as follows:

```
RestTemplate restTemplate = new RestTemplate(Arrays.asList(
    new FormHttpMessageConverter(),
    new OAuth2AccessTokenResponseHttpMessageConverter());

restTemplate.setErrorHandler(new OAuth2ErrorResponseErrorHandler());
```

Tip

Spring MVC `FormHttpMessageConverter` is required as it's used when sending the OAuth 2.0 Access Token Request.

`OAuth2AccessTokenResponseHttpMessageConverter` is a `HttpMessageConverter` for an OAuth 2.0 Access Token Response. You can provide `OAuth2AccessTokenResponseHttpMessageConverter.setTokenResponseConverter()` with a custom `Converter<Map<String, String>, OAuth2AccessTokenResponse>` that is used for converting the OAuth 2.0 Access Token Response parameters to an `OAuth2AccessTokenResponse`.

`OAuth2ErrorResponseErrorHandler` is a `ResponseErrorHandler` that can handle an OAuth 2.0 Error, eg. 400 Bad Request. It uses an `OAuth2ErrorHttpMessageConverter` for converting the OAuth 2.0 Error parameters to an `OAuth2Error`.

Whether you customize `DefaultRefreshTokenTokenResponseClient` or provide your own implementation of `OAuth2AccessTokenResponseClient`, you'll need to configure it as shown in the following example:

```
// Customize
OAuth2AccessTokenResponseClient<OAuth2RefreshTokenGrantRequest> refreshTokenTokenResponseClient = ...

OAuth2AuthorizedClientProvider authorizedClientProvider =
    OAuth2AuthorizedClientProviderBuilder.builder()
        .authorizationCode()
        .refreshToken(configurer ->
            configurer.accessTokenResponseClient(refreshTokenTokenResponseClient))
        .build();

...

authorizedClientManager.setAuthorizedClientProvider(authorizedClientProvider);
```

Note

`OAuth2AuthorizedClientProviderBuilder.builder().refreshToken()` configures a `RefreshTokenOAuth2AuthorizedClientProvider`, which is an implementation of an `OAuth2AuthorizedClientProvider` for the Refresh Token grant.

The `OAuth2RefreshToken` may optionally be returned in the Access Token Response for the `authorization_code` and `password` grant types. If the `OAuth2AuthorizedClient.getRefreshToken()` is available and the `OAuth2AuthorizedClient.getAccessToken()` is expired, it will automatically be refreshed by the `RefreshTokenOAuth2AuthorizedClientProvider`.

Client Credentials

Note

Please refer to the OAuth 2.0 Authorization Framework for further details on the [Client Credentials](#) grant.

Requesting an Access Token

Note

Please refer to the [Access Token Request/Response](#) protocol flow for the Client Credentials grant.

The default implementation of `OAuth2AccessTokenResponseClient` for the Client Credentials grant is `DefaultClientCredentialsTokenResponseClient`, which uses a `RestOperations` when requesting an access token at the Authorization Server's Token Endpoint.

The `DefaultClientCredentialsTokenResponseClient` is quite flexible as it allows you to customize the pre-processing of the Token Request and/or post-handling of the Token Response.

Customizing the Access Token Request

If you need to customize the pre-processing of the Token Request, you can provide `DefaultClientCredentialsTokenResponseClient.setRequestEntityConverter()` with a custom `Converter<OAuth2ClientCredentialsGrantRequest, RequestEntity<?>>`. The default implementation `OAuth2ClientCredentialsGrantRequestEntityConverter` builds a `RequestEntity` representation of a standard [OAuth 2.0 Access Token Request](#). However, providing a custom `Converter`, would allow you to extend the standard Token Request and add custom parameter(s).

Important

The custom `Converter` must return a valid `RequestEntity` representation of an OAuth 2.0 Access Token Request that is understood by the intended OAuth 2.0 Provider.

Customizing the Access Token Response

On the other end, if you need to customize the post-handling of the Token Response, you will need to provide `DefaultClientCredentialsTokenResponseClient.setRestOperations()` with a custom configured `RestOperations`. The default `RestOperations` is configured as follows:

```
RestTemplate restTemplate = new RestTemplate(Arrays.asList(
    new FormHttpMessageConverter(),
    new OAuth2AccessTokenResponseHttpMessageConverter());

restTemplate.setErrorHandler(new OAuth2ErrorResponseErrorHandler());
```

Tip

Spring MVC `FormHttpMessageConverter` is required as it's used when sending the OAuth 2.0 Access Token Request.

`OAuth2AccessTokenResponseHttpMessageConverter` is a `HttpMessageConverter` for an OAuth 2.0 Access Token Response. You can provide `OAuth2AccessTokenResponseHttpMessageConverter.setTokenResponseConverter()` with a custom `Converter<Map<String, String>, OAuth2AccessTokenResponse>` that is used for converting the OAuth 2.0 Access Token Response parameters to an `OAuth2AccessTokenResponse`.

`OAuth2ErrorResponseErrorHandler` is a `ResponseErrorHandler` that can handle an OAuth 2.0 Error, eg. 400 Bad Request. It uses an `OAuth2ErrorHttpMessageConverter` for converting the OAuth 2.0 Error parameters to an `OAuth2Error`.

Whether you customize `DefaultClientCredentialsTokenResponseClient` or provide your own implementation of `OAuth2AccessTokenResponseClient`, you'll need to configure it as shown in the following example:

```
// Customize
OAuth2AccessTokenResponseClient<OAuth2ClientCredentialsGrantRequest>
clientCredentialsTokenResponseClient = ...

OAuth2AuthorizedClientProvider authorizedClientProvider =
    OAuth2AuthorizedClientProviderBuilder.builder()
        .clientCredentials(configurer ->
            configurer.accessTokenResponseClient(clientCredentialsTokenResponseClient))
        .build();

...

authorizedClientManager.setAuthorizedClientProvider(authorizedClientProvider);
```

Note

`OAuth2AuthorizedClientProviderBuilder.builder().clientCredentials()` configures a `ClientCredentialsOAuth2AuthorizedClientProvider`, which is an implementation of an `OAuth2AuthorizedClientProvider` for the Client Credentials grant.

Using the Access Token

Given the following Spring Boot 2.x properties for an OAuth 2.0 Client registration:

```
spring:
  security:
    oauth2:
      client:
        registration:
          okta:
            client-id: okta-client-id
            client-secret: okta-client-secret
            authorization-grant-type: client_credentials
            scope: read, write
        provider:
          okta:
            token-uri: https://dev-1234.oktapreview.com/oauth2/v1/token
```

...and the `OAuth2AuthorizedClientManager` @Bean:

```
@Bean
public OAuth2AuthorizedClientManager authorizedClientManager(
    ClientRegistrationRepository clientRegistrationRepository,
    OAuth2AuthorizedClientRepository authorizedClientRepository) {

    OAuth2AuthorizedClientProvider authorizedClientProvider =
        OAuth2AuthorizedClientProviderBuilder.builder()
            .clientCredentials()
            .build();

    DefaultOAuth2AuthorizedClientManager authorizedClientManager =
        new DefaultOAuth2AuthorizedClientManager(
            clientRegistrationRepository, authorizedClientRepository);
    authorizedClientManager.setAuthorizedClientProvider(authorizedClientProvider);

    return authorizedClientManager;
}
```

You may obtain the `OAuth2AccessToken` as follows:

```

@Controller
public class OAuth2ClientController {

    @Autowired
    private OAuth2AuthorizedClientManager authorizedClientManager;

    @GetMapping("/")
    public String index(Authentication authentication,
        HttpServletRequest servletRequest,
        HttpServletResponse servletResponse) {

        OAuth2AuthorizeRequest authorizeRequest =
        OAuth2AuthorizeRequest.withClientRegistrationId("okta")
            .principal(authentication)
            .attributes(attrs -> {
                attrs.put(HttpServletRequest.class.getName(), servletRequest);
                attrs.put(HttpServletResponse.class.getName(), servletResponse);
            })
            .build();
        OAuth2AuthorizedClient authorizedClient
        = this.authorizedClientManager.authorize(authorizeRequest);

        OAuth2AccessToken accessToken = authorizedClient.getAccessToken();

        ...

        return "index";
    }
}

```

Note

`HttpServletRequest` and `HttpServletResponse` are both OPTIONAL attributes. If not provided, it will default to `ServletRequestAttributes` using `RequestContextHolder.getRequestAttributes()`.

Resource Owner Password Credentials**Note**

Please refer to the OAuth 2.0 Authorization Framework for further details on the [Resource Owner Password Credentials](#) grant.

Requesting an Access Token**Note**

Please refer to the [Access Token Request/Response](#) protocol flow for the Resource Owner Password Credentials grant.

The default implementation of `OAuth2AccessTokenResponseClient` for the Resource Owner Password Credentials grant is `DefaultPasswordTokenResponseClient`, which uses a `RestOperations` when requesting an access token at the Authorization Server's Token Endpoint.

The `DefaultPasswordTokenResponseClient` is quite flexible as it allows you to customize the pre-processing of the Token Request and/or post-handling of the Token Response.

Customizing the Access Token Request

If you need to customize the pre-processing of the Token Request, you can provide `DefaultPasswordTokenResponseClient.setRequestEntityConverter()` with a custom `Converter<OAuth2PasswordGrantRequest, RequestEntity<?>>`. The default implementation `OAuth2PasswordGrantRequestEntityConverter` builds a `RequestEntity` representation of a standard [OAuth 2.0 Access Token Request](#). However, providing a custom `Converter`, would allow you to extend the standard Token Request and add custom parameter(s).

Important

The custom `Converter` must return a valid `RequestEntity` representation of an OAuth 2.0 Access Token Request that is understood by the intended OAuth 2.0 Provider.

Customizing the Access Token Response

On the other end, if you need to customize the post-handling of the Token Response, you will need to provide `DefaultPasswordTokenResponseClient.setRestOperations()` with a custom configured `RestOperations`. The default `RestOperations` is configured as follows:

```
RestTemplate restTemplate = new RestTemplate(Arrays.asList(
    new FormHttpMessageConverter(),
    new OAuth2AccessTokenResponseHttpMessageConverter()));

restTemplate.setErrorHandler(new OAuth2ErrorResponseErrorHandler());
```

Tip

Spring MVC `FormHttpMessageConverter` is required as it's used when sending the OAuth 2.0 Access Token Request.

`OAuth2AccessTokenResponseHttpMessageConverter` is a `HttpMessageConverter` for an OAuth 2.0 Access Token Response. You can provide `OAuth2AccessTokenResponseHttpMessageConverter.setTokenResponseConverter()` with a custom `Converter<Map<String, String>, OAuth2AccessTokenResponse>` that is used for converting the OAuth 2.0 Access Token Response parameters to an `OAuth2AccessTokenResponse`.

`OAuth2ErrorResponseErrorHandler` is a `ResponseErrorHandler` that can handle an OAuth 2.0 Error, eg. 400 Bad Request. It uses an `OAuth2ErrorHttpMessageConverter` for converting the OAuth 2.0 Error parameters to an `OAuth2Error`.

Whether you customize `DefaultPasswordTokenResponseClient` or provide your own implementation of `OAuth2AccessTokenResponseClient`, you'll need to configure it as shown in the following example:

```
// Customize
OAuth2AccessTokenResponseClient<OAuth2PasswordGrantRequest> passwordTokenResponseClient = ...

OAuth2AuthorizedClientProvider authorizedClientProvider =
    OAuth2AuthorizedClientProviderBuilder.builder()
        .password(configurer ->
            configurer.accessTokenResponseClient(passwordTokenResponseClient))
        .refreshToken()
        .build();

...

authorizedClientManager.setAuthorizedClientProvider(authorizedClientProvider);
```

Note

`OAuth2AuthorizedClientProviderBuilder.builder().password()` configures a `PasswordOAuth2AuthorizedClientProvider`, which is an implementation of an `OAuth2AuthorizedClientProvider` for the Resource Owner Password Credentials grant.

Using the Access Token

Given the following Spring Boot 2.x properties for an OAuth 2.0 Client registration:

```
spring:
  security:
    oauth2:
      client:
        registration:
          okta:
            client-id: okta-client-id
            client-secret: okta-client-secret
            authorization-grant-type: password
            scope: read, write
        provider:
          okta:
            token-uri: https://dev-1234.oktapreview.com/oauth2/v1/token
```

...and the `OAuth2AuthorizedClientManager` @Bean:

```

@Bean
public OAuth2AuthorizedClientManager authorizedClientManager(
    ClientRegistrationRepository clientRegistrationRepository,
    OAuth2AuthorizedClientRepository authorizedClientRepository) {

    OAuth2AuthorizedClientProvider authorizedClientProvider =
        OAuth2AuthorizedClientProviderBuilder.builder()
            .password()
            .refreshToken()
            .build();

    DefaultOAuth2AuthorizedClientManager authorizedClientManager =
        new DefaultOAuth2AuthorizedClientManager(
            clientRegistrationRepository, authorizedClientRepository);
    authorizedClientManager.setAuthorizedClientProvider(authorizedClientProvider);

    // Assuming the `username` and `password` are supplied as `HttpServletRequest` parameters,
    // map the `HttpServletRequest` parameters to `OAuth2AuthorizationContext.getAttributes()`
    authorizedClientManager.setContextAttributesMapper(contextAttributesMapper());

    return authorizedClientManager;
}

private Function<OAuth2AuthorizeRequest, Map<String, Object>> contextAttributesMapper() {
    return authorizeRequest -> {
        Map<String, Object> contextAttributes = Collections.emptyMap();
        HttpServletRequest servletRequest =
            authorizeRequest.getAttribute(HttpServletRequest.class.getName());
        String username = servletRequest.getParameter(OAuth2ParameterNames.USERNAME);
        String password = servletRequest.getParameter(OAuth2ParameterNames.PASSWORD);
        if (StringUtils.hasText(username) && StringUtils.hasText(password)) {
            contextAttributes = new HashMap<>();

            // `PasswordOAuth2AuthorizedClientProvider` requires both attributes
            contextAttributes.put(OAuth2AuthorizationContext.USERNAME_ATTRIBUTE_NAME, username);
            contextAttributes.put(OAuth2AuthorizationContext.PASSWORD_ATTRIBUTE_NAME, password);
        }
        return contextAttributes;
    };
}

```

You may obtain the `OAuth2AccessToken` as follows:

```

@Controller
public class OAuth2ClientController {

    @Autowired
    private OAuth2AuthorizedClientManager authorizedClientManager;

    @GetMapping("/")
    public String index(Authentication authentication,
        HttpServletRequest servletRequest,
        HttpServletResponse servletResponse) {

        OAuth2AuthorizeRequest authorizeRequest =
        OAuth2AuthorizeRequest.withClientRegistrationId("okta")
            .principal(authentication)
            .attributes(attrs -> {
                attrs.put(HttpServletRequest.class.getName(), servletRequest);
                attrs.put(HttpServletResponse.class.getName(), servletResponse);
            })
            .build();
        OAuth2AuthorizedClient authorizedClient
        = this.authorizedClientManager.authorize(authorizeRequest);

        OAuth2AccessToken accessToken = authorizedClient.getAccessToken();

        ...

        return "index";
    }
}

```

Note

`HttpServletRequest` and `HttpServletResponse` are both **OPTIONAL** attributes. If not provided, it will default to `ServletRequestAttributes` using `RequestContextHolder.getRequestAttributes()`.

Additional Features**Resolving an Authorized Client**

The `@RegisteredOAuth2AuthorizedClient` annotation provides the capability of resolving a method parameter to an argument value of type `OAuth2AuthorizedClient`. This is a convenient alternative compared to accessing the `OAuth2AuthorizedClient` using the `OAuth2AuthorizedClientManager` or `OAuth2AuthorizedClientService`.

```

@Controller
public class OAuth2ClientController {

    @GetMapping("/")
    public String index(@RegisteredOAuth2AuthorizedClient("okta") OAuth2AuthorizedClient
        authorizedClient) {
        OAuth2AccessToken accessToken = authorizedClient.getAccessToken();

        ...

        return "index";
    }
}

```

The `@RegisteredOAuth2AuthorizedClient` annotation is handled by `OAuth2AuthorizedClientArgumentResolver`, which directly uses an [OAuth2AuthorizedClientManager](#) and therefore inherits its capabilities.

WebClient integration for Servlet Environments

The OAuth 2.0 Client support integrates with `WebClient` using an `ExchangeFilterFunction`.

The `ServletOAuth2AuthorizedClientExchangeFilterFunction` provides a simple mechanism for requesting protected resources by using an `OAuth2AuthorizedClient` and including the associated `OAuth2AccessToken` as a Bearer Token. It directly uses an [OAuth2AuthorizedClientManager](#) and therefore inherits the following capabilities:

- An `OAuth2AccessToken` will be requested if the client has not yet been authorized.
 - `authorization_code` - triggers the Authorization Request redirect to initiate the flow
 - `client_credentials` - the access token is obtained directly from the Token Endpoint
 - `password` - the access token is obtained directly from the Token Endpoint
- If the `OAuth2AccessToken` is expired, it will be refreshed (or renewed) if an `OAuth2AuthorizedClientProvider` is available to perform the authorization

The following code shows an example of how to configure `WebClient` with OAuth 2.0 Client support:

```
@Bean
WebClient webClient(OAuth2AuthorizedClientManager authorizedClientManager) {
    ServletOAuth2AuthorizedClientExchangeFilterFunction oauth2Client =
        new ServletOAuth2AuthorizedClientExchangeFilterFunction(authorizedClientManager);
    return WebClient.builder()
        .apply(oauth2Client.oauth2Configuration())
        .build();
}
```

Providing the Authorized Client

The `ServletOAuth2AuthorizedClientExchangeFilterFunction` determines the client to use (for a request) by resolving the `OAuth2AuthorizedClient` from the `ClientRequest.attributes()` (request attributes).

The following code shows how to set an `OAuth2AuthorizedClient` as a request attribute:

```
@GetMapping("/")
public String index(@RegisteredOAuth2AuthorizedClient("okta") OAuth2AuthorizedClient authorizedClient) {
    String resourceUri = ...

    String body = webClient
        .get()
        .uri(resourceUri)
        .attributes(oauth2AuthorizedClient(authorizedClient))
        .retrieve()
        .bodyToMono(String.class)
        .block();

    ...

    return "index";
}
```

- ❶ `oauth2AuthorizedClient()` is a static method in `ServletOAuth2AuthorizedClientExchangeFilterFunction`.

The following code shows how to set the `ClientRegistration.getRegistrationId()` as a request attribute:

```

@GetMapping("/")
public String index() {
    String resourceUri = ...

    String body = webClient
        .get()
        .uri(resourceUri)
        .attributes(clientRegistrationId("okta")) ❶
        .retrieve()
        .bodyToMono(String.class)
        .block();

    ...

    return "index";
}

```

- ❶ `clientRegistrationId()` is a static method in `ServletOAuth2AuthorizedClientExchangeFilterFunction`.

Defaulting the Authorized Client

If neither `OAuth2AuthorizedClient` or `ClientRegistration.getRegistrationId()` is provided as a request attribute, the `ServletOAuth2AuthorizedClientExchangeFilterFunction` can determine the *default* client to use depending on its configuration.

If `setDefaultOAuth2AuthorizedClient(true)` is configured and the user has authenticated using `HttpSecurity.oauth2Login()`, the `OAuth2AccessToken` associated with the current `OAuth2AuthenticationToken` is used.

The following code shows the specific configuration:

```

@Bean
WebClient webClient(OAuth2AuthorizedClientManager authorizedClientManager) {
    ServletOAuth2AuthorizedClientExchangeFilterFunction oauth2Client =
        new ServletOAuth2AuthorizedClientExchangeFilterFunction(authorizedClientManager);
    oauth2Client.setDefaultOAuth2AuthorizedClient(true);
    return WebClient.builder()
        .apply(oauth2Client.oauth2Configuration())
        .build();
}

```

Warning

It is recommended to be cautious with this feature since all HTTP requests will receive the access token.

Alternatively, if `setDefaultClientRegistrationId("okta")` is configured with a valid `ClientRegistration`, the `OAuth2AccessToken` associated with the `OAuth2AuthorizedClient` is used.

The following code shows the specific configuration:

```

@Bean
WebClient webClient(OAuth2AuthorizedClientManager authorizedClientManager) {
    ServletOAuth2AuthorizedClientExchangeFilterFunction oauth2Client =
        new ServletOAuth2AuthorizedClientExchangeFilterFunction(authorizedClientManager);
    oauth2Client.setDefaultClientRegistrationId("okta");
    return WebClient.builder()
        .apply(oauth2Client.oauth2Configuration())
        .build();
}

```

Warning

It is recommended to be cautious with this feature since all HTTP requests will receive the access token.

12.3 OAuth 2.0 Resource Server

Spring Security supports protecting endpoints using two forms of OAuth 2.0 [Bearer Tokens](#):

- [JWT](#)
- Opaque Tokens

This is handy in circumstances where an application has delegated its authority management to an [authorization server](#) (for example, Okta or Ping Identity). This authorization server can be consulted by resource servers to authorize requests.

Note

Working samples for both [JWTs](#) and [Opaque Tokens](#) are available in the [Spring Security repository](#).

Dependencies

Most Resource Server support is collected into `spring-security-oauth2-resource-server`. However, the support for decoding and verifying JWTs is in `spring-security-oauth2-jose`, meaning that both are necessary in order to have a working resource server that supports JWT-encoded Bearer Tokens.

Minimal Configuration for JWTs

When using [Spring Boot](#), configuring an application as a resource server consists of two basic steps. First, include the needed dependencies and second, indicate the location of the authorization server.

Specifying the Authorization Server

In a Spring Boot application, to specify which authorization server to use, simply do:

```

spring:
  security:
    oauth2:
      resourceserver:
        jwt:
          issuer-uri: https://idp.example.com/issuer

```

Where <https://idp.example.com/issuer> is the value contained in the `iss` claim for JWT tokens that the authorization server will issue. Resource Server will use this property to further self-configure, discover the authorization server's public keys, and subsequently validate incoming JWTs.

Note

To use the `issuer-uri` property, it must also be true that one of <https://idp.example.com/issuer/.well-known/openid-configuration>, <https://idp.example.com/.well-known/openid-configuration/issuer>, or <https://idp.example.com/.well-known/oauth-authorization-server/issuer> is a supported endpoint for the authorization server. This endpoint is referred to as a [Provider Configuration](#) endpoint or a [Authorization Server Metadata](#) endpoint.

And that's it!

Startup Expectations

When this property and these dependencies are used, Resource Server will automatically configure itself to validate JWT-encoded Bearer Tokens.

It achieves this through a deterministic startup process:

1. Hit the Provider Configuration or Authorization Server Metadata endpoint, processing the response for the `jwtks_url` property
2. Configure the validation strategy to query `jwtks_url` for valid public keys
3. Configure the validation strategy to validate each JWTs `iss` claim against <https://idp.example.com>.

A consequence of this process is that the authorization server must be up and receiving requests in order for Resource Server to successfully start up.

Note

If the authorization server is down when Resource Server queries it (given appropriate timeouts), then startup will fail.

Runtime Expectations

Once the application is started up, Resource Server will attempt to process any request containing an `Authorization: Bearer` header:

```
GET / HTTP/1.1
Authorization: Bearer some-token-value # Resource Server will process this
```

So long as this scheme is indicated, Resource Server will attempt to process the request according to the Bearer Token specification.

Given a well-formed JWT, Resource Server will:

1. Validate its signature against a public key obtained from the `jwtks_url` endpoint during startup and matched against the JWTs header
2. Validate the JWTs `exp` and `nbf` timestamps and the JWTs `iss` claim, and

3. Map each scope to an authority with the prefix `SCOPE_`.

Note

As the authorization server makes available new keys, Spring Security will automatically rotate the keys used to validate the JWT tokens.

The resulting `Authentication#getPrincipal`, by default, is a Spring Security `Jwt` object, and `Authentication#getName` maps to the JWT's `sub` property, if one is present.

From here, consider jumping to:

[How to Configure without Tying Resource Server startup to an authorization server's availability](#)

[How to Configure without Spring Boot](#)

Specifying the Authorization Server JWK Set Uri Directly

If the authorization server doesn't support any configuration endpoints, or if Resource Server must be able to start up independently from the authorization server, then the `jwt-set-uri` can be supplied as well:

```
spring:
  security:
    oauth2:
      resourceserver:
        jwt:
          issuer-uri: https://idp.example.com
          jwt-set-uri: https://idp.example.com/.well-known/jwks.json
```

Note

The JWK Set uri is not standardized, but can typically be found in the authorization server's documentation

Consequently, Resource Server will not ping the authorization server at startup. We still specify the `issuer-uri` so that Resource Server still validates the `iss` claim on incoming JWTs.

Note

This property can also be supplied directly on the [DSL](#).

Overriding or Replacing Boot Auto Configuration

There are two `@Beans` that Spring Boot generates on Resource Server's behalf.

The first is a `WebSecurityConfigurerAdapter` that configures the app as a resource server. When including `spring-security-oauth2-jose`, this `WebSecurityConfigurerAdapter` looks like:

```
protected void configure(HttpSecurity http) {
    http
        .authorizeRequests()
            .anyRequest().authenticated()
            .and()
            .oauth2ResourceServer(OAuth2ResourceServerConfigurer::jwt)
}
```

If the application doesn't expose a `WebSecurityConfigurerAdapter` bean, then Spring Boot will expose the above default one.

Replacing this is as simple as exposing the bean within the application:

```
@EnableWebSecurity
public class MyCustomSecurityConfiguration extends WebSecurityConfigurerAdapter {
    protected void configure(HttpSecurity http) {
        http
            .authorizeRequests()
                .mvcMatchers("/messages/**").hasAuthority("SCOPE_message:read")
                .anyRequest().authenticated()
                .and()
            .oauth2ResourceServer()
                .jwt()
                    .jwtAuthenticationConverter(myConverter());
    }
}
```

The above requires the scope of `message:read` for any URL that starts with `/messages/`.

Methods on the `oauth2ResourceServer` DSL will also override or replace auto configuration.

For example, the second `@Bean` Spring Boot creates is a `JwtDecoder`, which decodes `String` tokens into validated instances of `Jwt`:

```
@Bean
public JwtDecoder jwtDecoder() {
    return JwtDecoders.fromIssuerLocation(issuerUri);
}
```

Note

Calling [JwtDecoders#fromIssuerLocation](#) is what invokes the Provider Configuration or Authorization Server Metadata endpoint in order to derive the JWK Set Uri.

If the application doesn't expose a `JwtDecoder` bean, then Spring Boot will expose the above default one.

And its configuration can be overridden using `jwtSetUri()` or replaced using `decoder()`.

Using `jwtSetUri()`

An authorization server's JWK Set Uri can be configured [as a configuration property](#) or it can be supplied in the DSL:

```
@EnableWebSecurity
public class DirectlyConfiguredJwtSetUri extends WebSecurityConfigurerAdapter {
    protected void configure(HttpSecurity http) {
        http
            .authorizeRequests()
                .anyRequest().authenticated()
                .and()
            .oauth2ResourceServer()
                .jwt()
                    .jwtSetUri("https://idp.example.com/.well-known/jwks.json");
    }
}
```

Using `jwtSetUri()` takes precedence over any configuration property.

Using `decoder()`

More powerful than `jwtSetUri()` is `decoder()`, which will completely replace any Boot auto configuration of `JwtDecoder`:

```
@EnableWebSecurity
public class DirectlyConfiguredJwtDecoder extends WebSecurityConfigurerAdapter {
    protected void configure(HttpSecurity http) {
        http
            .authorizeRequests()
                .anyRequest().authenticated()
                .and()
            .oauth2ResourceServer()
                .jwt()
                    .decoder(myCustomDecoder());
    }
}
```

This is handy when deeper configuration, like [validation](#), [mapping](#), or [request timeouts](#), is necessary.

Exposing a `JwtDecoder` @Bean

Or, exposing a `JwtDecoder` @Bean has the same effect as `decoder()`:

```
@Bean
public JwtDecoder jwtDecoder() {
    return NimbusJwtDecoder.withJwkSetUri(jwkSetUri).build();
}
```

Configuring Trusted Algorithms

By default, `NimbusJwtDecoder`, and hence Resource Server, will only trust and verify tokens using RS256.

You can customize this via [Spring Boot, the NimbusJwtDecoder builder](#), or from the [JWK Set response](#).

Via Spring Boot

The simplest way to set the algorithm is as a property:

```
spring:
  security:
    oauth2:
      resourceserver:
        jwt:
          jws-algorithm: RS512
          jwk-set-uri: https://idp.example.org/.well-known/jwks.json
```

Using a Builder

For greater power, though, we can use a builder that ships with `NimbusJwtDecoder`:

```
@Bean
JwtDecoder jwtDecoder() {
    return NimbusJwtDecoder.fromJwkSetUri(this.jwkSetUri)
        .jwsAlgorithm(RS512).build();
}
```

Calling `jwsAlgorithm` more than once will configure `NimbusJwtDecoder` to trust more than one algorithm, like so:

```
@Bean
JwtDecoder jwtDecoder() {
    return NimbusJwtDecoder.fromJwkSetUri(this.jwkSetUri)
        .jwsAlgorithm(RS512).jwsAlgorithm(EC512).build();
}
```

Or, you can call `jwsAlgorithms`:

```
@Bean
JwtDecoder jwtDecoder() {
    return NimbusJwtDecoder.fromJwkSetUri(this.jwkSetUri)
        .jwsAlgorithms(algorithms -> {
            algorithms.add(RS512);
            algorithms.add(EC512);
        }).build();
}
```

From JWK Set response

Since Spring Security's JWT support is based off of Nimbus, you can use all its great features as well.

For example, Nimbus has a `JWSKeySelector` implementation that will select the set of algorithms based on the JWK Set URI response. You can use it to generate a `NimbusJwtDecoder` like so:

```
@Bean
public JwtDecoder jwtDecoder() {
    // makes a request to the JWK Set endpoint
    JWSKeySelector<SecurityContext> jwsKeySelector =
        JWSAlgorithmFamilyJWSKeySelector.fromJWKSetURL(this.jwkSetUrl);

    DefaultJWTProcessor<SecurityContext> jwtProcessor =
        new DefaultJWTProcessor<>();
    jwtProcessor.setJWSKeySelector(jwsKeySelector);

    return new NimbusJwtDecoder(jwtProcessor);
}
```

Trusting a Single Asymmetric Key

Simpler than backing a Resource Server with a JWK Set endpoint is to hard-code an RSA public key. The public key can be provided via [Spring Boot](#) or by [Using a Builder](#).

Via Spring Boot

Specifying a key via Spring Boot is quite simple. The key's location can be specified like so:

```
spring:
  security:
    oauth2:
      resourceserver:
        jwt:
          public-key-location: classpath:my-key.pub
```

Or, to allow for a more sophisticated lookup, you can post-process the `RsaKeyConversionServicePostProcessor`:

```
@Bean
BeanFactoryPostProcessor conversionServiceCustomizer() {
    return beanFactory ->
        beanFactory.getBean(RsaKeyConversionServicePostProcessor.class)
            .setResourceLoader(new CustomResourceLoader());
}
```


Specify your key's location:

```
key.location: hdfs://my-key.pub
```

And then autowire the value:

```
@Value("${key.location}")
RSAPublicKey key;
```

Using a Builder

To wire an `RSAPublicKey` directly, you can simply use the appropriate `NimbusJwtDecoder` builder, like so:

```
@Bean
public JwtDecoder jwtDecoder() {
    return NimbusJwtDecoder.withPublicKey(this.key).build();
}
```

Trusting a Single Symmetric Key

Using a single symmetric key is also simple. You can simply load in your `SecretKey` and use the appropriate `NimbusJwtDecoder` builder, like so:

```
@Bean
public JwtDecoder jwtDecoder() {
    return NimbusJwtDecoder.withSecretKey(this.key).build();
}
```

Configuring Authorization

A JWT that is issued from an OAuth 2.0 Authorization Server will typically either have a `scope` or `scp` attribute, indicating the scopes (or authorities) it's been granted, for example:

```
{ ..., "scope" : "messages contacts" }
```

When this is the case, Resource Server will attempt to coerce these scopes into a list of granted authorities, prefixing each scope with the string `"SCOPE_"`.

This means that to protect an endpoint or method with a scope derived from a JWT, the corresponding expressions should include this prefix:

```
@EnableWebSecurity
public class DirectlyConfiguredJwkSetUri extends WebSecurityConfigurerAdapter {
    protected void configure(HttpSecurity http) {
        http
            .authorizeRequests(authorizeRequests -> authorizeRequests
                .mvcMatchers("/contacts/**").hasAuthority("SCOPE_contacts")
                .mvcMatchers("/messages/**").hasAuthority("SCOPE_messages")
                .anyRequest().authenticated()
            )
            .oauth2ResourceServer(OAuth2ResourceServerConfigurer::jwt);
    }
}
```

Or similarly with method security:

```
@PreAuthorize("hasAuthority('SCOPE_messages')")
public List<Message> getMessages(...) {}
```

Extracting Authorities Manually

However, there are a number of circumstances where this default is insufficient. For example, some authorization servers don't use the `scope` attribute, but instead have their own custom attribute. Or, at other times, the resource server may need to adapt the attribute or a composition of attributes into internalized authorities.

To this end, the DSL exposes `jwtAuthenticationConverter()`:

```
@EnableWebSecurity
public class DirectlyConfiguredJwkSetUri extends WebSecurityConfigurerAdapter {
    protected void configure(HttpSecurity http) {
        http
            .authorizeRequests()
                .anyRequest().authenticated()
                .and()
            .oauth2ResourceServer()
                .jwt()
                    .jwtAuthenticationConverter(grantedAuthoritiesExtractor());
    }
}

Converter<Jwt, AbstractAuthenticationToken> grantedAuthoritiesExtractor() {
    JwtAuthenticationConverter jwtAuthenticationConverter =
        new JwtAuthenticationConverter();
    jwtAuthenticationConverter.setJwtGrantedAuthoritiesConverter
        (new GrantedAuthoritiesExtractor());
    return jwtAuthenticationConverter;
}
```

which is responsible for converting a `Jwt` into an `Authentication`. As part of its configuration, we can supply a subsidiary converter to go from `Jwt` to a `Collection` of granted authorities.

That final converter might be something like `GrantedAuthoritiesExtractor` below:

```
static class GrantedAuthoritiesExtractor
    implements Converter<Jwt, Collection<GrantedAuthority>> {

    public Collection<GrantedAuthority> convert(Jwt jwt) {
        Collection<String> authorities = (Collection<String>)
            jwt.getClaims().get("mycustomclaim");

        return authorities.stream()
            .map(SimpleGrantedAuthority::new)
            .collect(Collectors.toList());
    }
}
```

For more flexibility, the DSL supports entirely replacing the converter with any class that implements `Converter<Jwt, AbstractAuthenticationToken>`:

```
static class CustomAuthenticationConverter implements Converter<Jwt, AbstractAuthenticationToken> {
    public AbstractAuthenticationToken convert(Jwt jwt) {
        return new CustomAuthenticationToken(jwt);
    }
}
```

Configuring Validation

Using [minimal Spring Boot configuration](#), indicating the authorization server's issuer uri, Resource Server will default to verifying the `iss` claim as well as the `exp` and `nbf` timestamp claims.

In circumstances where validation needs to be customized, Resource Server ships with two standard validators and also accepts custom `OAuth2TokenValidator` instances.

Customizing Timestamp Validation

JWT's typically have a window of validity, with the start of the window indicated in the `nbf` claim and the end indicated in the `exp` claim.

However, every server can experience clock drift, which can cause tokens to appear expired to one server, but not to another. This can cause some implementation heartburn as the number of collaborating servers increases in a distributed system.

Resource Server uses `JwtTimestampValidator` to verify a token's validity window, and it can be configured with a `clockSkew` to alleviate the above problem:

```
@Bean
JwtDecoder jwtDecoder() {
    NimbusJwtDecoder jwtDecoder = (NimbusJwtDecoder)
        JwtDecoders.fromIssuerLocation(issuerUri);

    OAuth2TokenValidator<Jwt> withClockSkew = new DelegatingOAuth2TokenValidator<>(
        new JwtTimestampValidator(Duration.ofSeconds(60)),
        new IssuerValidator(issuerUri));

    jwtDecoder.setJwtValidator(withClockSkew);

    return jwtDecoder;
}
```

Note

By default, Resource Server configures a clock skew of 30 seconds.

Configuring a Custom Validator

Adding a check for the `aud` claim is simple with the `OAuth2TokenValidator` API:

```
public class AudienceValidator implements OAuth2TokenValidator<Jwt> {
    OAuth2Error error = new OAuth2Error("invalid_token", "The required audience is missing", null);

    public OAuth2TokenValidatorResult validate(Jwt jwt) {
        if (jwt.getAudience().contains("messaging")) {
            return OAuth2TokenValidatorResult.success();
        } else {
            return OAuth2TokenValidatorResult.failure(error);
        }
    }
}
```

Then, to add into a resource server, it's a matter of specifying the `JwtDecoder` instance:

```
@Bean
JwtDecoder jwtDecoder() {
    NimbusJwtDecoder jwtDecoder = (NimbusJwtDecoder)
        JwtDecoders.fromIssuerLocation(issuerUri);

    OAuth2TokenValidator<Jwt> audienceValidator = new AudienceValidator();
    OAuth2TokenValidator<Jwt> withIssuer = JwtValidators.createDefaultWithIssuer(issuerUri);
    OAuth2TokenValidator<Jwt> withAudience = new DelegatingOAuth2TokenValidator<>(withIssuer,
        audienceValidator);

    jwtDecoder.setJwtValidator(withAudience);

    return jwtDecoder;
}
```

Configuring Claim Set Mapping

Spring Security uses the [Nimbus](#) library for parsing JWTs and validating their signatures. Consequently, Spring Security is subject to Nimbus's interpretation of each field value and how to coerce each into a Java type.

For example, because Nimbus remains Java 7 compatible, it doesn't use `Instant` to represent timestamp fields.

And it's entirely possible to use a different library or for JWT processing, which may make its own coercion decisions that need adjustment.

Or, quite simply, a resource server may want to add or remove claims from a JWT for domain-specific reasons.

For these purposes, Resource Server supports mapping the JWT claim set with `MappedJwtClaimSetConverter`.

Customizing the Conversion of a Single Claim

By default, `MappedJwtClaimSetConverter` will attempt to coerce claims into the following types:

Claim	Java Type
aud	<code>Collection<String></code>
exp	<code>Instant</code>
iat	<code>Instant</code>
iss	<code>String</code>
jti	<code>String</code>
nbf	<code>Instant</code>
sub	<code>String</code>

An individual claim's conversion strategy can be configured using `MappedJwtClaimSetConverter.withDefaults()`:

```
@Bean
JwtDecoder jwtDecoder() {
    NimbusJwtDecoder jwtDecoder = NimbusJwtDecoder.withJwkSetUri(jwkSetUri).build();

    MappedJwtClaimSetConverter converter = MappedJwtClaimSetConverter
        .withDefaults(Collections.singletonMap("sub", this::lookupUserIdBySub));
    jwtDecoder.setClaimSetConverter(converter);

    return jwtDecoder;
}
```

This will keep all the defaults, except it will override the default claim converter for `sub`.

Adding a Claim

`MappedJwtClaimSetConverter` can also be used to add a custom claim, for example, to adapt to an existing system:

```
MappedJwtClaimSetConverter.withDefaults(Collections.singletonMap("custom", custom -> "value"));
```

Removing a Claim

And removing a claim is also simple, using the same API:

```
MappedJwtClaimSetConverter.withDefaults(Collections.singletonMap("legacyclaim", legacy -> null));
```

Renaming a Claim

In more sophisticated scenarios, like consulting multiple claims at once or renaming a claim, Resource Server accepts any class that implements `Converter<Map<String, Object>, Map<String, Object>>`:

```
public class UsernameSubClaimAdapter implements Converter<Map<String, Object>, Map<String, Object>> {
    private final MappedJwtClaimSetConverter delegate =
        MappedJwtClaimSetConverter.withDefaults(Collections.emptyMap());

    public Map<String, Object> convert(Map<String, Object> claims) {
        Map<String, Object> convertedClaims = this.delegate.convert(claims);

        String username = (String) convertedClaims.get("user_name");
        convertedClaims.put("sub", username);

        return convertedClaims;
    }
}
```

And then, the instance can be supplied like normal:

```
@Bean
JwtDecoder jwtDecoder() {
    NimbusJwtDecoder jwtDecoder = NimbusJwtDecoder.withJwkSetUri(jwkSetUri).build();
    jwtDecoder.setClaimSetConverter(new UsernameSubClaimAdapter());
    return jwtDecoder;
}
```

Configuring Timeouts

By default, Resource Server uses connection and socket timeouts of 30 seconds each for coordinating with the authorization server.

This may be too short in some scenarios. Further, it doesn't take into account more sophisticated patterns like back-off and discovery.

To adjust the way in which Resource Server connects to the authorization server, `NimbusJwtDecoder` accepts an instance of `RestOperations`:

```
@Bean
public JwtDecoder jwtDecoder(RestTemplateBuilder builder) {
    RestOperations rest = builder
        .setConnectionTimeout(60000)
        .setReadTimeout(60000)
        .build();

    NimbusJwtDecoder jwtDecoder =
        NimbusJwtDecoder.withJwkSetUri(jwkSetUri).restOperations(rest).build();
    return jwtDecoder;
}
```

Minimal Configuration for Introspection

Typically, an opaque token can be verified via an [OAuth 2.0 Introspection Endpoint](#), hosted by the authorization server. This can be handy when revocation is a requirement.

When using [Spring Boot](#), configuring an application as a resource server that uses introspection consists of two basic steps. First, include the needed dependencies and second, indicate the introspection endpoint details.

Specifying the Authorization Server

To specify where the introspection endpoint is, simply do:

```
security:
  oauth2:
    resourceserver:
      opaque-token:
        introspection-uri: https://idp.example.com/introspect
        client-id: client
        client-secret: secret
```

Where <https://idp.example.com/introspect> is the introspection endpoint hosted by your authorization server and `client-id` and `client-secret` are the credentials needed to hit that endpoint.

Resource Server will use these properties to further self-configure and subsequently validate incoming JWTs.

Note

When using introspection, the authorization server's word is the law. If the authorization server responds that the token is valid, then it is.

And that's it!

Startup Expectations

When this property and these dependencies are used, Resource Server will automatically configure itself to validate Opaque Bearer Tokens.

This startup process is quite a bit simpler than for JWTs since no endpoints need to be discovered and no additional validation rules get added.

Runtime Expectations

Once the application is started up, Resource Server will attempt to process any request containing an `Authorization: Bearer` header:

```
GET / HTTP/1.1
Authorization: Bearer some-token-value # Resource Server will process this
```

So long as this scheme is indicated, Resource Server will attempt to process the request according to the Bearer Token specification.

Given an Opaque Token, Resource Server will

1. Query the provided introspection endpoint using the provided credentials and the token
2. Inspect the response for an `{ 'active' : true }` attribute
3. Map each scope to an authority with the prefix `SCOPE_`

The resulting `Authentication#getPrincipal`, by default, is a Spring Security [OAuth2AuthenticatedPrincipal](#) object, and `Authentication#getName` maps to the token's sub property, if one is present.

From here, you may want to jump to:

- [Looking Up Attributes Post-Authentication](#)
- [Extracting Authorities Manually](#)
- [Using Introspection with JWTs](#)

Looking Up Attributes Post-Authentication

Once a token is authenticated, an instance of `BearerTokenAuthentication` is set in the `SecurityContext`.

This means that it's available in `@Controller` methods when using `@EnableWebMvc` in your configuration:

```
@GetMapping("/foo")
public String foo(BearerTokenAuthentication authentication) {
    return authentication.getTokenAttributes().get("sub") + " is the subject";
}
```

Since `BearerTokenAuthentication` holds an `OAuth2AuthenticatedPrincipal`, that also means that it's available to controller methods, too:

```
@GetMapping("/foo")
public String foo(@AuthenticationPrincipal OAuth2AuthenticatedPrincipal principal) {
    return principal.getAttribute("sub") + " is the subject";
}
```

Looking Up Attributes Via SpEL

Of course, this also means that attributes can be accessed via SpEL.

For example, if using `@EnableGlobalMethodSecurity` so that you can use `@PreAuthorize` annotations, you can do:

```
@PreAuthorize("principal?.attributes['sub'] == 'foo'")
public String forFoosEyesOnly() {
    return "foo";
}
```

Overriding or Replacing Boot Auto Configuration

There are two `@Beans` that Spring Boot generates on Resource Server's behalf.

The first is a `WebSecurityConfigurerAdapter` that configures the app as a resource server. When use Opaque Token, this `WebSecurityConfigurerAdapter` looks like:

```
protected void configure(HttpSecurity http) {
    http
        .authorizeRequests()
            .anyRequest().authenticated()
            .and()
            .oauth2ResourceServer(OAuth2ResourceServerConfigurer::opaqueToken)
}
```

If the application doesn't expose a `WebSecurityConfigurerAdapter` bean, then Spring Boot will expose the above default one.

Replacing this is as simple as exposing the bean within the application:

```
@EnableWebSecurity
public class MyCustomSecurityConfiguration extends WebSecurityConfigurerAdapter {
    protected void configure(HttpSecurity http) {
        http
            .authorizeRequests()
                .mvcMatchers("/messages/**").hasAuthority("SCOPE_message:read")
                .anyRequest().authenticated()
                .and()
            .oauth2ResourceServer()
                .opaqueToken()
                    .introspector(myIntrospector());
    }
}
```

The above requires the scope of `message:read` for any URL that starts with `/messages/`.

Methods on the `oauth2ResourceServer` DSL will also override or replace auto configuration.

For example, the second `@Bean` Spring Boot creates is an `OpaqueTokenIntrospector`, which decodes String tokens into validated instances of `OAuth2AuthenticatedPrincipal`:

```
@Bean
public OpaqueTokenIntrospector introspector() {
    return new NimbusOpaqueTokenIntrospector(introspectionUri, clientId, clientSecret);
}
```

If the application doesn't expose a `OpaqueTokenIntrospector` bean, then Spring Boot will expose the above default one.

And its configuration can be overridden using `introspectionUri()` and `introspectionClientCredentials()` or replaced using `introspector()`.

Using `introspectionUri()`

An authorization server's Introspection Uri can be configured [as a configuration property](#) or it can be supplied in the DSL:

```
@EnableWebSecurity
public class DirectlyConfiguredIntrospectionUri extends WebSecurityConfigurerAdapter {
    protected void configure(HttpSecurity http) {
        http
            .authorizeRequests()
                .anyRequest().authenticated()
                .and()
            .oauth2ResourceServer()
                .opaqueToken()
                    .introspectionUri("https://idp.example.com/introspect")
                    .introspectionClientCredentials("client", "secret");
    }
}
```

Using `introspectionUri()` takes precedence over any configuration property.

Using `introspector()`

More powerful than `introspectionUri()` is `introspector()`, which will completely replace any Boot auto configuration of `OpaqueTokenIntrospector`:


```

@EnableWebSecurity
public class DirectlyConfiguredIntrospector extends WebSecurityConfigurerAdapter {
    protected void configure(HttpSecurity http) {
        http
            .authorizeRequests()
                .anyRequest().authenticated()
                .and()
            .oauth2ResourceServer()
                .opaqueToken()
                    .introspector(myCustomIntrospector());
    }
}

```

This is handy when deeper configuration, like [authority mapping](#), [JWT revocation](#), or [request timeouts](#), is necessary.

Exposing a OpaqueTokenIntrospector @Bean

Or, exposing a OpaqueTokenIntrospector @Bean has the same effect as introspector():

```

@Bean
public OpaqueTokenIntrospector introspector() {
    return new NimbusOpaqueTokenIntrospector(introspectionUri, clientId, clientSecret);
}

```

Configuring Authorization

An OAuth 2.0 Introspection endpoint will typically return a `scope` attribute, indicating the scopes (or authorities) it's been granted, for example:

```
{ ..., "scope" : "messages contacts" }
```

When this is the case, Resource Server will attempt to coerce these scopes into a list of granted authorities, prefixing each scope with the string "SCOPE_".

This means that to protect an endpoint or method with a scope derived from an Opaque Token, the corresponding expressions should include this prefix:

```

@EnableWebSecurity
public class MappedAuthorities extends WebSecurityConfigurerAdapter {
    protected void configure(HttpSecurity http) {
        http
            .authorizeRequests(authorizeRequests -> authorizeRequests
                .mvcMatchers("/contacts/**").hasAuthority("SCOPE_contacts")
                .mvcMatchers("/messages/**").hasAuthority("SCOPE_messages")
                .anyRequest().authenticated()
            )
            .oauth2ResourceServer(OAuth2ResourceServerConfigurer::opaqueToken);
    }
}

```

Or similarly with method security:

```

@PreAuthorize("hasAuthority('SCOPE_messages')")
public List<Message> getMessages(...) {}

```

Extracting Authorities Manually

By default, Opaque Token support will extract the scope claim from an introspection response and parse it into individual `GrantedAuthority` instances.

For example, if the introspection response were:

```
{
  "active" : true,
  "scope" : "message:read message:write"
}
```

Then Resource Server would generate an Authentication with two authorities, one for `message:read` and the other for `message:write`.

This can, of course, be customized using a custom `OpaqueTokenIntrospector` that takes a look at the attribute set and converts in its own way:

```
public class CustomAuthoritiesOpaqueTokenIntrospector implements OpaqueTokenIntrospector {
    private OpaqueTokenIntrospector delegate =
        new NimbusOpaqueTokenIntrospector("https://idp.example.org/introspect", "client", "secret");

    public OAuth2AuthenticatedPrincipal introspect(String token) {
        OAuth2AuthenticatedPrincipal principal = this.delegate.introspect(token);
        return new DefaultOAuth2AuthenticatedPrincipal(
            principal.getName(), principal.getAttributes(), extractAuthorities(principal));
    }

    private Collection<GrantedAuthority> extractAuthorities(OAuth2AuthenticatedPrincipal principal) {
        List<String> scopes = principal.getAttribute(OAuth2IntrospectionClaimNames.SCOPE);
        return scopes.stream()
            .map(SimpleGrantedAuthority::new)
            .collect(Collectors.toList());
    }
}
```

Thereafter, this custom introspector can be configured simply by exposing it as a `@Bean`:

```
@Bean
public OpaqueTokenIntrospector introspector() {
    return new CustomAuthoritiesOpaqueTokenIntrospector();
}
```

Configuring Timeouts

By default, Resource Server uses connection and socket timeouts of 30 seconds each for coordinating with the authorization server.

This may be too short in some scenarios. Further, it doesn't take into account more sophisticated patterns like back-off and discovery.

To adjust the way in which Resource Server connects to the authorization server, `NimbusOpaqueTokenIntrospector` accepts an instance of `RestOperations`:

```
@Bean
public OpaqueTokenIntrospector introspector(RestTemplateBuilder builder) {
    RestOperations rest = builder
        .basicAuthentication(clientId, clientSecret)
        .setConnectionTimeout(60000)
        .setReadTimeout(60000)
        .build();

    return new NimbusOpaqueTokenIntrospector(introspectionUri, rest);
}
```

Using Introspection with JWTs

A common question is whether or not introspection is compatible with JWTs. Spring Security's Opaque Token support has been designed to not care about the format of the token — it will gladly pass any token to the introspection endpoint provided.

So, let's say that you've got a requirement that requires you to check with the authorization server on each request, in case the JWT has been revoked.

Even though you are using the JWT format for the token, your validation method is introspection, meaning you'd want to do:

```
spring:
  security:
    oauth2:
      resourceserver:
        opaque-token:
          introspection-uri: https://idp.example.org/introspection
          client-id: client
          client-secret: secret
```

In this case, the resulting `Authentication` would be `BearerTokenAuthentication`. Any attributes in the corresponding `OAuth2AuthenticatedPrincipal` would be whatever was returned by the introspection endpoint.

But, let's say that, oddly enough, the introspection endpoint only returns whether or not the token is active. Now what?

In this case, you can create a custom `OpaqueTokenIntrospector` that still hits the endpoint, but then updates the returned principal to have the JWT's claims as the attributes:

```
public class JwtOpaqueTokenIntrospector implements OpaqueTokenIntrospector {
    private OpaqueTokenIntrospector delegate =
        new NimbusOpaqueTokenIntrospector("https://idp.example.org/introspect", "client", "secret");
    private JwtDecoder jwtDecoder = new NimbusJwtDecoder(new ParseOnlyJWTProcessor());

    public OAuth2AuthenticatedPrincipal introspect(String token) {
        OAuth2AuthenticatedPrincipal principal = this.delegate.introspect(token);
        try {
            Jwt jwt = this.jwtDecoder.decode(token);
            return new DefaultOAuth2AuthenticatedPrincipal(jwt.getClaims(), NO_AUTHORITIES);
        } catch (JwtException e) {
            throw new OAuth2IntrospectionException(e);
        }
    }

    private static class ParseOnlyJWTProcessor extends DefaultJWTProcessor<SecurityContext> {
        JWTClaimsSet process(SignedJWT jwt, SecurityContext context)
            throws JOSEException {
            return jwt.getJWTClaimSet();
        }
    }
}
```

Thereafter, this custom introspector can be configured simply by exposing it as a `@Bean`:

```
@Bean
public OpaqueTokenIntrospector introspector() {
    return new JwtOpaqueTokenIntrospector();
}
```

Calling a `/userinfo` Endpoint

Generally speaking, a Resource Server doesn't care about the underlying user, but instead about the authorities that have been granted.

That said, at times it can be valuable to tie the authorization statement back to a user.

If an application is also using `spring-security-oauth2-client`, having set up the appropriate `ClientRegistrationRepository`, then this is quite simple with a custom `OpaqueTokenIntrospector`. This implementation below does three things:

- Delegates to the introspection endpoint, to affirm the token's validity
- Looks up the appropriate client registration associated with the `/userinfo` endpoint
- Invokes and returns the response from the `/userinfo` endpoint

```
public class UserInfoOpaqueTokenIntrospector implements OpaqueTokenIntrospector {
    private final OpaqueTokenIntrospector delegate =
        new NimbusOpaqueTokenIntrospector("https://idp.example.org/introspect", "client", "secret");
    private final OAuth2UserService oauth2UserService = new DefaultOAuth2UserService();

    private final ClientRegistrationRepository repository;

    // ... constructor

    @Override
    public OAuth2AuthenticatedPrincipal introspect(String token) {
        OAuth2AuthenticatedPrincipal authorized = this.delegate.introspect(token);
        Instant issuedAt = authorized.getAttribute(ISSUED_AT);
        Instant expiresAt = authorized.getAttribute(EXPIRES_AT);
        ClientRegistration clientRegistration = this.repository.findByRegistrationId("registration-id");
        OAuth2AccessToken token = new OAuth2AccessToken(BEARER, token, issuedAt, expiresAt);
        OAuth2UserRequest oauth2UserRequest = new OAuth2UserRequest(clientRegistration, token);
        return this.oauth2UserService.loadUser(oauth2UserRequest);
    }
}
```

If you aren't using `spring-security-oauth2-client`, it's still quite simple. You will simply need to invoke the `/userinfo` with your own instance of `WebClient`:

```
public class UserInfoOpaqueTokenIntrospector implements OpaqueTokenIntrospector {
    private final OpaqueTokenIntrospector delegate =
        new NimbusOpaqueTokenIntrospector("https://idp.example.org/introspect", "client", "secret");
    private final WebClient rest = WebClient.create();

    @Override
    public OAuth2AuthenticatedPrincipal introspect(String token) {
        OAuth2AuthenticatedPrincipal authorized = this.delegate.introspect(token);
        return makeUserInfoRequest(authorized);
    }
}
```

Either way, having created your `OpaqueTokenIntrospector`, you should publish it as a `@Bean` to override the defaults:

```
@Bean
OpaqueTokenIntrospector introspector() {
    return new UserInfoOpaqueTokenIntrospector(...);
}
```

Supporting both JWT and Opaque Token

In some cases, you may have a need to access both kinds of tokens. For example, you may support more than one tenant where one tenant issues JWTs and the other issues opaque tokens.

If this decision must be made at request-time, then you can use an `AuthenticationManagerResolver` to achieve it, like so:

```

@Bean
AuthenticationManagerResolver<HttpServletRequest> tokenAuthenticationManagerResolver() {
    BearerTokenResolver bearerToken = new DefaultBearerTokenResolver();
    JwtAuthenticationProvider jwt = jwt();
    OpaqueTokenAuthenticationProvider opaqueToken = opaqueToken();

    return request -> {
        String token = bearerToken.resolve(request);
        if (isAJwt(token)) {
            return jwt::authenticate;
        } else {
            return opaqueToken::authenticate;
        }
    }
}

```

And then specify this `AuthenticationManagerResolver` in the DSL:

```

http
    .authorizeRequests()
        .anyRequest().authenticated()
        .and()
    .oauth2ResourceServer()
        .authenticationManagerResolver(this.tokenAuthenticationManagerResolver);

```

Multi-tenancy

A resource server is considered multi-tenant when there are multiple strategies for verifying a bearer token, keyed by some tenant identifier.

For example, your resource server may accept bearer tokens from two different authorization servers. Or, your authorization server may represent a multiplicity of issuers.

In each case, there are two things that need to be done and trade-offs associated with how you choose to do them:

1. Resolve the tenant
2. Propagate the tenant

Resolving the Tenant By Request Material

Resolving the tenant by request material can be done by implementing an `AuthenticationManagerResolver`, which determines the `AuthenticationManager` at runtime, like so:

```

@Component
public class TenantAuthenticationManagerResolver
    implements AuthenticationManagerResolver<HttpServletRequest> {
    private final BearerTokenResolver resolver = new DefaultBearerTokenResolver();
    private final TenantRepository tenants; ❶

    private final Map<String, AuthenticationManager> authenticationManagers = new
    ConcurrentHashMap<>(); ❷

    public TenantAuthenticationManagerResolver(TenantRepository tenants) {
        this.tenants = tenants;
    }

    @Override
    public AuthenticationManager resolve(HttpServletRequest request) {
        return this.authenticationManagers.computeIfAbsent(toTenant(request), this::fromTenant);
    }

    private String toTenant(HttpServletRequest request) {
        String[] pathParts = request.getRequestURI().split("/");
        return pathParts.length > 0 ? pathParts[1] : null;
    }

    private AuthenticationManager fromTenant(String tenant) {
        return Optional.ofNullable(this.tenants.get(tenant)) ❸
            .map(JwtDecoders::fromIssuerLocation) ❹
            .map(JwtAuthenticationProvider::new)
            .orElseThrow(() -> new IllegalArgumentException("unknown tenant"))::authenticate;
    }
}

```

- ❶ A hypothetical source for tenant information
- ❷ A cache for `AuthenticationManager`s, keyed by tenant identifier
- ❸ Looking up the tenant is more secure than simply computing the issuer location on the fly - the lookup acts as a tenant whitelist
- ❹ Create a `JwtDecoder` via the discovery endpoint - the lazy lookup here means that you don't need to configure all tenants at startup

And then specify this `AuthenticationManagerResolver` in the DSL:

```

http
    .authorizeRequests()
        .anyRequest().authenticated()
        .and()
    .oauth2ResourceServer()
        .authenticationManagerResolver(this.tenantAuthenticationManagerResolver);

```

Resolving the Tenant By Claim

Resolving the tenant by claim is similar to doing so by request material. The only real difference is the `toTenant` method implementation:

```

@Component
public class TenantAuthenticationManagerResolver implements
AuthenticationManagerResolver<HttpServletRequest> {
    private final BearerTokenResolver resolver = new DefaultBearerTokenResolver();
    private final TenantRepository tenants; ❶

    private final Map<String, AuthenticationManager> authenticationManagers = new
ConcurrentHashMap<>(); ❷

    public TenantAuthenticationManagerResolver(TenantRepository tenants) {
        this.tenants = tenants;
    }

    @Override
    public AuthenticationManager resolve(HttpServletRequest request) {
        return this.authenticationManagers.computeIfAbsent(toTenant(request), this::fromTenant); ❸
    }

    private String toTenant(HttpServletRequest request) {
        try {
            String token = this.resolver.resolve(request);
            return (String) JWTParser.parse(token).getJWTClaimsSet().getIssuer();
        } catch (Exception e) {
            throw new IllegalArgumentException(e);
        }
    }

    private AuthenticationManager fromTenant(String tenant) {
        return Optional.ofNullable(this.tenants.get(tenant)) ❹
            .map(JwtDecoders::fromIssuerLocation) ❺
            .map(JwtAuthenticationProvider::new)
            .orElseThrow(() -> new IllegalArgumentException("unknown tenant"))::authenticate;
    }
}

```

- ❶ A hypothetical source for tenant information
- ❷ A cache for `AuthenticationManager`s, keyed by tenant identifier
- ❸❹ Looking up the tenant is more secure than simply computing the issuer location on the fly - the lookup acts as a tenant whitelist
- ❺ Create a `JwtDecoder` via the discovery endpoint - the lazy lookup here means that you don't need to configure all tenants at startup

```

http
    .authorizeRequests()
        .anyRequest().authenticated()
        .and()
    .oauth2ResourceServer()
        .authenticationManagerResolver(this.tenantAuthenticationManagerResolver);

```

Parsing the Claim Only Once

You may have observed that this strategy, while simple, comes with the trade-off that the JWT is parsed once by the `AuthenticationManagerResolver` and then again by the `JwtDecoder`.

This extra parsing can be alleviated by configuring the `JwtDecoder` directly with a `JWTClaimSetAwareJWSKeySelector` from Nimbus:

```

@Component
public class TenantJWSKeySelector
    implements JWTClaimSetAwareJWSKeySelector<SecurityContext> {

    private final TenantRepository tenants; ❶
    private final Map<String, JWSKeySelector<SecurityContext>> selectors = new ConcurrentHashMap<>(); ❷

    public TenantJWSKeySelector(TenantRepository tenants) {
        this.tenants = tenants;
    }

    @Override
    public List<? extends Key> selectKeys(JWSHeader jwsHeader, JWTClaimsSet jwtClaimsSet,
        SecurityContext securityContext)
        throws KeySourceException {
        return this.selectors.computeIfAbsent(toTenant(jwtClaimsSet), this::fromTenant)
            .selectJWSKeys(jwsHeader, securityContext);
    }

    private String toTenant(JWTClaimsSet claimSet) {
        return (String) claimSet.getClaim("iss");
    }

    private JWSKeySelector<SecurityContext> fromTenant(String tenant) {
        return Optional.ofNullable(this.tenantRepository.findById(tenant)) ❸
            .map(t -> t.getAttribute("jwks_uri"))
            .map(this::fromUri)
            .orElseThrow(() -> new IllegalArgumentException("unknown tenant"));
    }

    private JWSKeySelector<SecurityContext> fromUri(String uri) {
        try {
            return JWSAlgorithmFamilyJWSKeySelector.fromJWKSetURL(new URL(uri)); ❹
        } catch (Exception e) {
            throw new IllegalArgumentException(e);
        }
    }
}

```

- ❶ A hypothetical source for tenant information
- ❷ A cache for `JWSKeySelector`s`, keyed by tenant identifier
- ❸ Looking up the tenant is more secure than simply calculating the JWK Set endpoint on the fly - the lookup acts as a tenant whitelist
- ❹ Create a `JWSKeySelector` via the types of keys that come back from the JWK Set endpoint - the lazy lookup here means that you don't need to configure all tenants at startup

The above key selector is a composition of many key selectors. It chooses which key selector to use based on the `iss` claim in the JWT.

Note

To use this approach, make sure that the authorization server is configured to include the claim set as part of the token's signature. Without this, you have no guarantee that the issuer hasn't been altered by a bad actor.

Next, we can construct a `JWTProcessor`:


```

@Bean
JWTProcessor jwtProcessor(JWTClaimSetJWSKeySelector keySelector) {
    ConfigurableJWTProcessor<SecurityContext> jwtProcessor =
        new DefaultJWTProcessor();
    jwtProcessor.setJWTClaimSetJWSKeySelector(keySelector);
    return jwtProcessor;
}

```

As you are already seeing, the trade-off for moving tenant-awareness down to this level is more configuration. We have just a bit more.

Next, we still want to make sure you are validating the issuer. But, since the issuer may be different per JWT, then you'll need a tenant-aware validator, too:

```

@Component
public class TenantJwtIssuerValidator implements OAuth2TokenValidator<Jwt> {
    private final TenantRepository tenants;
    private final Map<String, JwtIssuerValidator> validators = new ConcurrentHashMap<>();

    public TenantJwtIssuerValidator(TenantRepository tenants) {
        this.tenants = tenants;
    }

    @Override
    public OAuth2TokenValidatorResult validate(Jwt token) {
        return this.validators.computeIfAbsent(toTenant(token), this::fromTenant)
            .validate(token);
    }

    private String toTenant(Jwt jwt) {
        return jwt.getIssuer();
    }

    private JwtIssuerValidator fromTenant(String tenant) {
        return Optional.ofNullable(this.tenants.findById(tenant))
            .map(t -> t.getAttribute("issuer"))
            .map(JwtIssuerValidator::new)
            .orElseThrow(() -> new IllegalArgumentException("unknown tenant"));
    }
}

```

Now that we have a tenant-aware processor and a tenant-aware validator, we can proceed with creating our `JwtDecoder`:

```

@Bean
JwtDecoder jwtDecoder(JWTProcessor jwtProcessor, OAuth2TokenValidator<Jwt> jwtValidator) {
    NimbusJwtDecoder decoder = new NimbusJwtDecoder(processor);
    OAuth2TokenValidator<Jwt> validator = new DelegatingOAuth2TokenValidator<>
        (JwtValidators.createDefault(), this.jwtValidator);
    decoder.setJwtValidator(validator);
    return decoder;
}

```

We've finished talking about resolving the tenant.

If you've chosen to resolve the tenant by request material, then you'll need to make sure you address your downstream resource servers in the same way. For example, if you are resolving it by subdomain, you'll need to address the downstream resource server using the same subdomain.

However, if you resolve it by a claim in the bearer token, read on to learn about [Spring Security's support for bearer token propagation](#).

Bearer Token Resolution

By default, Resource Server looks for a bearer token in the `Authorization` header. This, however, can be customized in a couple of ways.

Reading the Bearer Token from a Custom Header

For example, you may have a need to read the bearer token from a custom header. To achieve this, you can wire a `HeaderBearerTokenResolver` instance into the DSL, as you can see in the following example:

```
http
    .oauth2ResourceServer()
        .bearerTokenResolver(new HeaderBearerTokenResolver("x-goog-iap-jwt-assertion"));
```

Reading the Bearer Token from a Form Parameter

Or, you may wish to read the token from a form parameter, which you can do by configuring the `DefaultBearerTokenResolver`, as you can see below:

```
DefaultBearerTokenResolver resolver = new DefaultBearerTokenResolver();
resolver.setAllowFormEncodedBodyParameter(true);
http
    .oauth2ResourceServer()
        .bearerTokenResolver(resolver);
```

Bearer Token Propagation

Now that you're in possession of a bearer token, it might be handy to pass that to downstream services. This is quite simple with [ServletBearerExchangeFilterFunction](#), which you can see in the following example:

```
@Bean
public WebClient rest() {
    return WebClient.builder()
        .filter(new ServletBearerExchangeFilterFunction())
        .build();
}
```

When the above `WebClient` is used to perform requests, Spring Security will look up the current Authentication and extract any [AbstractOAuth2Token](#) credential. Then, it will propagate that token in the `Authorization` header.

For example:

```
this.rest.get()
    .uri("https://other-service.example.com/endpoint")
    .retrieve()
    .bodyToMono(String.class)
    .block();
```

Will invoke the <https://other-service.example.com/endpoint>, adding the bearer token `Authorization` header for you.

In places where you need to override this behavior, it's a simple matter of supplying the header yourself, like so:

```
this.rest.get()
    .uri("https://other-service.example.com/endpoint")
    .headers(headers -> headers.setBearerAuth(overridingToken))
    .retrieve()
    .bodyToMono(String.class)
    .block()
```

In this case, the filter will fall back and simply forward the request onto the rest of the web filter chain.

Note

Unlike the [OAuth 2.0 Client filter function](#), this filter function makes no attempt to renew the token, should it be expired. To obtain this level of support, please use the OAuth 2.0 Client filter.

RestTemplate support

There is no dedicated support for `RestTemplate` at the moment, but you can achieve propagation quite simply with your own interceptor:

```
@Bean
RestTemplate rest() {
    RestTemplate rest = new RestTemplate();
    rest.getInterceptors().add((request, body, execution) -> {
        Authentication authentication = SecurityContextHolder.getContext().getAuthentication();
        if (authentication == null) {
            return execution.execute(request, body);
        }

        if (!(authentication.getCredentials() instanceof AbstractOAuth2Token)) {
            return execution.execute(request, body);
        }

        AbstractOAuth2Token token = (AbstractOAuth2Token) authentication.getCredentials();
        request.getHeaders().setBearerAuth(token.getTokenValue());
        return execution.execute(request, body);
    });
    return rest;
}
```

13. SAML2

13.1 SAML 2.0 Login

The SAML 2.0 Login, `saml2Login()`, feature provides an application with the capability to have users log in to the application by using their existing account at an SAML 2.0 Identity Provider (Okta, ADFS, etc).

Note

SAML 2.0 Login is implemented by using the **Web Browser SSO Profile**, as specified in [SAML 2 Profiles](#). Our implementation is currently limited to a simple authentication scheme.

SAML 2 Support in Spring Security

SAML 2 Service Provider, SP a.k.a. a relying party, support existed as an [independent project](#) since 2009. The 1.0.x branch is still in use, including in the [Cloud Foundry User Account and Authentication Server](#) that also created a SAML 2.0 Identity Provider implementation based on the SP implementation.

In 2018 we experimented with creating an updated implementation of both a [Service Provider and Identity Provider](#) as a standalone library. After careful, and lengthy, deliberation we, the Spring Security team, decided to discontinue that effort. While this effort created a replacement for that standalone 1.0.x library we didn't feel that we should build a library on top of another library.

Instead we opted to provide framework support for SAML 2 authentication as part of [core Spring Security](#) instead.

Saml 2 Login - High Level Concepts

`saml2Login()` is aimed to support a fraction of the [SAML 2 feature set](#) with a focus on authentication being a Service Provider, SP, a relying party, receiving XML assertions from an Identity Provider, aka an asserting party.

A SAML 2 login, or authentication, is the concept that the SP receives and validates an XML message called an assertion from an IDP.

There are currently two supported authentication flows

1. IDP Initiated flow - example: You login in directly to Okta, and then select a web application to be authenticated for. Okta, the IDP, sends an assertion to the web application, the SP.
2. SP Initiated flow - example: You access a web application, a SP, the application sends an authentication request to the IDP requesting an assertion. Upon successful authentication on the IDP, the IDP sends an assertion to the SP.

Saml 2 Login - Current Feature Set

1. Service Provider (SP/Relying Party) is identified by `entityId = {baseUrl}/saml2/service-provider-metadata/{registrationId}`
2. Receive assertion embedded in a SAML response via Http-POST or Http-Redirect at `{baseUrl}/login/saml2/sso/{registrationId}`

3. Requires the assertion to be signed, unless the response is signed
4. Supports encrypted assertions
5. Supports encrypted NameId elements
6. Allows for extraction of assertion attributes into authorities using a `Converter<Assertion, Collection<? extends GrantedAuthority>>`
7. Allows mapping and white listing of authorities using a `GrantedAuthoritiesMapper`
8. Public keys in `java.security.cert.X509Certificate` format.
9. SP Initiated Authentication via an `AuthNRequest`

Saml 2 Login - Not Yet Supported

1. Mappings assertion conditions and attributes to session features (timeout, tracking, etc)
2. Single logout
3. Dynamic metadata generation
4. Receiving and validating standalone assertion (not wrapped in a response object)

Saml 2 Login - Introduction to Java Configuration

To add `saml2Login()` to a Spring Security filter chain, the minimal Java configuration requires a configuration repository, the `RelyingPartyRegistrationRepository`, that contains the SAML configuration and the invocation of the `HttpSecurity.saml2Login()` method:

```
@EnableWebSecurity
public class SecurityConfig extends WebSecurityConfigurerAdapter {

    @Bean
    public RelyingPartyRegistrationRepository relyingPartyRegistrationRepository() {
        //SAML configuration
        //Mapping this application to one or more Identity Providers
        return new InMemoryRelyingPartyRegistrationRepository(...);
    }

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .authorizeRequests()
                .anyRequest().authenticated()
                .and()
                .saml2Login()
            ;
    }
}
```

The bean declaration is a convenient, but optional, approach. You can directly wire up the repository using a method call

```

@EnableWebSecurity
public class SecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .authorizeRequests()
                .anyRequest().authenticated()
                .and()
            .saml2Login()
                .relyingPartyRegistrationRepository(...)
            ;
    }
}

```

RelyingPartyRegistration

The [RelyingPartyRegistration](#) object represents the mapping between this application, the SP, and the asserting party, the IDP.

URI Patterns

URI patterns are frequently used to automatically generate URIs based on an incoming request. The URI patterns in `saml2Login` can contain the following variables

- `baseUrl`
- `registrationId`
- `baseScheme`
- `baseHost`
- `basePort`

For example:

```
{baseUrl}/login/saml2/sso/{registrationId}
```

Relying Party

- `registrationId` - (required) a unique identifier for this configuration mapping. This identifier may be used in URI paths, so care should be taken that no URI encoding is required.
- `localEntityIdTemplate` - (optional) A URI pattern that creates an entity ID for this application based on the incoming request. The default is `{baseUrl}/saml2/service-provider-metadata/{registrationId}` and for a small sample application it would look like

```
http://localhost:8080/saml2/service-provider-metadata/my-test-configuration
```

There is no requirement that this configuration option is a pattern, it can be a fixed URI value.

- `remoteIdpEntityId` - (required) the entity ID of the Identity Provider. Always a fixed URI value or string, no patterns allowed.
- `assertionConsumerServiceUrlTemplate` - (optional) A URI pattern that denotes the assertion consumer service URI to be sent with any `AuthNRequest` from the SP to the IDP during the SP initiated flow. While this can be a pattern the actual URI must resolve to the ACS endpoint on the

SP. The default value is `{baseUrl}/login/saml2/sso/{registrationId}` and maps directly to the [Saml2WebSsoAuthenticationFilter](#) endpoint

- `idpWebSsoUrl` - (required) a fixed URI value for the IDP Single Sign On endpoint where the SP sends the `AuthNRequest` messages.
- `credentials` - A list of credentials, private keys and x509 certificates, used for message signing, verification, encryption and decryption. This list can contain redundant credentials to allow for easy rotation of credentials. For example
 - [0] - `X509Certificate{VERIFICATION,ENCRYPTION}` - The IDP's first public key used for verification and encryption.
 - [1] - `X509Certificate/{VERIFICATION,ENCRYPTION}` - The IDP's second verification key used for verification. Encryption is always done using the first `ENCRYPTION` key in the list.
 - [2] - `PrivateKey/X509Certificate{SIGNING,DECRYPTION}` - The SP's first signing and decryption credential.
 - [3] - `PrivateKey/X509Certificate{SIGNING,DECRYPTION}` - The SP's second decryption credential. Signing is always done using the first `SIGNING` key in the list.

When an incoming message is received, signatures are always required, the system will first attempt to validate the signature using the certificate at index [0] and only move to the second credential if the first one fails.

In a similar fashion, the SP configured private keys are used for decryption and attempted in the same order. The first SP credential (`type=SIGNING`) will be used when messages to the IDP are signed.

Duplicated Relying Party Configurations

In the use case where an application uses multiple identity providers it becomes obvious that some configuration is duplicated between two `RelyingPartyRegistration` objects

- `localEntityIdTemplate`
- `credentials` (all SP credentials, IDP credentials change)
- `assertionConsumerServiceUrlTemplate`

While there is some drawback in duplicating configuration values the back end configuration repository does not need to replicate this data storage model.

There is a benefit that comes with this setup. Credentials may be more easily rotated for some identity providers vs others. This object model can ensure that there is no disruption when configuration is changed in a multi IDP use case and you're not able to rotate credentials on all the identity providers.

Service Provider Metadata

The Spring Security SAML 2 implementation does not yet provide an endpoint for downloading SP metadata in XML format. The minimal pieces that are exchanged

- **entity ID** - defaults to `{baseUrl}/saml2/service-provider-metadata/{registrationId}`
Other known configuration names that also use this same value
 - Audience Restriction

- **single signon URL** - defaults to `{baseUrl}/login/saml2/sso/{registrationId}` Other known configuration names that also use this same value
 - Recipient URL
 - Destination URL
 - Assertion Consumer Service URL
- X509Certificate - the certificate that you configure as part of your `{SIGNING,DECRYPTION}` credentials must be shared with the Identity Provider

Authentication Requests - SP Initiated Flow

To initiate an authentication from the web application, a simple redirect to

```
{baseUrl}/saml2/authenticate/{registrationId}
```

The endpoint will generate an `AuthNRequest` by invoking the `createAuthenticationRequest` method on a configurable factory. Just expose the `Saml2AuthenticationRequestFactory` as a bean in your configuration.

```
public interface Saml2AuthenticationRequestFactory {
    String createAuthenticationRequest(Saml2AuthenticationRequest request);
}
```

Spring Boot 2.x Sample

We are currently working with the the Spring Boot team on the [Auto Configuration for Spring Security SAML Login](#). In the meantime, we have provided a Spring Boot sample that supports a Yaml configuration.

To run the sample, follow these three steps

1. Launch the Spring Boot application

- `./gradlew :spring-security-samples-boot-saml2login:bootRun`

2. Open a browser

- <http://localhost:8080/>

3. This will take you to an identity provider, log in using:

- User: `user`
- Password: `password`

Multiple Identity Provider Sample

It's very simple to use multiple providers, but there are some defaults that may trip you up if you don't pay attention. In our SAML configuration of `RelyingPartyRegistration` objects, we default an SP entity ID to

```
{baseUrl}/saml2/service-provider-metadata/{registrationId}
```

That means in our two provider configuration, our system would look like


```

registration-1 (Identity Provider 1) - Our local SP Entity ID is:
http://localhost:8080/saml2/service-provider-metadata/registration-1

registration-2 (Identity Provider 2) - Our local SP Entity ID is:
http://localhost:8080/saml2/service-provider-metadata/registration-2

```

In this configuration, illustrated in the sample below, to the outside world, we have actually created two virtual Service Provider identities hosted within the same application.

```

spring:
  security:
    saml2:
      login:
        relying-parties:
          - entity-id: &idp-entity-id https://simplesaml-for-spring-saml.cfapps.io/saml2/idp/
            metadata.php
            registration-id: simplesamlphp
            web-sso-url: &idp-sso-url https://simplesaml-for-spring-saml.cfapps.io/saml2/idp/
              SSOService.php
            signing-credentials: &service-provider-credentials
              - private-key: |
                  -----BEGIN PRIVATE KEY-----
                  MIIcCeAIBADANBgkqhkiG9w0BAQEFAASCAmIwggJeAgEAAoGBANG7v8QjQGU3MwQE
                  .....SHORTENED FOR READ ABILITY.....
                  INRtuLp4YHbgklmi
                  -----END PRIVATE KEY-----
                certificate: |
                  -----BEGIN CERTIFICATE-----
                  MIIcCgTCCAcCCQCuVzyqFgMSyDANBgkqhkiG9w0BAQsFADCBhDELMAkGA1UEBhMC
                  .....SHORTENED FOR READ ABILITY.....
                  RZ/nbTJ7VTeZOSyRoVn5XHhpuJ0B
                  -----END CERTIFICATE-----
            verification-credentials: &idp-certificates
              - |
                  -----BEGIN CERTIFICATE-----
                  MIIIEEzCCAvugAwIBAgIJAIClqzLrv+5nMA0GCSqGSIb3DQEBCwUAMIGfMQswCQYD
                  .....SHORTENED FOR READ ABILITY.....
                  lx13Y1YlQ4/tlpgTgfIJxKV6nyPiLoK0nywbMG+vpAirDt2Oc+hk
                  -----END CERTIFICATE-----
            - entity-id: *idp-entity-id
              registration-id: simplesamlphp2
              web-sso-url: *idp-sso-url
              signing-credentials: *service-provider-credentials
              verification-credentials: *idp-certificates

```

If this is not desirable, you can manually override the local SP entity ID by using the

```

localEntityIdTemplate = {baseUrl}/saml2/service-provider-metadata

```

If we change our local SP entity ID to this value, it is still important that we give out the correct single sign on URL (the assertion consumer service URL) for each registered identity provider based on the registration Id. `{baseUrl}/login/saml2/sso/{registrationId}`

14. Protection Against Exploits

14.1 Cross Site Request Forgery (CSRF) for Servlet Environments

This section discusses Spring Security's [Cross Site Request Forgery \(CSRF\)](#) support for servlet environments.

Using Spring Security CSRF Protection

The steps to using Spring Security's CSRF protection are outlined below:

- [Use proper HTTP verbs](#)
- [Configure CSRF Protection](#)
- [Include the CSRF Token](#)

Use proper HTTP verbs

The first step to protecting against CSRF attacks is to ensure your website uses proper HTTP verbs. This is covered in detail in [Safe Methods Must be Idempotent](#).

Configure CSRF Protection

The next step is to configure Spring Security's CSRF protection within your application. Spring Security's CSRF protection is enabled by default, but you may need to customize the configuration. Below are a few common customizations.

Custom CsrfTokenRepository

By default Spring Security stores the expected CSRF token in the `HttpSession` using `HttpSessionCsrfTokenRepository`. There can be cases where users will want to configure a custom `CsrfTokenRepository`. For example, it might be desirable to persist the `CsrfToken` in a cookie to [support a JavaScript based application](#).

By default the `CookieCsrfTokenRepository` will write to a cookie named `XSRF-TOKEN` and read it from a header named `X-XSRF-TOKEN` or the HTTP parameter `_csrf`. These defaults come from [AngularJS](#)

You can configure `CookieCsrfTokenRepository` in XML using the following:

```
<http>
  <!-- ... -->
  <csrf token-repository-ref="tokenRepository"/>
</http>
<b:bean id="tokenRepository"
  class="org.springframework.security.web.csrf.CookieCsrfTokenRepository"
  p:cookieHttpOnly="false"/>
```

Example 14.1 Store CSRF Token in a Cookie with XML Configuration

Note

The sample explicitly sets `cookieHttpOnly=false`. This is necessary to allow JavaScript (i.e. AngularJS) to read it. If you do not need the ability to read the cookie with JavaScript directly, it is recommended to omit `cookieHttpOnly=false` to improve security.

You can configure `CookieCsrfTokenRepository` in Java Configuration using:

```
@EnableWebSecurity
public class WebSecurityConfig extends
    WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) {
        http
            .csrf(csrf ->
                csrf
                    .csrfTokenRepository(CookieCsrfTokenRepository.withHttpOnlyFalse())
            );
    }
}
```

Example 14.2 Store CSRF Token in a Cookie with Java Configuration

Note

The sample explicitly sets `cookieHttpOnly=false`. This is necessary to allow JavaScript (i.e. AngularJS) to read it. If you do not need the ability to read the cookie with JavaScript directly, it is recommended to omit `cookieHttpOnly=false` (by using new `CookieCsrfTokenRepository()` instead) to improve security.

Disable CSRF Protection

CSRF protection is enabled by default. However, it is simple to disable CSRF protection if it [makes sense for your application](#).

The XML configuration below will disable CSRF protection.

```
<http>
  <!-- ... -->
  <csrf disabled="true"/>
</http>
```

Example 14.3 Disable CSRF XML Configuration

The Java configuration below will disable CSRF protection.

```
@Configuration
@EnableWebSecurity
public class WebSecurityConfig extends
    WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) {
        http
            .csrf(csrf ->
                csrf.disable()
            );
    }
}
```

Example 14.4 Disable CSRF Java Configuration

Include the CSRF Token

In order for the [synchronizer token pattern](#) to protect against CSRF attacks, we must include the actual CSRF token in the HTTP request. This must be included in a part of the request (i.e. form parameter, HTTP header, etc) that is not automatically included in the HTTP request by the browser.

Spring Security's [CsrfFilter](#) exposes a [CsrfToken](#) as an `HttpServletRequest` attribute named `_csrf`. This means that any view technology can access the `CsrfToken` to expose the expected token as either a [form](#) or [meta tag](#). Fortunately, there are integrations listed below that make including the token in [form](#) and [ajax](#) requests even easier.

Form URL Encoded

In order to post an HTML form the CSRF token must be included in the form as a hidden input. For example, the rendered HTML might look like:

```
<input type="hidden"
  name="_csrf"
  value="4bfd1575-3ad1-4d21-96c7-4ef2d9f86721"/>
```

Example 14.5 CSRF Token HTML

Next we will discuss various ways of including the CSRF token in a form as a hidden input.

Automatic CSRF Token Inclusion

Spring Security's CSRF support provides integration with Spring's [RequestDataValueProcessor](#) via its [CsrfRequestDataValueProcessor](#). This means that if you leverage [Spring's form tag library](#), [Thymleaf](#), or any other view technology that integrates with `RequestDataValueProcessor`, then forms that have an unsafe HTTP method (i.e. post) will automatically include the actual CSRF token.

csrfInput Tag

If you are using JSPs, then you can use [Spring's form tag library](#). However, if that is not an option, you can also easily include the token with the [csrfInput](#) tag.

CsrfToken Request Attribute

If the [other options](#) for including the actual CSRF token in the request do not work, you can take advantage of the fact that the `CsrfToken` [is exposed](#) as an `HttpServletRequest` attribute named `_csrf`.

An example of doing this with a JSP is shown below:

```
<c:url var="logoutUrl" value="/logout"/>
<form action="${logoutUrl}"
  method="post">
<input type="submit"
  value="Log out" />
<input type="hidden"
  name="${_csrf.parameterName}"
  value="${_csrf.token}"/>
</form>
```

Example 14.6 CSRF Token in Form with Request Attribute

Ajax and JSON Requests

If you are using JSON, then it is not possible to submit the CSRF token within an HTTP parameter. Instead you can submit the token within a HTTP header.

In the following sections we will discuss various ways of including the CSRF token as an HTTP request header in JavaScript based applications.

Automatic Inclusion

Spring Security can easily be [configured](#) to store the expected CSRF token in a cookie. By storing the expected CSRF in a cookie, JavaScript frameworks like [AngularJS](#) will automatically include the actual CSRF token in the HTTP request headers.

Meta tags

An alternative pattern to [exposing the CSRF in a cookie](#) is to include the CSRF token within your meta tags. The HTML might look something like this:

```
<html>
<head>
  <meta name="_csrf" content="4bfd1575-3ad1-4d21-96c7-4ef2d9f86721"/>
  <meta name="_csrf_header" content="X-CSRF-TOKEN"/>
  <!-- ... -->
</head>
<!-- ... -->
```

Example 14.7 CSRF meta tag HTML

Once the meta tags contained the CSRF token, the JavaScript code would read the meta tags and include the CSRF token as a header. If you were using jQuery, this could be done with the following:

```
$(function () {
  var token = $("meta[name='_csrf']").attr("content");
  var header = $("meta[name='_csrf_header']").attr("content");
  $(document).ajaxSend(function(e, xhr, options) {
    xhr.setRequestHeader(header, token);
  });
});
```

Example 14.8 AJAX send CSRF Token

csrfMeta tag

If you are using JSPs a simple way to write the CSRF token to the meta tags is by leveraging the [csrfMeta](#) tag.

CsrfToken Request Attribute

If the [other options](#) for including the actual CSRF token in the request do not work, you can take advantage of the fact that the `CsrfToken` is [exposed](#) as an `HttpServletRequest` attribute named `_csrf`. An example of doing this with a JSP is shown below:

```
<html>
<head>
  <meta name="_csrf" content="${_csrf.token}"/>
  <!-- default header name is X-CSRF-TOKEN -->
  <meta name="_csrf_header" content="${_csrf.headerName}"/>
  <!-- ... -->
</head>
<!-- ... -->
```

Example 14.9 CSRF meta tag JSP

CSRF Considerations

There are a few special considerations to consider when implementing protection against CSRF attacks. This section discusses those considerations as it pertains to servlet environments. Refer to the section called “CSRF Considerations” for a more general discussion.

Logging In

It is important to [require CSRF for log in](#) requests to protect against forging log in attempts. Spring Security's servlet support does this out of the box.

Logging Out

It is important to [require CSRF for log out](#) requests to protect against forging log out attempts. If CSRF protection is enabled (default), Spring Security's `LogoutFilter` to only process HTTP POST. This ensures that log out requires a CSRF token and that a malicious user cannot forcibly log out your users.

The easiest approach is to use a form to log out. If you really want a link, you can use JavaScript to have the link perform a POST (i.e. maybe on a hidden form). For browsers with JavaScript that is disabled, you can optionally have the link take the user to a log out confirmation page that will perform the POST.

If you really want to use HTTP GET with logout you can do so, but remember this is generally not recommended. For example, the following Java Configuration will perform logout with the URL `/logout` is requested with any HTTP method:

```
@EnableWebSecurity
public class WebSecurityConfig extends
    WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) {
        http
            .logout(logout ->
                logout
                    .logoutRequestMatcher(new AntPathRequestMatcher("/logout"))
            );
    }
}
```

Example 14.10 Log out with HTTP GET

CSRF and Session Timeouts

By default Spring Security stores the CSRF token in the `HttpSession`. This can lead to a situation where the session expires which means there is not an expected CSRF token to validate against.

We've already discussed [general solutions](#) to session timeouts. This section discusses the specifics of CSRF timeouts as it pertains to the servlet support.

It is simple to change storage of the expected CSRF token to be in a cookie. For details, refer to the the section called "Custom `CsrfTokenRepository`" section.

If a token does expire, you might want to customize how it is handled by specifying a custom `AccessDeniedHandler`. The custom `AccessDeniedHandler` can process the `InvalidCsrfTokenException` any way you like. For an example of how to customize the `AccessDeniedHandler` refer to the provided links for both [xml](#) and [Java configuration](#).

Multipart (file upload)

We have [already discussed](#) how protecting multipart requests (file uploads) from CSRF attacks causes a [chicken and the egg](#) problem. This section discusses how to implement placing the CSRF token in the [body](#) and [url](#) within a servlet application.

Note

More information about using multipart forms with Spring can be found within the [1.1.11. Multipart Resolver](#) section of the Spring reference and the [MultipartFile javadoc](#).

Place CSRF Token in the Body

We have [already discussed](#) the tradeoffs of placing the CSRF token in the body. In this section we will discuss how to configure Spring Security to read the CSRF from the body.

In order to read the CSRF token from the body, the `MultipartFile` is specified before the Spring Security filter. Specifying the `MultipartFile` before the Spring Security filter means that there is no authorization for invoking the `MultipartFile` which means anyone can place temporary files on your server. However, only authorized users will be able to submit a File that is processed by your application. In general, this is the recommended approach because the temporary file upload should have a negligible impact on most servers.

To ensure `MultipartFile` is specified before the Spring Security filter with java configuration, users can override `beforeSpringSecurityFilterChain` as shown below:

```
public class SecurityApplicationInitializer extends AbstractSecurityWebApplicationInitializer {

    @Override
    protected void beforeSpringSecurityFilterChain(ServletContext servletContext) {
        insertFilters(servletContext, new MultipartFile());
    }
}
```

Example 14.11 Initializer MultipartFile

To ensure `MultipartFile` is specified before the Spring Security filter with XML configuration, users can ensure the `<filter-mapping>` element of the `MultipartFile` is placed before the `springSecurityFilterChain` within the `web.xml` as shown below:

```
<filter>
  <filter-name>MultipartFile</filter-name>
  <filter-class>org.springframework.web.multipart.support.MultipartFile</filter-class>
</filter>
<filter>
  <filter-name>springSecurityFilterChain</filter-name>
  <filter-class>org.springframework.web.filter.DelegatingFilterProxy</filter-class>
</filter>
<filter-mapping>
  <filter-name>MultipartFile</filter-name>
  <url-pattern>/*</url-pattern>
</filter-mapping>
<filter-mapping>
  <filter-name>springSecurityFilterChain</filter-name>
  <url-pattern>/*</url-pattern>
</filter-mapping>
```

*Example 14.12 web.xml - MultipartFile***Include CSRF Token in URL**

If allowing unauthorized users to upload temporary files is not acceptable, an alternative is to place the `MultipartFile` after the Spring Security filter and include the CSRF as a query parameter in the action attribute of the form. Since the `CsrfToken` is exposed as an `HttpServletRequest` [request attribute](#), we can use that to create an action with the CSRF token in it. An example with a jsp is shown below

```
<form method="post"
      action="./upload?${_csrf.parameterName}=${_csrf.token}"
      enctype="multipart/form-data">
```

Example 14.13 CSRF Token in Action

HiddenHttpMethodFilter

We have [already discussed](#) the trade-offs of placing the CSRF token in the body.

In Spring's Servlet support, overriding the HTTP method is done using [HiddenHttpMethodFilter](#). More information can be found in [HTTP Method Conversion](#) section of the reference documentation.

14.2 Security HTTP Response Headers

This section discusses Spring Security's support for adding various security headers to the response.

Default Security Headers

Spring Security allows users to easily inject the default security headers to assist in protecting their application. The default for Spring Security is to include the following headers:

```
Cache-Control: no-cache, no-store, max-age=0, must-revalidate
Pragma: no-cache
Expires: 0
X-Content-Type-Options: nosniff
Strict-Transport-Security: max-age=31536000 ; includeSubDomains
X-Frame-Options: DENY
X-XSS-Protection: 1; mode=block
```

Note

Strict-Transport-Security is only added on HTTPS requests

For additional details on each of these headers, refer to the corresponding sections:

- [Cache Control](#)
- [Content Type Options](#)
- [HTTP Strict Transport Security](#)
- [X-Frame-Options](#)
- [X-XSS-Protection](#)

While each of these headers are considered best practice, it should be noted that not all clients utilize the headers, so additional testing is encouraged.

You can customize specific headers. For example, assume that you want your HTTP response headers to look like the following:

```
Cache-Control: no-cache, no-store, max-age=0, must-revalidate
Pragma: no-cache
Expires: 0
X-Content-Type-Options: nosniff
X-Frame-Options: SAMEORIGIN
X-XSS-Protection: 1; mode=block
```


Specifically, you want all of the default headers with the following customizations:

- [X-Frame-Options](#) to allow any request from same domain
- [HTTP Strict Transport Security \(HSTS\)](#) will not be added to the response

You can easily do this with the following Java Configuration:

```
@EnableWebSecurity
public class WebSecurityConfig extends
    WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .frameOptions(frameOptions ->
                        frameOptions.sameOrigin()
                    )
                    .httpStrictTransportSecurity(hsts ->
                        hsts.disable()
                    )
            );
    }
}
```

Alternatively, if you are using Spring Security XML Configuration, you can use the following:

```
<http>
  <!-- ... -->

  <headers>
    <frame-options policy="SAMEORIGIN" />
    <hsts disable="true"/>
  </headers>
</http>
```

If you do not want the defaults to be added and want explicit control over what should be used, you can disable the defaults. An example for both Java and XML based configuration is provided below:

If you are using Spring Security's Java Configuration the following will only add [Cache Control](#).

```
@EnableWebSecurity
public class WebSecurityConfig extends
    WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    // do not use any default headers unless explicitly listed
                    .defaultsDisabled()
                    .cacheControl(withDefaults())
            );
    }
}
```

The following XML will only add [Cache Control](#).

```
<http>
  <!-- ... -->

  <headers defaults-disabled="true">
    <cache-control/>
  </headers>
</http>
```

If necessary, you can disable all of the HTTP Security response headers with the following Java Configuration:

```
@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers.disable()
            );
    }
}
```

If necessary, you can disable all of the HTTP Security response headers with the following XML configuration below:

```
<http>
  <!-- ... -->

  <headers disabled="true" />
</http>
```

Cache Control

In the past Spring Security required you to provide your own cache control for your web application. This seemed reasonable at the time, but browser caches have evolved to include caches for secure connections as well. This means that a user may view an authenticated page, log out, and then a malicious user can use the browser history to view the cached page. To help mitigate this Spring Security has added cache control support which will insert the following headers into you response.

```
Cache-Control: no-cache, no-store, max-age=0, must-revalidate
Pragma: no-cache
Expires: 0
```

Simply adding the [<headers>](#) element with no child elements will automatically add Cache Control and quite a few other protections. However, if you only want cache control, you can enable this feature using Spring Security's XML namespace with the [<cache-control>](#) element and the [headers@defaults-disabled](#) attribute.

```
<http>
  <!-- ... -->

  <headers defaults-disable="true">
    <cache-control />
  </headers>
</http>
```

Similarly, you can enable only cache control within Java Configuration with the following:

```

@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .defaultsDisabled()
                    .cacheControl(withDefaults())
            );
    }
}

```

If you actually want to cache specific responses, your application can selectively invoke [HttpServletRequest.setHeader\(String,String\)](#) to override the header set by Spring Security. This is useful to ensure things like CSS, JavaScript, and images are properly cached.

When using Spring Web MVC, this is typically done within your configuration. For example, the following configuration will ensure that the cache headers are set for all of your resources:

```

@EnableWebMvc
public class WebMvcConfiguration implements WebMvcConfigurer {

    @Override
    public void addResourceHandlers(ResourceHandlerRegistry registry) {
        registry
            .addResourceHandler("/resources/**")
            .addResourceLocations("/resources/")
            .setCachePeriod(31556926);
    }

    // ...
}

```

Content Type Options

Historically browsers, including Internet Explorer, would try to guess the content type of a request using [content sniffing](#). This allowed browsers to improve the user experience by guessing the content type on resources that had not specified the content type. For example, if a browser encountered a JavaScript file that did not have the content type specified, it would be able to guess the content type and then execute it.

Note

=== There are many additional things one should do (i.e. only display the document in a distinct domain, ensure Content-Type header is set, sanitize the document, etc) when allowing content to be uploaded. However, these measures are out of the scope of what Spring Security provides. It is also important to point out when disabling content sniffing, you must specify the content type in order for things to work properly. ===

The problem with content sniffing is that this allowed malicious users to use polyglots (i.e. a file that is valid as multiple content types) to execute XSS attacks. For example, some sites may allow users to submit a valid postscript document to a website and view it. A malicious user might create a [postscript document that is also a valid JavaScript file](#) and execute a XSS attack with it.

Content sniffing can be disabled by adding the following header to our response:

```
X-Content-Type-Options: nosniff
```

Just as with the cache control element, the nosniff directive is added by default when using the `<headers>` element with no child elements. However, if you want more control over which headers are added you can use the `<content-type-options>` element and the `headers@defaults-disabled` attribute as shown below:

```
<http>
  <!-- ... -->

  <headers defaults-disabled="true">
    <content-type-options />
  </headers>
</http>
```

The X-Content-Type-Options header is added by default with Spring Security Java configuration. If you want more control over the headers, you can explicitly specify the content type options with the following:

```
@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .defaultsDisabled()
                    .contentTypeOptions(withDefaults())
            );
    }
}
```

HTTP Strict Transport Security (HSTS)

When you type in your bank's website, do you enter `mybank.example.com` or do you enter <https://mybank.example.com>? If you omit the https protocol, you are potentially vulnerable to [Man in the Middle attacks](#). Even if the website performs a redirect to <https://mybank.example.com> a malicious user could intercept the initial HTTP request and manipulate the response (i.e. redirect to <https://mibank.example.com> and steal their credentials).

Many users omit the https protocol and this is why [HTTP Strict Transport Security \(HSTS\)](#) was created. Once `mybank.example.com` is added as a [HSTS host](#), a browser can know ahead of time that any request to `mybank.example.com` should be interpreted as <https://mybank.example.com>. This greatly reduces the possibility of a Man in the Middle attack occurring.

Note

=== In accordance with [RFC6797](#), the HSTS header is only injected into HTTPS responses. In order for the browser to acknowledge the header, the browser must first trust the CA that signed the SSL certificate used to make the connection (not just the SSL certificate). ===

One way for a site to be marked as a HSTS host is to have the host preloaded into the browser. Another is to add the "Strict-Transport-Security" header to the response. For example the following would instruct the browser to treat the domain as an HSTS host for a year (there are approximately 31536000 seconds in a year):

```
Strict-Transport-Security: max-age=31536000 ; includeSubDomains ; preload
```

The optional `includeSubDomains` directive instructs Spring Security that subdomains (i.e. `secure.mybank.example.com`) should also be treated as an HSTS domain.

The optional `preload` directive instructs Spring Security that domain should be preloaded in browser as HSTS domain. For more details on HSTS preload please see <https://hstspreload.org>.

As with the other headers, Spring Security adds HSTS by default. You can customize HSTS headers with the `<hsts>` element as shown below:

```
<http>
  <!-- ... -->

  <headers>
    <hsts
      include-subdomains="true"
      max-age-seconds="31536000" preload="true" />
    </headers>
  </http>
```

Similarly, you can enable only HSTS headers with Java Configuration:

```
@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .httpStrictTransportSecurity(hsts ->
                        hsts
                            .includeSubDomains(true)
                            .preload(true)
                            .maxAgeInSeconds(31536000)
                        )
                )
            );
    }
}
```

HTTP Public Key Pinning (HPKP)

HTTP Public Key Pinning (HPKP) is a security feature that tells a web client to associate a specific cryptographic public key with a certain web server to prevent Man in the Middle (MITM) attacks with forged certificates.

To ensure the authenticity of a server's public key used in TLS sessions, this public key is wrapped into a X.509 certificate which is usually signed by a certificate authority (CA). Web clients such as browsers trust a lot of these CAs, which can all create certificates for arbitrary domain names. If an attacker is able to compromise a single CA, they can perform MITM attacks on various TLS connections. HPKP can circumvent this threat for the HTTPS protocol by telling the client which public key belongs to a certain web server. HPKP is a Trust on First Use (TOFU) technique. The first time a web server tells a client via a special HTTP header which public keys belong to it, the client stores this information for a given period of time. When the client visits the server again, it expects a certificate containing a public key whose fingerprint is already known via HPKP. If the server delivers an unknown public key, the client should present a warning to the user.

Note

=== Because the user-agent needs to validate the pins against the SSL certificate chain, the HPKP header is only injected into HTTPS responses. ===

Enabling this feature for your site is as simple as returning the Public-Key-Pins HTTP header when your site is accessed over HTTPS. For example, the following would instruct the user-agent to only report pin validation failures to a given URI (via the [report-uri](#) directive) for 2 pins:

```
Public-Key-Pins-Report-Only: max-age=5184000 ; pin-
sha256="d6qzRu9zOECb90Uez27xWltNsjoelMd7GkYYkVoZWmM=" ; pin-sha256="E9CZ9INdbd
+2eRQozYqqbQ2yXLVKB9+xcprMF+44Ulg=" ; report-uri="https://example.net/pkp-report" ; includeSubDomains
```

A [pin validation failure report](#) is a standard JSON structure that can be captured either by the web application's own API or by a publicly hosted HPKP reporting service, such as, [REPORT-URI](#).

The optional `includeSubDomains` directive instructs the browser to also validate subdomains with the given pins.

Opposed to the other headers, Spring Security does not add HPKP by default. You can customize HPKP headers with the `<hpkp>` element as shown below:

```
<http>
  <!-- ... -->

  <headers>
    <hpkp
      include-subdomains="true"
      report-uri="https://example.net/pkp-report">
      <pins>
        <pin algorithm="sha256">d6qzRu9zOECb90Uez27xWltNsjoelMd7GkYYkVoZWmM=</pin>
        <pin algorithm="sha256">E9CZ9INdbd+2eRQozYqqbQ2yXLVKB9+xcprMF+44Ulg=</pin>
      </pins>
    </hpkp>
  </headers>
</http>
```

Similarly, you can enable HPKP headers with Java Configuration:

```
@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .httpPublicKeyPinning(hpkp ->
                        hpkp
                            .includeSubDomains(true)
                            .reportUri("https://example.net/pkp-report")
                            .addSha256Pins("d6qzRu9zOECb90Uez27xWltNsjoelMd7GkYYkVoZWmM=", "E9CZ9INdbd
+2eRQozYqqbQ2yXLVKB9+xcprMF+44Ulg=")
                    )
                );
    }
}
```

X-Frame-Options

Allowing your website to be added to a frame can be a security issue. For example, using clever CSS styling users could be tricked into clicking on something that they were not intending ([video demo](#)). For example, a user that is logged into their bank might click a button that grants access to other users. This sort of attack is known as [Clickjacking](#).

Note

=== Another modern approach to dealing with clickjacking is to use the section called “Content Security Policy (CSP)”. ===

There are a number ways to mitigate clickjacking attacks. For example, to protect legacy browsers from clickjacking attacks you can use [frame breaking code](#). While not perfect, the frame breaking code is the best you can do for the legacy browsers.

A more modern approach to address clickjacking is to use [X-Frame-Options](#) header:

```
X-Frame-Options: DENY
```

The X-Frame-Options response header instructs the browser to prevent any site with this header in the response from being rendered within a frame. By default, Spring Security disables rendering within an iframe.

You can customize X-Frame-Options with the [frame-options](#) element. For example, the following will instruct Spring Security to use "X-Frame-Options: SAMEORIGIN" which allows iframes within the same domain:

```
<http>
  <!-- ... -->

  <headers>
    <frame-options
      policy="SAMEORIGIN" />
  </headers>
</http>
```

Similarly, you can customize frame options to use the same origin within Java Configuration using the following:

```
@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .frameOptions(frameOptions ->
                        frameOptions
                            .sameOrigin()
                    )
            );
    }
}
```

X-XSS-Protection

Some browsers have built in support for filtering out [reflected XSS attacks](#). This is by no means foolproof, but does assist in XSS protection.

The filtering is typically enabled by default, so adding the header typically just ensures it is enabled and instructs the browser what to do when a XSS attack is detected. For example, the filter might try to change the content in the least invasive way to still render everything. At times, this type of replacement can become a [XSS vulnerability in itself](#). Instead, it is best to block the content rather than attempt to fix it. To do this we can add the following header:

```
X-XSS-Protection: 1; mode=block
```

This header is included by default. However, we can customize it if we wanted. For example:

```
<http>
  <!-- ... -->

  <headers>
    <xss-protection block="false"/>
  </headers>
</http>
```

Similarly, you can customize XSS protection within Java Configuration with the following:

```
@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .xssProtection(xssProtection ->
                        xssProtection
                            .block(false)
                    )
            );
    }
}
```

Content Security Policy (CSP)

[Content Security Policy \(CSP\)](#) is a mechanism that web applications can leverage to mitigate content injection vulnerabilities, such as cross-site scripting (XSS). CSP is a declarative policy that provides a facility for web application authors to declare and ultimately inform the client (user-agent) about the sources from which the web application expects to load resources.

Note

=== Content Security Policy is not intended to solve all content injection vulnerabilities. Instead, CSP can be leveraged to help reduce the harm caused by content injection attacks. As a first line of defense, web application authors should validate their input and encode their output. ===

A web application may employ the use of CSP by including one of the following HTTP headers in the response:

- **Content-Security-Policy**
- **Content-Security-Policy-Report-Only**

Each of these headers are used as a mechanism to deliver a **security policy** to the client. A security policy contains a set of **security policy directives** (for example, *script-src* and *object-src*), each responsible for declaring the restrictions for a particular resource representation.

For example, a web application can declare that it expects to load scripts from specific, trusted sources, by including the following header in the response:

```
Content-Security-Policy: script-src https://trustedscripts.example.com
```

An attempt to load a script from another source other than what is declared in the *script-src* directive will be blocked by the user-agent. Additionally, if the [report-uri](#) directive is declared in the security policy, then the violation will be reported by the user-agent to the declared URL.

For example, if a web application violates the declared security policy, the following response header will instruct the user-agent to send violation reports to the URL specified in the policy's *report-uri* directive.

```
Content-Security-Policy: script-src https://trustedscripts.example.com; report-uri /csp-report-endpoint/
```

[Violation reports](#) are standard JSON structures that can be captured either by the web application's own API or by a publicly hosted CSP violation reporting service, such as, [REPORT-URI](#).

The **Content-Security-Policy-Report-Only** header provides the capability for web application authors and administrators to monitor security policies, rather than enforce them. This header is typically used when experimenting and/or developing security policies for a site. When a policy is deemed effective, it can be enforced by using the *Content-Security-Policy* header field instead.

Given the following response header, the policy declares that scripts may be loaded from one of two possible sources.

```
Content-Security-Policy-Report-Only: script-src 'self' https://trustedscripts.example.com; report-uri /csp-report-endpoint/
```

If the site violates this policy, by attempting to load a script from *evil.com*, the user-agent will send a violation report to the declared URL specified by the *report-uri* directive, but still allow the violating resource to load nevertheless.

Configuring Content Security Policy

It's important to note that Spring Security **does not add** Content Security Policy by default. The web application author must declare the security policy(s) to enforce and/or monitor for the protected resources.

For example, given the following security policy:

```
script-src 'self' https://trustedscripts.example.com; object-src https://trustedplugins.example.com;
report-uri /csp-report-endpoint/
```

You can enable the CSP header using XML configuration with the [<content-security-policy>](#) element as shown below:

```

<http>
  <!-- ... -->

  <headers>
    <content-security-policy
      policy-directives="script-src 'self' https://trustedscripts.example.com; object-src https://
trustedplugins.example.com; report-uri /csp-report-endpoint/" />
    </headers>
  </http>

```

To enable the CSP *'report-only'* header, configure the element as follows:

```

<http>
  <!-- ... -->

  <headers>
    <content-security-policy
      policy-directives="script-src 'self' https://trustedscripts.example.com; object-src https://
trustedplugins.example.com; report-uri /csp-report-endpoint/"
      report-only="true" />
    </headers>
  </http>

```

Similarly, you can enable the CSP header using Java configuration as shown below:

```

@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .contentSecurityPolicy(csp ->
                        csp
                            .policyDirectives("script-src 'self' https://trustedscripts.example.com;
object-src https://trustedplugins.example.com; report-uri /csp-report-endpoint/")
                    )
            );
    }
}

```

To enable the CSP *'report-only'* header, provide the following Java configuration:

```

@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .contentSecurityPolicy(csp ->
                        csp
                            .policyDirectives("script-src 'self' https://trustedscripts.example.com;
object-src https://trustedplugins.example.com; report-uri /csp-report-endpoint/")
                            .reportOnly()
                    )
            );
    }
}

```

Additional Resources

Applying Content Security Policy to a web application is often a non-trivial undertaking. The following resources may provide further assistance in developing effective security policies for your site.

[An Introduction to Content Security Policy](#)

[CSP Guide - Mozilla Developer Network](#)

[W3C Candidate Recommendation](#)

Referrer Policy

[Referrer Policy](#) is a mechanism that web applications can leverage to manage the referrer field, which contains the last page the user was on.

Spring Security's approach is to use [Referrer Policy](#) header, which provides different [policies](#):

```
Referrer-Policy: same-origin
```

The Referrer-Policy response header instructs the browser to let the destination know the source where the user was previously.

Configuring Referrer Policy

Spring Security **doesn't add** Referrer Policy header by default.

You can enable the Referrer-Policy header using XML configuration with the `<referrer-policy>` element as shown below:

```
<http>
  <!-- ... -->

  <headers>
    <referrer-policy policy="same-origin" />
  </headers>
</http>
```

Similarly, you can enable the Referrer Policy header using Java configuration as shown below:

```
@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .referrerPolicy(referrerPolicy ->
                        referrerPolicy
                            .policy(ReferrerPolicy.SAME_ORIGIN)
                    )
            )
    };
}
```

Feature Policy

[Feature Policy](#) is a mechanism that allows web developers to selectively enable, disable, and modify the behavior of certain APIs and web features in the browser.

```
Feature-Policy: geolocation 'self'
```

With Feature Policy, developers can opt-in to a set of "policies" for the browser to enforce on specific features used throughout your site. These policies restrict what APIs the site can access or modify the browser's default behavior for certain features.

Configuring Feature Policy

Spring Security **doesn't add** Feature Policy header by default.

You can enable the Feature-Policy header using XML configuration with the [<feature-policy>](#) element as shown below:

```
<http>
  <!-- ... -->

  <headers>
    <feature-policy policy-directives="geolocation 'self'" />
  </headers>
</http>
```

Similarly, you can enable the Feature Policy header using Java configuration as shown below:

```
@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .featurePolicy("geolocation 'self'")
            );
    }
}
```

Clear Site Data

[Clear Site Data](#) is a mechanism by which any browser-side data - cookies, local storage, and the like - can be removed when an HTTP response contains this header:

```
Clear-Site-Data: "cache", "cookies", "storage", "executionContexts"
```

This is a nice clean-up action to perform on logout.

Configuring Clear Site Data

Spring Security **doesn't add** the Clear Site Data header by default.

You can configure your application to send down this header on logout like so:

```

@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .logout()
                .addLogoutHandler(new HeaderWriterLogoutHandler(new ClearSiteDataHeaderWriter(CACHE,
COOKIES)));
    }
}

```

Note

It's not recommended that you configure this header writer via the `headers()` directive. The reason for this is that any session state, say the `JSESSIONID` cookie, would be removed, effectively logging the user out.

Custom Headers

Spring Security has mechanisms to make it convenient to add the more common security headers to your application. However, it also provides hooks to enable adding custom headers.

Static Headers

There may be times you wish to inject custom security headers into your application that are not supported out of the box. For example, given the following custom security header:

```
X-Custom-Security-Header: header-value
```

When using the XML namespace, these headers can be added to the response using the [<header>](#) element as shown below:

```

<http>
  <!-- ... -->

  <headers>
    <header name="X-Custom-Security-Header" value="header-value"/>
  </headers>
</http>

```

Similarly, the headers could be added to the response using Java Configuration as shown in the following:

```

@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .addHeaderWriter(new StaticHeadersWriter("X-Custom-Security-Header", "header-value"));
    }
}

```

Headers Writer

When the namespace or Java configuration does not support the headers you want, you can create a custom `HeadersWriter` instance or even provide a custom implementation of the `HeadersWriter`.

Let's take a look at an example of using a custom instance of `XFrameOptionsHeaderWriter`. Perhaps you want to allow framing of content for the same origin. This is easily supported by setting the `policy` attribute to "SAMEORIGIN", but let's take a look at a more explicit example using the `ref` attribute.

```
<http>
  <!-- ... -->

  <headers>
    <header ref="frameOptionsWriter"/>
  </headers>
</http>
<!-- Requires the c-namespace.
See https://docs.spring.io/spring/docs/current/spring-framework-reference/htmlsingle/#beans-c-namespace
-->
<beans:bean id="frameOptionsWriter"
  class="org.springframework.security.web.header.writers.frameoptions.XFrameOptionsHeaderWriter"
  c:frameOptionsMode="SAMEORIGIN"/>
```

We could also restrict framing of content to the same origin with Java configuration:

```
@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .addHeaderWriter(new XFrameOptionsHeaderWriter(XFrameOptionsMode.SAMEORIGIN))
            );
    }
}
```

DelegatingRequestMatcherHeaderWriter

At times you may want to only write a header for certain requests. For example, perhaps you want to only protect your log in page from being framed. You could use the `DelegatingRequestMatcherHeaderWriter` to do so. When using the XML namespace configuration, this can be done with the following:

```

<http>
  <!-- ... -->

  <headers>
    <frame-options disabled="true"/>
    <header ref="headerWriter"/>
  </headers>
</http>

<beans:bean id="headerWriter"
  class="org.springframework.security.web.header.writers.DelegatingRequestMatcherHeaderWriter">
  <beans:constructor-arg>
    <bean class="org.springframework.security.web.util.matcher.AntPathRequestMatcher"
      c:pattern="/login"/>
    </beans:constructor-arg>
  <beans:constructor-arg>
    <beans:bean
      class="org.springframework.security.web.header.writers.frameoptions.XFrameOptionsHeaderWriter"/>
    </beans:constructor-arg>
  </beans:bean>

```

We could also prevent framing of content to the log in page using java configuration:

```

@EnableWebSecurity
public class WebSecurityConfig extends
WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        RequestMatcher matcher = new AntPathRequestMatcher("/login");
        DelegatingRequestMatcherHeaderWriter headerWriter =
            new DelegatingRequestMatcherHeaderWriter(matcher, new XFrameOptionsHeaderWriter());
        http
            // ...
            .headers(headers ->
                headers
                    .frameOptions(frameOptions ->
                        frameOptions.disable()
                    )
                    .addHeaderWriter(headerWriter)
            );
    }
}

```

14.3 HTTPS

Adding HTTP/HTTPS Channel Security

If your application supports both HTTP and HTTPS, and you require that particular URLs can only be accessed over HTTPS, then this is directly supported using the `requires-channel` attribute on `<intercept-url>`:

```

<http>
<intercept-url pattern="/secure/**" access="ROLE_USER" requires-channel="https"/>
<intercept-url pattern="/**" access="ROLE_USER" requires-channel="any"/>
...
</http>

```

With this configuration in place, if a user attempts to access anything matching the `/secure/**` pattern using HTTP, they will first be redirected to an HTTPS URL ⁴⁴. The available options are "http", "https" or "any". Using the value "any" means that either HTTP or HTTPS can be used.

⁴⁴For more details on how channel-processing is implemented, see the Javadoc for `ChannelProcessingFilter` and related classes.

If your application uses non-standard ports for HTTP and/or HTTPS, you can specify a list of port mappings as follows:

```
<http>
...
<port-mappings>
  <port-mapping http="9080" https="9443"/>
</port-mappings>
</http>
```

Note that in order to be truly secure, an application should not use HTTP at all or switch between HTTP and HTTPS. It should start in HTTPS (with the user entering an HTTPS URL) and use a secure connection throughout to avoid any possibility of man-in-the-middle attacks.

15. Integrations

15.1 Servlet API integration

This section describes how Spring Security is integrated with the Servlet API. The [servletapi-xml](#) sample application demonstrates the usage of each of these methods.

Servlet 2.5+ Integration

`HttpServletRequest.getRemoteUser()`

The [HttpServletRequest.getRemoteUser\(\)](#) will return the result of `SecurityContextHolder.getContext().getAuthentication().getName()` which is typically the current username. This can be useful if you want to display the current username in your application. Additionally, checking if this is null can be used to indicate if a user has authenticated or is anonymous. Knowing if the user is authenticated or not can be useful for determining if certain UI elements should be shown or not (i.e. a log out link should only be displayed if the user is authenticated).

`HttpServletRequest.getUserPrincipal()`

The [HttpServletRequest.getUserPrincipal\(\)](#) will return the result of `SecurityContextHolder.getContext().getAuthentication()`. This means it is an `Authentication` which is typically an instance of `UsernamePasswordAuthenticationToken` when using username and password based authentication. This can be useful if you need additional information about your user. For example, you might have created a custom `UserDetailsService` that returns a custom `UserDetails` containing a first and last name for your user. You could obtain this information with the following:

```
Authentication auth = httpServletRequest.getUserPrincipal();
// assume integrated custom UserDetails called MyCustomUserDetails
// by default, typically instance of UserDetails
MyCustomUserDetails userDetails = (MyCustomUserDetails) auth.getPrincipal();
String firstName = userDetails.getFirstName();
String lastName = userDetails.getLastName();
```

Note

It should be noted that it is typically bad practice to perform so much logic throughout your application. Instead, one should centralize it to reduce any coupling of Spring Security and the Servlet API's.

`HttpServletRequest.isUserInRole(String)`

The [HttpServletRequest.isUserInRole\(String\)](#) will determine if `SecurityContextHolder.getContext().getAuthentication().getAuthorities()` contains a `GrantedAuthority` with the role passed into `isUserInRole(String)`. Typically users should not pass in the "ROLE_" prefix into this method since it is added automatically. For example, if you want to determine if the current user has the authority "ROLE_ADMIN", you could use the following:

```
boolean isAdmin = httpServletRequest.isUserInRole("ADMIN");
```

This might be useful to determine if certain UI components should be displayed. For example, you might display admin links only if the current user is an admin.

Servlet 3+ Integration

The following section describes the Servlet 3 methods that Spring Security integrates with.

`HttpServletRequest.authenticate(HttpServletRequest, HttpServletResponse)`

The [HttpServletRequest.authenticate\(HttpServletRequest, HttpServletResponse\)](#) method can be used to ensure that a user is authenticated. If they are not authenticated, the configured `AuthenticationEntryPoint` will be used to request the user to authenticate (i.e. redirect to the login page).

`HttpServletRequest.login(String, String)`

The [HttpServletRequest.login\(String, String\)](#) method can be used to authenticate the user with the current `AuthenticationManager`. For example, the following would attempt to authenticate with the username "user" and password "password":

```
try {
    httpServletRequest.login("user", "password");
} catch (ServletException e) {
    // fail to authenticate
}
```

Note

It is not necessary to catch the `ServletException` if you want Spring Security to process the failed authentication attempt.

`HttpServletRequest.logout()`

The [HttpServletRequest.logout\(\)](#) method can be used to log the current user out.

Typically this means that the `SecurityContextHolder` will be cleared out, the `HttpSession` will be invalidated, any "Remember Me" authentication will be cleaned up, etc. However, the configured `LogoutHandler` implementations will vary depending on your Spring Security configuration. It is important to note that after `HttpServletRequest.logout()` has been invoked, you are still in charge of writing a response out. Typically this would involve a redirect to the welcome page.

`AsyncContext.start(Runnable)`

The [AsyncContext.start\(Runnable\)](#) method that ensures your credentials will be propagated to the new Thread. Using Spring Security's concurrency support, Spring Security overrides the `AsyncContext.start(Runnable)` to ensure that the current `SecurityContext` is used when processing the `Runnable`. For example, the following would output the current user's Authentication:

```
final AsyncContext async = httpServletRequest.startAsync();
async.start(new Runnable() {
    public void run() {
        Authentication authentication = SecurityContextHolder.getContext().getAuthentication();
        try {
            final HttpServletResponse asyncResponse = (HttpServletResponse) async.getResponse();
            asyncResponse.setStatus(HttpServletResponse.SC_OK);
            asyncResponse.getWriter().write(String.valueOf(authentication));
            async.complete();
        } catch (Exception e) {
            throw new RuntimeException(e);
        }
    }
});
```

Async Servlet Support

If you are using Java Based configuration, you are ready to go. If you are using XML configuration, there are a few updates that are necessary. The first step is to ensure you have updated your web.xml to use at least the 3.0 schema as shown below:

```
<web-app xmlns="http://java.sun.com/xml/ns/javaee"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://java.sun.com/xml/ns/javaee https://java.sun.com/xml/ns/javaee/web-
app_3_0.xsd"
version="3.0">

</web-app>
```

Next you need to ensure that your `springSecurityFilterChain` is setup for processing asynchronous requests.

```
<filter>
<filter-name>springSecurityFilterChain</filter-name>
<filter-class>
    org.springframework.web.filter.DelegatingFilterProxy
</filter-class>
<async-supported>true</async-supported>
</filter>
<filter-mapping>
<filter-name>springSecurityFilterChain</filter-name>
<url-pattern>/*</url-pattern>
<dispatcher>REQUEST</dispatcher>
<dispatcher>ASYNC</dispatcher>
</filter-mapping>
```

That's it! Now Spring Security will ensure that your `SecurityContext` is propagated on asynchronous requests too.

So how does it work? If you are not really interested, feel free to skip the remainder of this section, otherwise read on. Most of this is built into the Servlet specification, but there is a little bit of tweaking that Spring Security does to ensure things work with asynchronous requests properly. Prior to Spring Security 3.2, the `SecurityContext` from the `SecurityContextHolder` was automatically saved as soon as the `HttpServletResponse` was committed. This can cause issues in an Async environment. For example, consider the following:

```
HttpServletRequest.startAsync();
new Thread("AsyncThread") {
    @Override
    public void run() {
        try {
            // Do work
            TimeUnit.SECONDS.sleep(1);

            // Write to and commit the HttpServletResponse
            HttpServletResponse.getOutputStream().flush();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}.start();
```

The issue is that this Thread is not known to Spring Security, so the `SecurityContext` is not propagated to it. This means when we commit the `HttpServletResponse` there is no `SecurityContext`. When Spring Security automatically saved the `SecurityContext` on committing the `HttpServletResponse` it would lose our logged in user.

Since version 3.2, Spring Security is smart enough to no longer automatically save the `SecurityContext` on committing the `HttpServletResponse` as soon as `HttpServletRequest.startAsync()` is invoked.

Servlet 3.1+ Integration

The following section describes the Servlet 3.1 methods that Spring Security integrates with.

`HttpServletRequest#changeSessionId()`

The [HttpServletRequest.changeSessionId\(\)](#) is the default method for protecting against [Session Fixation](#) attacks in Servlet 3.1 and higher.

15.2 Spring Data Integration

Spring Security provides Spring Data integration that allows referring to the current user within your queries. It is not only useful but necessary to include the user in the queries to support paged results since filtering the results afterwards would not scale.

Spring Data & Spring Security Configuration

To use this support, add `org.springframework.security:spring-security-data` dependency and provide a bean of type `SecurityEvaluationContextExtension`. In Java Configuration, this would look like:

```
@Bean
public SecurityEvaluationContextExtension securityEvaluationContextExtension() {
    return new SecurityEvaluationContextExtension();
}
```

In XML Configuration, this would look like:

```
<bean class="org.springframework.security.data.repository.query.SecurityEvaluationContextExtension"/>
```

Security Expressions within @Query

Now Spring Security can be used within your queries. For example:

```
@Repository
public interface MessageRepository extends PagingAndSortingRepository<Message, Long> {
    @Query("select m from Message m where m.to.id = ?#{ principal?.id }")
    Page<Message> findInbox(Pageable pageable);
}
```

This checks to see if the `Authentication.getPrincipal().getId()` is equal to the recipient of the `Message`. Note that this example assumes you have customized the principal to be an Object that has an `id` property. By exposing the `SecurityEvaluationContextExtension` bean, all of the [Common Security Expressions](#) are available within the Query.

15.3 Concurrency Support

In most environments, Security is stored on a per Thread basis. This means that when work is done on a new Thread, the `SecurityContext` is lost. Spring Security provides some infrastructure to help make this much easier for users. Spring Security provides low level abstractions for working with Spring Security in multi-threaded environments. In fact, this is what Spring Security builds on to integration with the section called “`AsyncContext.start(Runnable)`” and the section called “Spring MVC Async Integration”.

DelegatingSecurityContextRunnable

One of the most fundamental building blocks within Spring Security's concurrency support is the `DelegatingSecurityContextRunnable`. It wraps a delegate `Runnable` in order to initialize the `SecurityContextHolder` with a specified `SecurityContext` for the delegate. It then invokes the delegate `Runnable` ensuring to clear the `SecurityContextHolder` afterwards. The `DelegatingSecurityContextRunnable` looks something like this:

```
public void run() {
    try {
        SecurityContextHolder.setContext(securityContext);
        delegate.run();
    } finally {
        SecurityContextHolder.clearContext();
    }
}
```

While very simple, it makes it seamless to transfer the `SecurityContext` from one `Thread` to another. This is important since, in most cases, the `SecurityContextHolder` acts on a per `Thread` basis. For example, you might have used Spring Security's the section called “<global-method-security>” support to secure one of your services. You can now easily transfer the `SecurityContext` of the current `Thread` to the `Thread` that invokes the secured service. An example of how you might do this can be found below:

```
Runnable originalRunnable = new Runnable() {
    public void run() {
        // invoke secured service
    }
};

SecurityContext context = SecurityContextHolder.getContext();
DelegatingSecurityContextRunnable wrappedRunnable =
    new DelegatingSecurityContextRunnable(originalRunnable, context);

new Thread(wrappedRunnable).start();
```

The code above performs the following steps:

- Creates a `Runnable` that will be invoking our secured service. Notice that it is not aware of Spring Security
- Obtains the `SecurityContext` that we wish to use from the `SecurityContextHolder` and initializes the `DelegatingSecurityContextRunnable`
- Use the `DelegatingSecurityContextRunnable` to create a `Thread`
- Start the `Thread` we created

Since it is quite common to create a `DelegatingSecurityContextRunnable` with the `SecurityContext` from the `SecurityContextHolder` there is a shortcut constructor for it. The following code is the same as the code above:

```
Runnable originalRunnable = new Runnable() {
    public void run() {
        // invoke secured service
    }
};

DelegatingSecurityContextRunnable wrappedRunnable =
    new DelegatingSecurityContextRunnable(originalRunnable);

new Thread(wrappedRunnable).start();
```

The code we have is simple to use, but it still requires knowledge that we are using Spring Security. In the next section we will take a look at how we can utilize `DelegatingSecurityContextExecutor` to hide the fact that we are using Spring Security.

DelegatingSecurityContextExecutor

In the previous section we found that it was easy to use the `DelegatingSecurityContextRunnable`, but it was not ideal since we had to be aware of Spring Security in order to use it. Let's take a look at how `DelegatingSecurityContextExecutor` can shield our code from any knowledge that we are using Spring Security.

The design of `DelegatingSecurityContextExecutor` is very similar to that of `DelegatingSecurityContextRunnable` except it accepts a delegate `Executor` instead of a delegate `Runnable`. You can see an example of how it might be used below:

```
SecurityContext context = SecurityContextHolder.createEmptyContext();
Authentication authentication =
    new UsernamePasswordAuthenticationToken("user", "doesnotmatter",
        AuthorityUtils.createAuthorityList("ROLE_USER"));
context.setAuthentication(authentication);

SimpleAsyncTaskExecutor delegateExecutor =
    new SimpleAsyncTaskExecutor();
DelegatingSecurityContextExecutor executor =
    new DelegatingSecurityContextExecutor(delegateExecutor, context);

Runnable originalRunnable = new Runnable() {
    public void run() {
        // invoke secured service
    }
};

executor.execute(originalRunnable);
```

The code performs the following steps:

- Creates the `SecurityContext` to be used for our `DelegatingSecurityContextExecutor`. Note that in this example we simply create the `SecurityContext` by hand. However, it does not matter where or how we get the `SecurityContext` (i.e. we could obtain it from the `SecurityContextHolder` if we wanted).
- Creates a `delegateExecutor` that is in charge of executing submitted `Runnable`s
- Finally we create a `DelegatingSecurityContextExecutor` which is in charge of wrapping any `Runnable` that is passed into the `execute` method with a `DelegatingSecurityContextRunnable`. It then passes the wrapped `Runnable` to the `delegateExecutor`. In this instance, the same `SecurityContext` will be used for every `Runnable` submitted to our `DelegatingSecurityContextExecutor`. This is nice if we are running background tasks that need to be run by a user with elevated privileges.
- At this point you may be asking yourself "How does this shield my code of any knowledge of Spring Security?" Instead of creating the `SecurityContext` and the `DelegatingSecurityContextExecutor` in our own code, we can inject an already initialized instance of `DelegatingSecurityContextExecutor`.

```

@Autowired
private Executor executor; // becomes an instance of our DelegatingSecurityContextExecutor

public void submitRunnable() {
    Runnable originalRunnable = new Runnable() {
        public void run() {
            // invoke secured service
        }
    };
    executor.execute(originalRunnable);
}

```

Now our code is unaware that the `SecurityContext` is being propagated to the `Thread`, then the `originalRunnable` is executed, and then the `SecurityContextHolder` is cleared out. In this example, the same user is being used to execute each `Thread`. What if we wanted to use the user from `SecurityContextHolder` at the time we invoked `executor.execute(Runnable)` (i.e. the currently logged in user) to process `originalRunnable`? This can be done by removing the `SecurityContext` argument from our `DelegatingSecurityContextExecutor` constructor. For example:

```

SimpleAsyncTaskExecutor delegateExecutor = new SimpleAsyncTaskExecutor();
DelegatingSecurityContextExecutor executor =
    new DelegatingSecurityContextExecutor(delegateExecutor);

```

Now anytime `executor.execute(Runnable)` is executed the `SecurityContext` is first obtained by the `SecurityContextHolder` and then that `SecurityContext` is used to create our `DelegatingSecurityContextRunnable`. This means that we are executing our `Runnable` with the same user that was used to invoke the `executor.execute(Runnable)` code.

Spring Security Concurrency Classes

Refer to the Javadoc for additional integrations with both the Java concurrent APIs and the Spring Task abstractions. They are quite self-explanatory once you understand the previous code.

- `DelegatingSecurityContextCallable`
- `DelegatingSecurityContextExecutor`
- `DelegatingSecurityContextExecutorService`
- `DelegatingSecurityContextRunnable`
- `DelegatingSecurityContextScheduledExecutorService`
- `DelegatingSecurityContextSchedulingTaskExecutor`
- `DelegatingSecurityContextAsyncTaskExecutor`
- `DelegatingSecurityContextTaskExecutor`
- `DelegatingSecurityContextTaskScheduler`

15.4 Jackson Support

Spring Security has added Jackson Support for persisting Spring Security related classes. This can improve the performance of serializing Spring Security related classes when working with distributed sessions (i.e. session replication, Spring Session, etc).

To use it, register the `SecurityJackson2Modules.getModules(ClassLoader)` as [Jackson Modules](#).

```

ObjectMapper mapper = new ObjectMapper();
ClassLoader loader = getClass().getClassLoader();
List<Module> modules = SecurityJackson2Modules.getModules(loader);
mapper.registerModules(modules);

// ... use ObjectMapper as normally ...
SecurityContext context = new SecurityContextImpl();
// ...
String json = mapper.writeValueAsString(context);

```

15.5 Localization

Spring Security supports localization of exception messages that end users are likely to see. If your application is designed for English-speaking users, you don't need to do anything as by default all Security messages are in English. If you need to support other locales, everything you need to know is contained in this section.

All exception messages can be localized, including messages related to authentication failures and access being denied (authorization failures). Exceptions and logging messages that are focused on developers or system deployers (including incorrect attributes, interface contract violations, using incorrect constructors, startup time validation, debug-level logging) are not localized and instead are hard-coded in English within Spring Security's code.

Shipping in the `spring-security-core-xx.jar` you will find an `org.springframework.security` package that in turn contains a `messages.properties` file, as well as localized versions for some common languages. This should be referred to by your `ApplicationContext`, as Spring Security classes implement Spring's `MessageSourceAware` interface and expect the message resolver to be dependency injected at application context startup time. Usually all you need to do is register a bean inside your application context to refer to the messages. An example is shown below:

```

<bean id="messageSource"
      class="org.springframework.context.support.ReloadableResourceBundleMessageSource">
  <property name="basename" value="classpath:org/springframework/security/messages"/>
</bean>

```

The `messages.properties` is named in accordance with standard resource bundles and represents the default language supported by Spring Security messages. This default file is in English.

If you wish to customize the `messages.properties` file, or support other languages, you should copy the file, rename it accordingly, and register it inside the above bean definition. There are not a large number of message keys inside this file, so localization should not be considered a major initiative. If you do perform localization of this file, please consider sharing your work with the community by logging a JIRA task and attaching your appropriately-named localized version of `messages.properties`.

Spring Security relies on Spring's localization support in order to actually lookup the appropriate message. In order for this to work, you have to make sure that the locale from the incoming request is stored in Spring's `org.springframework.context.i18n.LocaleContextHolder`. Spring MVC's `DispatcherServlet` does this for your application automatically, but since Spring Security's filters are invoked before this, the `LocaleContextHolder` needs to be set up to contain the correct `Locale` before the filters are called. You can either do this in a filter yourself (which must come before the Spring Security filters in `web.xml`) or you can use Spring's `RequestContextFilter`. Please refer to the Spring Framework documentation for further details on using localization with Spring.

The "contacts" sample application is set up to use localized messages.

15.6 Spring MVC Integration

Spring Security provides a number of optional integrations with Spring MVC. This section covers the integration in further detail.

@EnableWebMvcSecurity

Note

As of Spring Security 4.0, `@EnableWebMvcSecurity` is deprecated. The replacement is `@EnableWebSecurity` which will determine adding the Spring MVC features based upon the classpath.

To enable Spring Security integration with Spring MVC add the `@EnableWebSecurity` annotation to your configuration.

Note

Spring Security provides the configuration using Spring MVC's [WebMvcConfigurer](#). This means that if you are using more advanced options, like integrating with `WebMvcConfigurationSupport` directly, then you will need to manually provide the Spring Security configuration.

MvcRequestMatcher

Spring Security provides deep integration with how Spring MVC matches on URLs with `MvcRequestMatcher`. This is helpful to ensure your Security rules match the logic used to handle your requests.

In order to use `MvcRequestMatcher` you must place the Spring Security Configuration in the same `ApplicationContext` as your `DispatcherServlet`. This is necessary because Spring Security's `MvcRequestMatcher` expects a `HandlerMappingIntrospector` bean with the name of `mvcHandlerMappingIntrospector` to be registered by your Spring MVC configuration that is used to perform the matching.

For a `web.xml` this means that you should place your configuration in the `DispatcherServlet.xml`.

```

<listener>
  <listener-class>org.springframework.web.context.ContextLoaderListener</listener-class>
</listener>

<!-- All Spring Configuration (both MVC and Security) are in /WEB-INF/spring/ -->
<context-param>
  <param-name>contextConfigLocation</param-name>
  <param-value>/WEB-INF/spring/*.xml</param-value>
</context-param>

<servlet>
  <servlet-name>spring</servlet-name>
  <servlet-class>org.springframework.web.servlet.DispatcherServlet</servlet-class>
  <!-- Load from the ContextLoaderListener -->
  <init-param>
    <param-name>contextConfigLocation</param-name>
    <param-value></param-value>
  </init-param>
</servlet>

<servlet-mapping>
  <servlet-name>spring</servlet-name>
  <url-pattern>/</url-pattern>
</servlet-mapping>

```

Below `WebSecurityConfiguration` is placed in the `DispatcherServlets` `ApplicationContext`.

```

public class SecurityInitializer extends
    AbstractAnnotationConfigDispatcherServletInitializer {

    @Override
    protected Class<?>[] getRootConfigClasses() {
        return null;
    }

    @Override
    protected Class<?>[] getServletConfigClasses() {
        return new Class[] { RootConfiguration.class,
            WebMvcConfiguration.class };
    }

    @Override
    protected String[] getServletMappings() {
        return new String[] { "/" };
    }
}

```

Note

It is always recommended to provide authorization rules by matching on the `HttpServletRequest` and method security.

Providing authorization rules by matching on `HttpServletRequest` is good because it happens very early in the code path and helps reduce the [attack surface](#). Method security ensures that if someone has bypassed the web authorization rules, that your application is still secured. This is what is known as [Defence in Depth](#)

Consider a controller that is mapped as follows:

```

@RequestMapping("/admin")
public String admin() {

```

If we wanted to restrict access to this controller method to admin users, a developer can provide authorization rules by matching on the `HttpServletRequest` with the following:

```
protected configure(HttpSecurity http) throws Exception {
    http
        .authorizeRequests(authorizeRequests ->
            authorizeRequests
                .antMatchers("/admin").hasRole("ADMIN")
        );
}
```

or in XML

```
<http>
  <intercept-url pattern="/admin" access="hasRole('ADMIN')"/>
</http>
```

With either configuration, the URL `/admin` will require the authenticated user to be an admin user. However, depending on our Spring MVC configuration, the URL `/admin.html` will also map to our `admin()` method. Additionally, depending on our Spring MVC configuration, the URL `/admin/` will also map to our `admin()` method.

The problem is that our security rule is only protecting `/admin`. We could add additional rules for all the permutations of Spring MVC, but this would be quite verbose and tedious.

Instead, we can leverage Spring Security's `MvcRequestMatcher`. The following configuration will protect the same URLs that Spring MVC will match on by using Spring MVC to match on the URL.

```
protected configure(HttpSecurity http) throws Exception {
    http
        .authorizeRequests(authorizeRequests ->
            authorizeRequests
                .mvcMatchers("/admin").hasRole("ADMIN")
        );
}
```

or in XML

```
<http request-matcher="mvc">
  <intercept-url pattern="/admin" access="hasRole('ADMIN')"/>
</http>
```

@AuthenticationPrincipal

Spring Security provides `AuthenticationPrincipalArgumentResolver` which can automatically resolve the current `Authentication.getPrincipal()` for Spring MVC arguments. By using `@EnableWebSecurity` you will automatically have this added to your Spring MVC configuration. If you use XML based configuration, you must add this yourself. For example:

```
<mvc:annotation-driven>
  <mvc:argument-resolvers>

  <bean class="org.springframework.security.web.method.annotation.AuthenticationPrincipalArgumentResolver"
  />

  </mvc:argument-resolvers>
</mvc:annotation-driven>
```

Once `AuthenticationPrincipalArgumentResolver` is properly configured, you can be entirely decoupled from Spring Security in your Spring MVC layer.

Consider a situation where a custom `UserDetailsService` that returns an `Object` that implements `UserDetails` and your own `CustomUser` `Object`. The `CustomUser` of the currently authenticated user could be accessed using the following code:

```
@RequestMapping("/messages/inbox")
public ModelAndView findMessagesForUser() {
    Authentication authentication =
        SecurityContextHolder.getContext().getAuthentication();
    CustomUser custom = (CustomUser) authentication == null ? null : authentication.getPrincipal();

    // .. find messages for this user and return them ...
}
```

As of Spring Security 3.2 we can resolve the argument more directly by adding an annotation. For example:

```
import org.springframework.security.core.annotation.AuthenticationPrincipal;

// ...

@RequestMapping("/messages/inbox")
public ModelAndView findMessagesForUser(@AuthenticationPrincipal CustomUser customUser) {

    // .. find messages for this user and return them ...
}
```

Sometimes it may be necessary to transform the principal in some way. For example, if `CustomUser` needed to be final it could not be extended. In this situation the `UserDetailsService` might return an `Object` that implements `UserDetails` and provides a method named `getCustomUser` to access `CustomUser`. For example, it might look like:

```
public class CustomUserUserDetails extends User {
    // ...
    public CustomUser getCustomUser() {
        return customUser;
    }
}
```

We could then access the `CustomUser` using a [SpEL expression](#) that uses `Authentication.getPrincipal()` as the root object:

```
import org.springframework.security.core.annotation.AuthenticationPrincipal;

// ...

@RequestMapping("/messages/inbox")
public ModelAndView findMessagesForUser(@AuthenticationPrincipal(expression = "customUser") CustomUser
    customUser) {

    // .. find messages for this user and return them ...
}
```

We can also refer to Beans in our SpEL expressions. For example, the following could be used if we were using JPA to manage our Users and we wanted to modify and save a property on the current user.

```
import org.springframework.security.core.annotation.AuthenticationPrincipal;

// ...

@PutMapping("/users/self")
public ModelAndView updateName(@AuthenticationPrincipal(expression = "@jpaEntityManager.merge(#this)")
    CustomUser attachedCustomUser,
    @RequestParam String firstName) {

    // change the firstName on an attached instance which will be persisted to the database
    attachedCustomUser.setFirstName(firstName);

    // ...
}
```

We can further remove our dependency on Spring Security by making `@AuthenticationPrincipal` a meta annotation on our own annotation. Below we demonstrate how we could do this on an annotation named `@CurrentUser`.

Note

It is important to realize that in order to remove the dependency on Spring Security, it is the consuming application that would create `@CurrentUser`. This step is not strictly required, but assists in isolating your dependency to Spring Security to a more central location.

```
@Target({ElementType.PARAMETER, ElementType.TYPE})
@Retention(RetentionPolicy.RUNTIME)
@Documented
@AuthenticationPrincipal
public @interface CurrentUser {}
```

Now that `@CurrentUser` has been specified, we can use it to signal to resolve our `CustomUser` of the currently authenticated user. We have also isolated our dependency on Spring Security to a single file.

```
@RequestMapping("/messages/inbox")
public ModelAndView findMessagesForUser(@CurrentUser CustomUser customUser) {

    // .. find messages for this user and return them ...

}
```

Spring MVC Async Integration

Spring Web MVC 3.2+ has excellent support for [Asynchronous Request Processing](#). With no additional configuration, Spring Security will automatically setup the `SecurityContext` to the `Thread` that executes a `Callable` returned by your controllers. For example, the following method will automatically have its `Callable` executed with the `SecurityContext` that was available when the `Callable` was created:

```
@RequestMapping(method=RequestMethod.POST)
public Callable<String> processUpload(final MultipartFile file) {

    return new Callable<String>() {
        public Object call() throws Exception {
            // ...
            return "someView";
        }
    };
}
```

Associating SecurityContext to Callable's

More technically speaking, Spring Security integrates with `WebAsyncManager`. The `SecurityContext` that is used to process the `Callable` is the `SecurityContext` that exists on the `SecurityContextHolder` at the time `startCallableProcessing` is invoked.

There is no automatic integration with a `DeferredResult` that is returned by controllers. This is because `DeferredResult` is processed by the users and thus there is no way of automatically integrating with it. However, you can still use [Concurrency Support](#) to provide transparent integration with Spring Security.

Spring MVC and CSRF Integration

Automatic Token Inclusion

Spring Security will automatically [include the CSRF Token](#) within forms that use the [Spring MVC form tag](#). For example, the following JSP:

```
<jsp:root xmlns:jsp="http://java.sun.com/JSP/Page"
  xmlns:c="http://java.sun.com/jsp/jstl/core"
  xmlns:form="http://www.springframework.org/tags/form" version="2.0">
  <jsp:directive.page language="java" contentType="text/html" />
  <html xmlns="http://www.w3.org/1999/xhtml" lang="en" xml:lang="en">
    <!-- ... -->

    <c:url var="logoutUrl" value="/logout"/>
    <form:form action="${logoutUrl}"
      method="post">
      <input type="submit"
        value="Log out" />
      <input type="hidden"
        name="${_csrf.parameterName}"
        value="${_csrf.token}"/>
    </form:form>

    <!-- ... -->
  </html>
</jsp:root>
```

Will output HTML that is similar to the following:

```
<!-- ... -->

<form action="/context/logout" method="post">
<input type="submit" value="Log out"/>
<input type="hidden" name="_csrf" value="f81d4fae-7dec-11d0-a765-00a0c91e6bf6"/>
</form>

<!-- ... -->
```

Resolving the CsrfToken

Spring Security provides `CsrfTokenArgumentResolver` which can automatically resolve the current `CsrfToken` for Spring MVC arguments. By using [@EnableWebSecurity](#) you will automatically have this added to your Spring MVC configuration. If you use XML based configuration, you must add this yourself.

Once `CsrfTokenArgumentResolver` is properly configured, you can expose the `CsrfToken` to your static HTML based application.

```

@RestController
public class CsrfController {

    @RequestMapping("/csrf")
    public CsrfToken csrf(CsrfToken token) {
        return token;
    }
}

```

It is important to keep the `CsrfToken` a secret from other domains. This means if you are using [Cross Origin Sharing \(CORS\)](#), you should **NOT** expose the `CsrfToken` to any external domains.

15.7 WebSocket Security

Spring Security 4 added support for securing [Spring's WebSocket support](#). This section describes how to use Spring Security's WebSocket support.

Note

You can find a complete working sample of WebSocket security at <https://github.com/spring-projects/spring-session/tree/master/samples/boot/websocket>.

Direct JSR-356 Support

Spring Security does not provide direct JSR-356 support because doing so would provide little value. This is because the format is unknown, so there is [little Spring can do to secure an unknown format](#). Additionally, JSR-356 does not provide a way to intercept messages, so security would be rather invasive.

WebSocket Configuration

Spring Security 4.0 has introduced authorization support for WebSockets through the Spring Messaging abstraction. To configure authorization using Java Configuration, simply extend the `AbstractSecurityWebSocketMessageBrokerConfigurer` and configure the `MessageSecurityMetadataSourceRegistry`. For example:

```

@Configuration
public class WebSocketSecurityConfig
    extends AbstractSecurityWebSocketMessageBrokerConfigurer { ❶ ❷

    protected void configureInbound(MessageSecurityMetadataSourceRegistry messages) {
        messages
            .simpDestMatchers("/user/**").authenticated() ❸
    }
}

```

This will ensure that:

- ❶ Any inbound CONNECT message requires a valid CSRF token to enforce [Same Origin Policy](#)
- ❷ The `SecurityContextHolder` is populated with the user within the `simpUser` header attribute for any inbound request.
- ❸ Our messages require the proper authorization. Specifically, any inbound message that starts with `/user/` will require `ROLE_USER`. Additional details on authorization can be found in the section called "WebSocket Authorization"

Spring Security also provides [XML Namespace](#) support for securing WebSockets. A comparable XML based configuration looks like the following:

```
<websocket-message-broker> ❶ ❷
  ❸
  <intercept-message pattern="/user/**" access="hasRole('USER')" />
</websocket-message-broker>
```

This will ensure that:

- ❶ Any inbound CONNECT message requires a valid CSRF token to enforce [Same Origin Policy](#)
- ❷ The SecurityContextHolder is populated with the user within the simpUser header attribute for any inbound request.
- ❸ Our messages require the proper authorization. Specifically, any inbound message that starts with "/user/" will require ROLE_USER. Additional details on authorization can be found in the section called "WebSocket Authorization"

WebSocket Authentication

WebSockets reuse the same authentication information that is found in the HTTP request when the WebSocket connection was made. This means that the `Principal` on the `HttpServletRequest` will be handed off to WebSockets. If you are using Spring Security, the `Principal` on the `HttpServletRequest` is overridden automatically.

More concretely, to ensure a user has authenticated to your WebSocket application, all that is necessary is to ensure that you setup Spring Security to authenticate your HTTP based web application.

WebSocket Authorization

Spring Security 4.0 has introduced authorization support for WebSockets through the Spring Messaging abstraction. To configure authorization using Java Configuration, simply extend the `AbstractSecurityWebSocketMessageBrokerConfigurer` and configure the `MessageSecurityMetadataSourceRegistry`. For example:

```
@Configuration
public class WebSocketSecurityConfig extends AbstractSecurityWebSocketMessageBrokerConfigurer {

    @Override
    protected void configureInbound(MessageSecurityMetadataSourceRegistry messages) {
        messages
            .nullDestMatcher().authenticated() ❶
            .simpSubscribeDestMatchers("/user/queue/errors").permitAll() ❷
            .simpDestMatchers("/app/**").hasRole("USER") ❸
            .simpSubscribeDestMatchers("/user/**", "/topic/friends/**").hasRole("USER") ❹
            .simpTypeMatchers(MESSAGE, SUBSCRIBE).denyAll() ❺
            .anyMessage().denyAll(); ❻
    }
}
```

This will ensure that:

- ❶ Any message without a destination (i.e. anything other than Message type of MESSAGE or SUBSCRIBE) will require the user to be authenticated
- ❷ Anyone can subscribe to /user/queue/errors
- ❸ Any message that has a destination starting with "/app/" will be require the user to have the role ROLE_USER

- ④ Any message that starts with "/user/" or "/topic/friends/" that is of type SUBSCRIBE will require ROLE_USER
- ⑤ Any other message of type MESSAGE or SUBSCRIBE is rejected. Due to 6 we do not need this step, but it illustrates how one can match on specific message types.
- ⑥ Any other Message is rejected. This is a good idea to ensure that you do not miss any messages.

Spring Security also provides [XML Namespace](#) support for securing WebSockets. A comparable XML based configuration looks like the following:

```

<websocket-message-broker>
  ①
  <intercept-message type="CONNECT" access="permitAll" />
  <intercept-message type="UNSUBSCRIBE" access="permitAll" />
  <intercept-message type="DISCONNECT" access="permitAll" />

  <intercept-message pattern="/user/queue/errors" type="SUBSCRIBE" access="permitAll" /> ②
  <intercept-message pattern="/app/**" access="hasRole('USER')" /> ③

  ④
  <intercept-message pattern="/user/**" access="hasRole('USER')" />
  <intercept-message pattern="/topic/friends/**" access="hasRole('USER')" />

  ⑤
  <intercept-message type="MESSAGE" access="denyAll" />
  <intercept-message type="SUBSCRIBE" access="denyAll" />

  <intercept-message pattern="/**" access="denyAll" /> ⑥
</websocket-message-broker>

```

This will ensure that:

- ① Any message of type CONNECT, UNSUBSCRIBE, or DISCONNECT will require the user to be authenticated
- ② Anyone can subscribe to /user/queue/errors
- ③ Any message that has a destination starting with "/app/" will be require the user to have the role ROLE_USER
- ④ Any message that starts with "/user/" or "/topic/friends/" that is of type SUBSCRIBE will require ROLE_USER
- ⑤ Any other message of type MESSAGE or SUBSCRIBE is rejected. Due to 6 we do not need this step, but it illustrates how one can match on specific message types.
- ⑥ Any other message with a destination is rejected. This is a good idea to ensure that you do not miss any messages.

WebSocket Authorization Notes

In order to properly secure your application it is important to understand Spring's WebSocket support.

WebSocket Authorization on Message Types

It is important to understand the distinction between SUBSCRIBE and MESSAGE types of messages and how it works within Spring.

Consider a chat application.

- The system can send notifications MESSAGE to all users through a destination of "/topic/system/notifications"
- Clients can receive notifications by SUBSCRIBE to the "/topic/system/notifications".

While we want clients to be able to SUBSCRIBE to `"/topic/system/notifications"`, we do not want to enable them to send a MESSAGE to that destination. If we allowed sending a MESSAGE to `"/topic/system/notifications"`, then clients could send a message directly to that endpoint and impersonate the system.

In general, it is common for applications to deny any MESSAGE sent to a destination that starts with the [broker prefix](#) (i.e. `"/topic/"` or `"/queue/"`).

WebSocket Authorization on Destinations

It is also important to understand how destinations are transformed.

Consider a chat application.

- Users can send messages to a specific user by sending a message to the destination of `"/app/chat"`.
- The application sees the message, ensures that the "from" attribute is specified as the current user (we cannot trust the client).
- The application then sends the message to the recipient using `SimpMessageSendingOperations.convertAndSendToUser("toUser", "/queue/messages", message)`.
- The message gets turned into the destination of `"/queue/user/messages-<sessionid>"`

With the application above, we want to allow our client to listen to `"/user/queue"` which is transformed into `"/queue/user/messages-<sessionid>"`. However, we do not want the client to be able to listen to `"/queue/*"` because that would allow the client to see messages for every user.

In general, it is common for applications to deny any SUBSCRIBE sent to a message that starts with the [broker prefix](#) (i.e. `"/topic/"` or `"/queue/"`). Of course we may provide exceptions to account for things like

Outbound Messages

Spring contains a section titled [Flow of Messages](#) that describes how messages flow through the system. It is important to note that Spring Security only secures the `clientInboundChannel`. Spring Security does not attempt to secure the `clientOutboundChannel`.

The most important reason for this is performance. For every message that goes in, there are typically many more that go out. Instead of securing the outbound messages, we encourage securing the subscription to the endpoints.

Enforcing Same Origin Policy

It is important to emphasize that the browser does not enforce the [Same Origin Policy](#) for WebSocket connections. This is an extremely important consideration.

Why Same Origin?

Consider the following scenario. A user visits `bank.com` and authenticates to their account. The same user opens another tab in their browser and visits `evil.com`. The Same Origin Policy ensures that `evil.com` cannot read or write data to `bank.com`.

With WebSockets the Same Origin Policy does not apply. In fact, unless `bank.com` explicitly forbids it, `evil.com` can read and write data on behalf of the user. This means that anything the user can do over the `websocket` (i.e. transfer money), `evil.com` can do on that users behalf.

Since SockJS tries to emulate WebSockets it also bypasses the Same Origin Policy. This means developers need to explicitly protect their applications from external domains when using SockJS.

Spring WebSocket Allowed Origin

Fortunately, since Spring 4.1.5 Spring's WebSocket and SockJS support restricts access to the [current domain](#). Spring Security adds an additional layer of protection to provide [defence in depth](#).

Adding CSRF to Stomp Headers

By default Spring Security requires the [CSRF token](#) in any CONNECT message type. This ensures that only a site that has access to the CSRF token can connect. Since only the **Same Origin** can access the CSRF token, external domains are not allowed to make a connection.

Typically we need to include the CSRF token in an HTTP header or an HTTP parameter. However, SockJS does not allow for these options. Instead, we must include the token in the Stomp headers

Applications can [obtain a CSRF token](#) by accessing the request attribute named `_csrf`. For example, the following will allow accessing the `CsrfToken` in a JSP:

```
var headerName = "${_csrf.headerName}";
var token = "${_csrf.token}";
```

If you are using static HTML, you can expose the `CsrfToken` on a REST endpoint. For example, the following would expose the `CsrfToken` on the URL `/csrf`

```
@RestController
public class CsrfController {

    @RequestMapping("/csrf")
    public CsrfToken csrf(CsrfToken token) {
        return token;
    }
}
```

The JavaScript can make a REST call to the endpoint and use the response to populate the `headerName` and the `token`.

We can now include the token in our Stomp client. For example:

```
...
var headers = {};
headers[headerName] = token;
stompClient.connect(headers, function(frame) {
    ...
})
```

Disable CSRF within WebSockets

If you want to allow other domains to access your site, you can disable Spring Security's protection. For example, in Java Configuration you can use the following:

```
@Configuration
public class WebSocketSecurityConfig extends AbstractSecurityWebSocketMessageBrokerConfigurer {

    ...

    @Override
    protected boolean sameOriginDisabled() {
        return true;
    }
}
```

Working with SockJS

[SockJS](#) provides fallback transports to support older browsers. When using the fallback options we need to relax a few security constraints to allow SockJS to work with Spring Security.

SockJS & frame-options

SockJS may use an [transport that leverages an iframe](#). By default Spring Security will [deny](#) the site from being framed to prevent Clickjacking attacks. To allow SockJS frame based transports to work, we need to configure Spring Security to allow the same origin to frame the content.

You can customize X-Frame-Options with the [frame-options](#) element. For example, the following will instruct Spring Security to use "X-Frame-Options: SAMEORIGIN" which allows iframes within the same domain:

```
<http>
  <!-- ... -->

  <headers>
    <frame-options
      policy="SAMEORIGIN" />
  </headers>
</http>
```

Similarly, you can customize frame options to use the same origin within Java Configuration using the following:

```
@EnableWebSecurity
public class WebSecurityConfig extends
    WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // ...
            .headers(headers ->
                headers
                    .frameOptions(frameOptions ->
                        frameOptions
                            .sameOrigin()
                    )
            )
        );
    }
}
```

SockJS & Relaxing CSRF

SockJS uses a POST on the CONNECT messages for any HTTP based transport. Typically we need to include the CSRF token in an HTTP header or an HTTP parameter. However, SockJS does not allow for these options. Instead, we must include the token in the Stomp headers as described in the section called "Adding CSRF to Stomp Headers".

It also means we need to relax our CSRF protection with the web layer. Specifically, we want to disable CSRF protection for our connect URLs. We do NOT want to disable CSRF protection for every URL. Otherwise our site will be vulnerable to CSRF attacks.

We can easily achieve this by providing a CSRF RequestMatcher. Our Java Configuration makes this extremely easy. For example, if our stomp endpoint is "/chat" we can disable CSRF protection for only URLs that start with "/chat/" using the following configuration:

```

@Configuration
@EnableWebSecurity
public class WebSecurityConfig
    extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .csrf(csrf ->
                csrf
                // ignore our stomp endpoints since they are protected using Stomp headers
                .ignoringAntMatchers("/chat/**")
            )
            .headers(headers ->
                headers
                // allow same origin to frame our site to support iframe SockJS
                .frameOptions(frameOptions ->
                    frameOptions
                    .sameOrigin()
                )
            )
            .authorizeRequests(authorizeRequests ->
                ...
            )
            ...
    }
}

```

If we are using XML based configuration, we can use the [csrf@request-matcher-ref](#). For example:

```

<http ...>
  <csrf request-matcher-ref="csrfMatcher"/>

  <headers>
    <frame-options policy="SAMEORIGIN"/>
  </headers>

  ...
</http>

<b:bean id="csrfMatcher"
  class="AndRequestMatcher">
  <b:constructor-arg
  arg value="#{T(org.springframework.security.web.csrf.CsrfFilter).DEFAULT_CSRF_MATCHER}"/>
  <b:constructor-arg>
    <b:bean class="org.springframework.security.web.util.matcher.NegatedRequestMatcher">
      <b:bean class="org.springframework.security.web.util.matcher.AntPathRequestMatcher">
        <b:constructor-arg value="/chat/**"/>
      </b:bean>
    </b:bean>
  </b:constructor-arg>
</b:bean>

```

15.8 CORS

Spring Framework provides [first class support for CORS](#). CORS must be processed before Spring Security because the pre-flight request will not contain any cookies (i.e. the `JSESSIONID`). If the request does not contain any cookies and Spring Security is first, the request will determine the user is not authenticated (since there are no cookies in the request) and reject it.

The easiest way to ensure that CORS is handled first is to use the `CorsFilter`. Users can integrate the `CorsFilter` with Spring Security by providing a `CorsConfigurationSource` using the following:

```

@EnableWebSecurity
public class WebSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // by default uses a Bean by the name of corsConfigurationSource
            .cors(withDefaults())
            ...
    }

    @Bean
    CorsConfigurationSource corsConfigurationSource() {
        CorsConfiguration configuration = new CorsConfiguration();
        configuration.setAllowedOrigins(Arrays.asList("https://example.com"));
        configuration.setAllowedMethods(Arrays.asList("GET", "POST"));
        UrlBasedCorsConfigurationSource source = new UrlBasedCorsConfigurationSource();
        source.registerCorsConfiguration("/**", configuration);
        return source;
    }
}

```

or in XML

```

<http>
  <cors configuration-source-ref="corsSource"/>
  ...
</http>
<b:bean id="corsSource" class="org.springframework.web.cors.UrlBasedCorsConfigurationSource">
  ...
</b:bean>

```

If you are using Spring MVC's CORS support, you can omit specifying the `CorsConfigurationSource` and Spring Security will leverage the CORS configuration provided to Spring MVC.

```

@EnableWebSecurity
public class WebSecurityConfig extends WebSecurityConfigurerAdapter {

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            // if Spring MVC is on classpath and no CorsConfigurationSource is provided,
            // Spring Security will use CORS configuration provided to Spring MVC
            .cors(withDefaults())
            ...
    }
}

```

or in XML

```

<http>
  <!-- Default to Spring MVC's CORS configuration -->
  <cors />
  ...
</http>

```

15.9 JSP Tag Libraries

Spring Security has its own taglib which provides basic support for accessing security information and applying security constraints in JSPs.

Declaring the Taglib

To use any of the tags, you must have the security taglib declared in your JSP:

```
<%@ taglib prefix="sec" uri="http://www.springframework.org/security/tags" %>
```

The authorize Tag

This tag is used to determine whether its contents should be evaluated or not. In Spring Security 3.0, it can be used in two ways²⁹. The first approach uses a [web-security expression](#), specified in the `access` attribute of the tag. The expression evaluation will be delegated to the `SecurityExpressionHandler<FilterInvocation>` defined in the application context (you should have web expressions enabled in your `<http>` namespace configuration to make sure this service is available). So, for example, you might have

```
<sec:authorize access="hasRole('supervisor')">
```

This content will only be visible to users who have the "supervisor" authority in their list of `<tt>GrantedAuthority</tt>`s.

```
</sec:authorize>
```

When used in conjunction with Spring Security's `PermissionEvaluator`, the tag can also be used to check permissions. For example:

```
<sec:authorize access="hasPermission(#domain,'read') or hasPermission(#domain,'write')">
```

This content will only be visible to users who have read or write permission to the Object found as a request attribute named "domain".

```
</sec:authorize>
```

A common requirement is to only show a particular link, if the user is actually allowed to click it. How can we determine in advance whether something will be allowed? This tag can also operate in an alternative mode which allows you to define a particular URL as an attribute. If the user is allowed to invoke that URL, then the tag body will be evaluated, otherwise it will be skipped. So you might have something like

```
<sec:authorize url="/admin">
```

This content will only be visible to users who are authorized to send requests to the "/admin" URL.

```
</sec:authorize>
```

To use this tag there must also be an instance of `WebInvocationPrivilegeEvaluator` in your application context. If you are using the namespace, one will automatically be registered. This is an instance of `DefaultWebInvocationPrivilegeEvaluator`, which creates a dummy web request for the supplied URL and invokes the security interceptor to see whether the request would succeed or fail. This allows you to delegate to the access-control setup you defined using `intercept-url` declarations within the `<http>` namespace configuration and saves having to duplicate the information (such as the required roles) within your JSPs. This approach can also be combined with a `method` attribute, supplying the HTTP method, for a more specific match.

The Boolean result of evaluating the tag (whether it grants or denies access) can be stored in a page context scope variable by setting the `var` attribute to the variable name, avoiding the need for duplicating and re-evaluating the condition at other points in the page.

Disabling Tag Authorization for Testing

Hiding a link in a page for unauthorized users doesn't prevent them from accessing the URL. They could just type it into their browser directly, for example. As part of your testing process, you may

²⁹The legacy options from Spring Security 2.0 are also supported, but discouraged.

want to reveal the hidden areas in order to check that links really are secured at the back end. If you set the system property `spring.security.disableUISecurity` to `true`, the `authorize` tag will still run but will not hide its contents. By default it will also surround the content with `...` tags. This allows you to display "hidden" content with a particular CSS style such as a different background colour. Try running the "tutorial" sample application with this property enabled, for example.

You can also set the properties `spring.security.securedUIPrefix` and `spring.security.securedUISuffix` if you want to change surrounding text from the default `span` tags (or use empty strings to remove it completely).

The authentication Tag

This tag allows access to the current `Authentication` object stored in the security context. It renders a property of the object directly in the JSP. So, for example, if the `principal` property of the `Authentication` is an instance of Spring Security's `UserDetails` object, then using `<sec:authentication property="principal.username" />` will render the name of the current user.

Of course, it isn't necessary to use JSP tags for this kind of thing and some people prefer to keep as little logic as possible in the view. You can access the `Authentication` object in your MVC controller (by calling `SecurityContextHolder.getContext().getAuthentication()`) and add the data directly to your model for rendering by the view.

The accesscontrollist Tag

This tag is only valid when used with Spring Security's ACL module. It checks a comma-separated list of required permissions for a specified domain object. If the current user has all of those permissions, then the tag body will be evaluated. If they don't, it will be skipped. An example might be

Caution

In general this tag should be considered deprecated. Instead use the the section called "The `authorize` Tag".

```
<sec:accesscontrollist hasPermission="1,2" domainObject="${someObject}">
```

This will be shown if the user has all of the permissions represented by the values "1" or "2" on the given object.

```
</sec:accesscontrollist>
```

The permissions are passed to the `PermissionFactory` defined in the application context, converting them to `ACL Permission` instances, so they may be any format which is supported by the factory - they don't have to be integers, they could be strings like `READ` or `WRITE`. If no `PermissionFactory` is found, an instance of `DefaultPermissionFactory` will be used. The `AclService` from the application context will be used to load the `Acl` instance for the supplied object. The `Acl` will be invoked with the required permissions to check if all of them are granted.

This tag also supports the `var` attribute, in the same way as the `authorize` tag.

The csrfInput Tag

If CSRF protection is enabled, this tag inserts a hidden form field with the correct name and value for the CSRF protection token. If CSRF protection is not enabled, this tag outputs nothing.

Normally Spring Security automatically inserts a CSRF form field for any `<form:form>` tags you use, but if for some reason you cannot use `<form:form>`, `csrfInput` is a handy replacement.

You should place this tag within an HTML `<form></form>` block, where you would normally place other input fields. Do NOT place this tag within a Spring `<form:form></form:form>` block. Spring Security handles Spring forms automatically.

```
<form method="post" action="/do/something">
  <sec:csrfInput />
  Name:<br />
  <input type="text" name="name" />
  ...
</form>
```

The csrfMetaTags Tag

If CSRF protection is enabled, this tag inserts meta tags containing the CSRF protection token form field and header names and CSRF protection token value. These meta tags are useful for employing CSRF protection within JavaScript in your applications.

You should place `csrfMetaTags` within an HTML `<head></head>` block, where you would normally place other meta tags. Once you use this tag, you can access the form field name, header name, and token value easily using JavaScript. JQuery is used in this example to make the task easier.

```

<!DOCTYPE html>
<html>
  <head>
    <title>CSRF Protected JavaScript Page</title>
    <meta name="description" content="This is the description for this page" />
    <sec:csrfMetaTags />
    <script type="text/javascript" language="javascript">

      var csrfParameter = $("meta[name='_csrf_parameter']").attr("content");
      var csrfHeader = $("meta[name='_csrf_header']").attr("content");
      var csrfToken = $("meta[name='_csrf']").attr("content");

      // using XMLHttpRequest directly to send an x-www-form-urlencoded request
      var ajax = new XMLHttpRequest();
      ajax.open("POST", "https://www.example.org/do/something", true);
      ajax.setRequestHeader("Content-Type", "application/x-www-form-urlencoded data");
      ajax.send(csrfParameter + "=" + csrfToken + "&name=John&...");

      // using XMLHttpRequest directly to send a non-x-www-form-urlencoded request
      var ajax = new XMLHttpRequest();
      ajax.open("POST", "https://www.example.org/do/something", true);
      ajax.setRequestHeader(csrfHeader, csrfToken);
      ajax.send("...");

      // using JQuery to send an x-www-form-urlencoded request
      var data = {};
      data[csrfParameter] = csrfToken;
      data["name"] = "John";
      ...
      $.ajax({
        url: "https://www.example.org/do/something",
        type: "POST",
        data: data,
        ...
      });

      // using JQuery to send a non-x-www-form-urlencoded request
      var headers = {};
      headers[csrfHeader] = csrfToken;
      $.ajax({
        url: "https://www.example.org/do/something",
        type: "POST",
        headers: headers,
        ...
      });

    </script>
  </head>
  <body>
    ...
  </body>
</html>

```

If CSRF protection is not enabled, `csrfMetaTags` outputs nothing.

16. Java Configuration

General support for [Java Configuration](#) was added to Spring Framework in Spring 3.1. Since Spring Security 3.2 there has been Spring Security Java Configuration support which enables users to easily configure Spring Security without the use of any XML.

If you are familiar with the Chapter 17, *Security Namespace Configuration* then you should find quite a few similarities between it and the Security Java Configuration support.

Note

Spring Security provides [lots of sample applications](#) which demonstrate the use of Spring Security Java Configuration.

16.1 Hello Web Security Java Configuration

The first step is to create our Spring Security Java Configuration. The configuration creates a Servlet Filter known as the `springSecurityFilterChain` which is responsible for all the security (protecting the application URLs, validating submitted username and passwords, redirecting to the log in form, etc) within your application. You can find the most basic example of a Spring Security Java Configuration below:

```
import org.springframework.beans.factory.annotation.Autowired;

import org.springframework.context.annotation.*;
import org.springframework.security.config.annotation.authentication.builders.*;
import org.springframework.security.config.annotation.web.configuration.*;

@EnableWebSecurity
public class WebSecurityConfig {

    @Bean
    public UserDetailsService userDetailsService() {
        InMemoryUserDetailsManager manager = new InMemoryUserDetailsManager();

        manager.createUser(User.withDefaultPasswordEncoder().username("user").password("password").roles("USER").build());
        return manager;
    }
}
```

There really isn't much to this configuration, but it does a lot. You can find a summary of the features below:

- Require authentication to every URL in your application
- Generate a login form for you
- Allow the user with the **Username** *user* and the **Password** *password* to authenticate with form based authentication
- Allow the user to logout
- [CSRF attack](#) prevention
- [Session Fixation](#) protection
- Security Header integration

- [HTTP Strict Transport Security](#) for secure requests
- [X-Content-Type-Options](#) integration
- Cache Control (can be overridden later by your application to allow caching of your static resources)
- [X-XSS-Protection](#) integration
- X-Frame-Options integration to help prevent [Clickjacking](#)
- Integrate with the following Servlet API methods
 - [HttpServletRequest#getRemoteUser\(\)](#)
 - [HttpServletRequest#getUserPrincipal\(\)](#)
 - [HttpServletRequest#isUserInRole\(java.lang.String\)](#)
 - [HttpServletRequest#login\(java.lang.String, java.lang.String\)](#)
 - [HttpServletRequest#logout\(\)](#)

AbstractSecurityWebApplicationInitializer

The next step is to register the `springSecurityFilterChain` with the war. This can be done in Java Configuration with [Spring's WebApplicationInitializer support](#) in a Servlet 3.0+ environment. Not suprisingly, Spring Security provides a base class `AbstractSecurityWebApplicationInitializer` that will ensure the `springSecurityFilterChain` gets registered for you. The way in which we use `AbstractSecurityWebApplicationInitializer` differs depending on if we are already using Spring or if Spring Security is the only Spring component in our application.

- the section called “AbstractSecurityWebApplicationInitializer without Existing Spring” - Use these instructions if you are not using Spring already
- the section called “AbstractSecurityWebApplicationInitializer with Spring MVC” - Use these instructions if you are already using Spring

AbstractSecurityWebApplicationInitializer without Existing Spring

If you are not using Spring or Spring MVC, you will need to pass in the `WebSecurityConfig` into the superclass to ensure the configuration is picked up. You can find an example below:

```
import org.springframework.security.web.context.*;

public class SecurityWebApplicationInitializer
    extends AbstractSecurityWebApplicationInitializer {

    public SecurityWebApplicationInitializer() {
        super(WebSecurityConfig.class);
    }
}
```

The `SecurityWebApplicationInitializer` will do the following things:

- Automatically register the `springSecurityFilterChain` Filter for every URL in your application

- Add a ContextLoaderListener that loads the [WebSecurityConfig](#).

AbstractSecurityWebApplicationInitializer with Spring MVC

If we were using Spring elsewhere in our application we probably already had a `WebApplicationInitializer` that is loading our Spring Configuration. If we use the previous configuration we would get an error. Instead, we should register Spring Security with the existing `ApplicationContext`. For example, if we were using Spring MVC our `SecurityWebApplicationInitializer` would look something like the following:

```
import org.springframework.security.web.context.*;

public class SecurityWebApplicationInitializer
    extends AbstractSecurityWebApplicationInitializer {
}

```

This would simply only register the `springSecurityFilterChain` Filter for every URL in your application. After that we would ensure that `WebSecurityConfig` was loaded in our existing `ApplicationInitializer`. For example, if we were using Spring MVC it would be added in the `getRootConfigClasses()`

```
public class MvcWebApplicationInitializer extends
    AbstractAnnotationConfigDispatcherServletInitializer {

    @Override
    protected Class<?>[] getRootConfigClasses() {
        return new Class[] { WebSecurityConfig.class };
    }

    // ... other overrides ...
}

```

16.2 HttpSecurity

Thus far our [WebSecurityConfig](#) only contains information about how to authenticate our users. How does Spring Security know that we want to require all users to be authenticated? How does Spring Security know we want to support form based authentication? Actually, there is an configuration class that is being invoked behind the scenes called `WebSecurityConfigurerAdapter`. It has a method called `configure` with the following default implementation:

```
protected void configure(HttpSecurity http) throws Exception {
    http
        .authorizeRequests(authorizeRequests ->
            authorizeRequests
                .anyRequest().authenticated()
        )
        .formLogin(withDefaults())
        .httpBasic(withDefaults());
}

```

The default configuration above:

- Ensures that any request to our application requires the user to be authenticated
- Allows users to authenticate with form based login
- Allows users to authenticate with HTTP Basic authentication

You will notice that this configuration is quite similar the XML Namespace configuration:

```
<http>
  <intercept-url pattern="/**" access="authenticated"/>
  <form-login />
  <http-basic />
</http>
```

16.3 Multiple HttpSecurity

We can configure multiple `HttpSecurity` instances just as we can have multiple `<http>` blocks. The key is to extend the `WebSecurityConfigurerAdapter` multiple times. For example, the following is an example of having a different configuration for URL's that start with `/api/`.

```
@EnableWebSecurity
public class MultiHttpSecurityConfig {
    @Bean
    public UserDetailsService userDetailsService() throws Exception {
        // ensure the passwords are encoded properly
        UserBuilder users = User.withDefaultPasswordEncoder();
        InMemoryUserDetailsManager manager = new InMemoryUserDetailsManager();
        manager.createUser(users.username("user").password("password").roles("USER").build());
        manager.createUser(users.username("admin").password("password").roles("USER", "ADMIN").build());
        return manager;
    }

    @Configuration
    @Order(1)
    public static class ApiWebSecurityConfigurationAdapter extends WebSecurityConfigurerAdapter {
        protected void configure(HttpSecurity http) throws Exception {
            http
                .antMatcher("/api/**")
                .authorizeRequests(authorizeRequests ->
                    authorizeRequests
                        .anyRequest().hasRole("ADMIN")
                )
                .httpBasic(withDefaults());
        }
    }

    @Configuration
    public static class FormLoginWebSecurityConfigurerAdapter extends WebSecurityConfigurerAdapter {
        @Override
        protected void configure(HttpSecurity http) throws Exception {
            http
                .authorizeRequests(authorizeRequests ->
                    authorizeRequests
                        .anyRequest().authenticated()
                )
                .formLogin(withDefaults());
        }
    }
}
```

- ❶ Configure Authentication as normal
- ❷ Create an instance of `WebSecurityConfigurerAdapter` that contains `@Order` to specify which `WebSecurityConfigurerAdapter` should be considered first.
- ❸ The `http.antMatcher` states that this `HttpSecurity` will only be applicable to URLs that start with `/api/`
- ❹ Create another instance of `WebSecurityConfigurerAdapter`. If the URL does not start with `/api/` this configuration will be used. This configuration is considered after `ApiWebSecurityConfigurationAdapter` since it has an `@Order` value after 1 (no `@Order` defaults to last).

16.4 Custom DSLs

You can provide your own custom DSLs in Spring Security. For example, you might have something that looks like this:

```
public class MyCustomDsl extends AbstractHttpConfigurer<MyCustomDsl, HttpSecurity> {
    private boolean flag;

    @Override
    public void init(H http) throws Exception {
        // any method that adds another configurer
        // must be done in the init method
        http.csrf().disable();
    }

    @Override
    public void configure(H http) throws Exception {
        ApplicationContext context = http.getSharedObject(ApplicationContext.class);

        // here we lookup from the ApplicationContext. You can also just create a new instance.
        MyFilter myFilter = context.getBean(MyFilter.class);
        myFilter.setFlag(flag);
        http.addFilterBefore(myFilter, UsernamePasswordAuthenticationFilter.class);
    }

    public MyCustomDsl flag(boolean value) {
        this.flag = value;
        return this;
    }

    public static MyCustomDsl customDsl() {
        return new MyCustomDsl();
    }
}
```

Note

This is actually how methods like `HttpSecurity.authorizeRequests()` are implemented.

The custom DSL can then be used like this:

```
@EnableWebSecurity
public class Config extends WebSecurityConfigurerAdapter {
    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .apply(customDsl())
            .flag(true)
            .and()
            ...;
    }
}
```

The code is invoked in the following order:

- Code in `Config`'s` `configure` method is invoked
- Code in `MyCustomDsl`'s` `init` method is invoked
- Code in `MyCustomDsl`'s` `configure` method is invoked

If you want, you can have `WebSecurityConfigurerAdapter` add `MyCustomDsl` by default by using `SpringFactories`. For example, you would create a resource on the classpath named `META-INF/spring.factories` with the following contents:

META-INF/spring.factories.

```
org.springframework.security.config.annotation.web.configurers.AbstractHttpConfigurer =
sample.MyCustomDsl
```

Users wishing to disable the default can do so explicitly.

```
@EnableWebSecurity
public class Config extends WebSecurityConfigurerAdapter {
    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http
            .apply(customDsl()).disable()
            ...;
    }
}
```

16.5 Post Processing Configured Objects

Spring Security's Java Configuration does not expose every property of every object that it configures. This simplifies the configuration for a majority of users. Afterall, if every property was exposed, users could use standard bean configuration.

While there are good reasons to not directly expose every property, users may still need more advanced configuration options. To address this Spring Security introduces the concept of an `ObjectPostProcessor` which can be used to modify or replace many of the `Object` instances created by the Java Configuration. For example, if you wanted to configure the `filterSecurityPublishAuthorizationSuccess` property on `FilterSecurityInterceptor` you could use the following:

```
@Override
protected void configure(HttpSecurity http) throws Exception {
    http
        .authorizeRequests(authorizeRequests ->
            authorizeRequests
                .anyRequest().authenticated()
                .withObjectPostProcessor(new ObjectPostProcessor<FilterSecurityInterceptor>() {
                    public <O extends FilterSecurityInterceptor> O postProcess(
                        O fsi) {
                        fsi.setPublishAuthorizationSuccess(true);
                        return fsi;
                    }
                })
        );
}
```


17. Security Namespace Configuration

17.1 Introduction

Namespace configuration has been available since version 2.0 of the Spring Framework. It allows you to supplement the traditional Spring beans application context syntax with elements from additional XML schema. You can find more information in the Spring [Reference Documentation](#). A namespace element can be used simply to allow a more concise way of configuring an individual bean or, more powerfully, to define an alternative configuration syntax which more closely matches the problem domain and hides the underlying complexity from the user. A simple element may conceal the fact that multiple beans and processing steps are being added to the application context. For example, adding the following element from the security namespace to an application context will start up an embedded LDAP server for testing use within the application:

```
<security:ldap-server />
```

This is much simpler than wiring up the equivalent Apache Directory Server beans. The most common alternative configuration requirements are supported by attributes on the `ldap-server` element and the user is isolated from worrying about which beans they need to create and what the bean property names are.² Use of a good XML editor while editing the application context file should provide information on the attributes and elements that are available. We would recommend that you try out the [Spring Tool Suite](#) as it has special features for working with standard Spring namespaces.

To start using the security namespace in your application context, you need to have the `spring-security-config` jar on your classpath. Then all you need to do is add the schema declaration to your application context file:

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:security="http://www.springframework.org/schema/security"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    https://www.springframework.org/schema/beans/spring-beans-3.0.xsd
    http://www.springframework.org/schema/security
    https://www.springframework.org/schema/security/spring-security.xsd">
  ...
</beans>
```

In many of the examples you will see (and in the sample applications), we will often use "security" as the default namespace rather than "beans", which means we can omit the prefix on all the security namespace elements, making the content easier to read. You may also want to do this if you have your application context divided up into separate files and have most of your security configuration in one of them. Your security application context file would then start like this

```
<beans:beans xmlns="http://www.springframework.org/schema/security"
  xmlns:beans="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    https://www.springframework.org/schema/beans/spring-beans-3.0.xsd
    http://www.springframework.org/schema/security
    https://www.springframework.org/schema/security/spring-security.xsd">
  ...
</beans:beans>
```

We'll assume this syntax is being used from now on in this chapter.

²You can find out more about the use of the `ldap-server` element in the chapter on Section 10.3, "LDAP Authentication".

Design of the Namespace

The namespace is designed to capture the most common uses of the framework and provide a simplified and concise syntax for enabling them within an application. The design is based around the large-scale dependencies within the framework, and can be divided up into the following areas:

- *Web/HTTP Security* - the most complex part. Sets up the filters and related service beans used to apply the framework authentication mechanisms, to secure URLs, render login and error pages and much more.
- *Business Object (Method) Security* - options for securing the service layer.
- *AuthenticationManager* - handles authentication requests from other parts of the framework.
- *AccessDecisionManager* - provides access decisions for web and method security. A default one will be registered, but you can also choose to use a custom one, declared using normal Spring bean syntax.
- *AuthenticationProviders* - mechanisms against which the authentication manager authenticates users. The namespace provides supports for several standard options and also a means of adding custom beans declared using a traditional syntax.
- *UserDetailsService* - closely related to authentication providers, but often also required by other beans.

We'll see how to configure these in the following sections.

17.2 Getting Started with Security Namespace Configuration

In this section, we'll look at how you can build up a namespace configuration to use some of the main features of the framework. Let's assume you initially want to get up and running as quickly as possible and add authentication support and access control to an existing web application, with a few test logins. Then we'll look at how to change over to authenticating against a database or other security repository. In later sections we'll introduce more advanced namespace configuration options.

web.xml Configuration

The first thing you need to do is add the following filter declaration to your `web.xml` file:

```
<filter>
<filter-name>springSecurityFilterChain</filter-name>
<filter-class>org.springframework.web.filter.DelegatingFilterProxy</filter-class>
</filter>

<filter-mapping>
<filter-name>springSecurityFilterChain</filter-name>
<url-pattern>/*</url-pattern>
</filter-mapping>
```

This provides a hook into the Spring Security web infrastructure. `DelegatingFilterProxy` is a Spring Framework class which delegates to a filter implementation which is defined as a Spring bean in your application context. In this case, the bean is named "springSecurityFilterChain", which is an internal infrastructure bean created by the namespace to handle web security. Note that you should not use this bean name yourself. Once you've added this to your `web.xml`, you're ready to start editing your application context file. Web security services are configured using the `<http>` element.

A Minimal <http> Configuration

All you need to enable web security to begin with is

```
<http>
<intercept-url pattern="/**" access="hasRole('USER')" />
<form-login />
<logout />
</http>
```

Which says that we want all URLs within our application to be secured, requiring the role `ROLE_USER` to access them, we want to log in to the application using a form with username and password, and that we want a logout URL registered which will allow us to log out of the application. `<http>` element is the parent for all web-related namespace functionality. The `<intercept-url>` element defines a `pattern` which is matched against the URLs of incoming requests using an ant path style syntax ⁴. You can also use regular-expression matching as an alternative (see the namespace appendix for more details). The `access` attribute defines the access requirements for requests matching the given pattern. With the default configuration, this is typically a comma-separated list of roles, one of which a user must have to be allowed to make the request. The prefix "ROLE_" is a marker which indicates that a simple comparison with the user's authorities should be made. In other words, a normal role-based check should be used. Access-control in Spring Security is not limited to the use of simple roles (hence the use of the prefix to differentiate between different types of security attributes). We'll see later how the interpretation can vary ⁵. In Spring Security 3.0, the attribute can also be populated with an `#1#`.

Note

===

You can use multiple `<intercept-url>` elements to define different access requirements for different sets of URLs, but they will be evaluated in the order listed and the first match will be used. So you must put the most specific matches at the top. You can also add a `method` attribute to limit the match to a particular HTTP method (GET, POST, PUT etc.).

===

To add some users, you can define a set of test data directly in the namespace:

```
<authentication-manager>
<authentication-provider>
  <user-service>
    <!-- Password is prefixed with {noop} to indicate to DelegatingPasswordEncoder that
    NoOpPasswordEncoder should be used. This is not safe for production, but makes reading
    in samples easier. Normally passwords should be hashed using BCrypt -->
    <user name="jimi" password="{noop}jimispw" authorities="ROLE_USER, ROLE_ADMIN" />
    <user name="bob" password="{noop}bobspw" authorities="ROLE_USER" />
  </user-service>
</authentication-provider>
</authentication-manager>
```

This is an example of a secure way of storing the same passwords. The password is prefixed with `{bcrypt}` to instruct `DelegatingPasswordEncoder`, which supports any configured `PasswordEncoder` for matching, that the passwords are hashed using BCrypt:

⁴See the section on ??? in the Web Application Infrastructure chapter for more details on how matches are actually performed.

⁵The interpretation of the comma-separated values in the `access` attribute depends on the implementation of the [AccessDecisionManager](#) which is used.

```

<authentication-manager>
<authentication-provider>
  <user-service>
    <user name="jimi" password="{bcrypt}$2a$10$ddEWZU18aU0GdZPPpy7wbu82dvEw/pBpbRvDQRqA41y6mK1CoH00m"
      authorities="ROLE_USER, ROLE_ADMIN" />
    <user name="bob" password="{bcrypt}$2a$10$/elFpMBnAYYig6KRR5bvOOYeZr1ie1hSogJryg9qDlhza4oCw1Qka"
      authorities="ROLE_USER" />
    <user name="jimi" password="{noop}jimispassword" authorities="ROLE_USER, ROLE_ADMIN" />
    <user name="bob" password="{noop}bobspassword" authorities="ROLE_USER" />
  </user-service>
</authentication-provider>
</authentication-manager>

```

If you are familiar with pre-namespace versions of the framework, you can probably already guess roughly what's going on here. The `<http>` element is responsible for creating a `FilterChainProxy` and the filter beans which it uses. Common problems like incorrect filter ordering are no longer an issue as the filter positions are predefined.

The `<authentication-provider>` element creates a `DaoAuthenticationProvider` bean and the `<user-service>` element creates an `InMemoryDaoImpl`. All authentication-provider elements must be children of the `<authentication-manager>` element, which creates a `ProviderManager` and registers the authentication providers with it. You can find more detailed information on the beans that are created in the [namespace appendix](#). It's worth cross-checking this if you want to start understanding what the important classes in the framework are and how they are used, particularly if you want to customise things later.

The configuration above defines two users, their passwords and their roles within the application (which will be used for access control). It is also possible to load user information from a standard properties file using the `properties` attribute on `user-service`. See the section on [in-memory authentication](#) for more details on the file format. Using the `<authentication-provider>` element means that the user information will be used by the authentication manager to process authentication requests. You can have multiple `<authentication-provider>` elements to define different authentication sources and each will be consulted in turn.

At this point you should be able to start up your application and you will be required to log in to proceed. Try it out, or try experimenting with the "tutorial" sample application that comes with the project.

Setting a Default Post-Login Destination

If a form login isn't prompted by an attempt to access a protected resource, the `default-target-url` option comes into play. This is the URL the user will be taken to after successfully logging in, and defaults to `"/`. You can also configure things so that the user *always* ends up at this page (regardless of whether the login was "on-demand" or they explicitly chose to log in) by setting the `always-use-default-target` attribute to `"true"`. This is useful if your application always requires that the user starts at a "home" page, for example:

```

<http pattern="/login.htm*" security="none"/>
<http use-expressions="false">
<intercept-url pattern='/**' access='ROLE_USER' />
<form-login login-page='/login.htm' default-target-url='/home.htm'
  always-use-default-target='true' />
</http>

```

For even more control over the destination, you can use the `authentication-success-handler-ref` attribute as an alternative to `default-target-url`. The referenced bean should be an instance

of `AuthenticationSuccessHandler`. You'll find more on this in the [Core Filters](#) chapter and also in the namespace appendix, as well as information on how to customize the flow when authentication fails.

17.3 Advanced Web Features

Adding in Your Own Filters

If you've used Spring Security before, you'll know that the framework maintains a chain of filters in order to apply its services. You may want to add your own filters to the stack at particular locations or use a Spring Security filter for which there isn't currently a namespace configuration option (CAS, for example). Or you might want to use a customized version of a standard namespace filter, such as the `UsernamePasswordAuthenticationFilter` which is created by the `<form-login>` element, taking advantage of some of the extra configuration options which are available by using the bean explicitly. How can you do this with namespace configuration, since the filter chain is not directly exposed?

The order of the filters is always strictly enforced when using the namespace. When the application context is being created, the filter beans are sorted by the namespace handling code and the standard Spring Security filters each have an alias in the namespace and a well-known position.

Note

=== In previous versions, the sorting took place after the filter instances had been created, during post-processing of the application context. In version 3.0+ the sorting is now done at the bean metadata level, before the classes have been instantiated. This has implications for how you add your own filters to the stack as the entire filter list must be known during the parsing of the `<http>` element, so the syntax has changed slightly in 3.0. ===

The filters, aliases and namespace elements/attributes which create the filters are shown in Table 17.1, "Standard Filter Aliases and Ordering". The filters are listed in the order in which they occur in the filter chain.

Table 17.1. Standard Filter Aliases and Ordering

Alias	Filter Class	Namespace Element or Attribute
CHANNEL_FILTER	<code>ChannelProcessingFilter</code>	<code>http/intercept-url@requires-channel</code>
SECURITY_CONTEXT_FILTER	<code>SecurityContextPersistenceFilter</code>	<code>http</code>
CONCURRENT_SESSION_FILTER	<code>ConcurrentSessionFilter</code>	<code>session-management/concurrency-control</code>
HEADERS_FILTER	<code>HeaderWriterFilter</code>	<code>http/headers</code>
CSRF_FILTER	<code>CsrfFilter</code>	<code>http/csrf</code>
LOGOUT_FILTER	<code>LogoutFilter</code>	<code>http/logout</code>
X509_FILTER	<code>X509AuthenticationFilter</code>	<code>http/x509</code>
PRE_AUTH_FILTER	<code>AbstractPreAuthenticatedProcessingFilter</code> Subclasses	N/A

Alias	Filter Class	Namespace Element or Attribute
CAS_FILTER	CasAuthenticationFilter	N/A
FORM_LOGIN_FILTER	UsernamePasswordAuthenticationFilter	http/form-login
BASIC_AUTH_FILTER	BasicAuthenticationFilter	http/http-basic
SERVLET_API_SUPPORT_FILTER	SecurityContextHolderAwareRequestFilter	http/api-provision
JAAS_API_SUPPORT_FILTER	JaasApiIntegrationFilter	http/@jaas-api-provision
REMEMBER_ME_FILTER	RememberMeAuthenticationFilter	http/remember-me
ANONYMOUS_FILTER	AnonymousAuthenticationFilter	http/anonymous
SESSION_MANAGEMENT_FILTER	SessionManagementFilter	session-management
EXCEPTION_TRANSLATION_FILTER	ExceptionHandler	http
FILTER_SECURITY_INTERCEPTOR	FilterSecurityInterceptor	http
SWITCH_USER_FILTER	SwitchUserFilter	N/A

You can add your own filter to the stack, using the `custom-filter` element and one of these names to specify the position your filter should appear at:

```
<http>
<custom-filter position="FORM_LOGIN_FILTER" ref="myFilter" />
</http>

<beans:bean id="myFilter" class="com.mycompany.MySpecialAuthenticationFilter"/>
```

You can also use the `after` or `before` attributes if you want your filter to be inserted before or after another filter in the stack. The names "FIRST" and "LAST" can be used with the `position` attribute to indicate that you want your filter to appear before or after the entire stack, respectively.

Avoiding filter position conflicts

===

If you are inserting a custom filter which may occupy the same position as one of the standard filters created by the namespace then it's important that you don't include the namespace versions by mistake. Remove any elements which create filters whose functionality you want to replace.

Note that you can't replace filters which are created by the use of the `<http>` element itself - `SecurityContextPersistenceFilter`, `ExceptionHandler` or `FilterSecurityInterceptor`. Some other filters are added by default, but you can disable them. An `AnonymousAuthenticationFilter` is added by default and unless you have [session-fixation protection](#) disabled, a `SessionManagementFilter` will also be added to the filter chain.

===

If you're replacing a namespace filter which requires an authentication entry point (i.e. where the authentication process is triggered by an attempt by an unauthenticated user to access to a secured resource), you will need to add a custom entry point bean too.

17.4 Method Security

From version 2.0 onwards Spring Security has improved support substantially for adding security to your service layer methods. It provides support for JSR-250 annotation security as well as the framework's original `@Secured` annotation. From 3.0 you can also make use of new [expression-based annotations](#). You can apply security to a single bean, using the `intercept-methods` element to decorate the bean declaration, or you can secure multiple beans across the entire service layer using the AspectJ style pointcuts.

17.5 The Default AccessDecisionManager

This section assumes you have some knowledge of the underlying architecture for access-control within Spring Security. If you don't you can skip it and come back to it later, as this section is only really relevant for people who need to do some customization in order to use more than simple role-based security.

When you use a namespace configuration, a default instance of `AccessDecisionManager` is automatically registered for you and will be used for making access decisions for method invocations and web URL access, based on the access attributes you specify in your `intercept-url` and `protect-pointcut` declarations (and in annotations if you are using annotation secured methods).

The default strategy is to use an `AffirmativeBased AccessDecisionManager` with a `RoleVoter` and an `AuthenticatedVoter`. You can find out more about these in the chapter on [authorization](#).

Customizing the AccessDecisionManager

If you need to use a more complicated access control strategy then it is easy to set an alternative for both method and web security.

For method security, you do this by setting the `access-decision-manager-ref` attribute on `global-method-security` to the id of the appropriate `AccessDecisionManager` bean in the application context:

```
<global-method-security access-decision-manager-ref="myAccessDecisionManagerBean">
...
</global-method-security>
```

The syntax for web security is the same, but on the `http` element:

```
<http access-decision-manager-ref="myAccessDecisionManagerBean">
...
</http>
```

18. Testing

This section describes the testing support provided by Spring Security.

Tip

To use the Spring Security test support, you must include `spring-security-test-5.2.1.RELEASE.jar` as a dependency of your project.

18.1 Testing Method Security

This section demonstrates how to use Spring Security's Test support to test method based security. We first introduce a `MessageService` that requires the user to be authenticated in order to access it.

```
public class HelloMessageService implements MessageService {
    @PreAuthorize("authenticated")
    public String getMessage() {
        Authentication authentication = SecurityContextHolder.getContext()
            .getAuthentication();
        return "Hello " + authentication;
    }
}
```

The result of `getMessage` is a String saying "Hello" to the current Spring Security Authentication. An example of the output is displayed below.

```
Hello org.springframework.security.authentication.UsernamePasswordAuthenticationToken@ca25360:
Principal: org.springframework.security.core.userdetails.User@36ebcb: Username: user; Password:
[PROTECTED]; Enabled: true; AccountNonExpired: true; credentialsNonExpired: true; AccountNonLocked:
true; Granted Authorities: ROLE_USER; Credentials: [PROTECTED]; Authenticated: true; Details: null;
Granted Authorities: ROLE_USER
```

Security Test Setup

Before we can use Spring Security Test support, we must perform some setup. An example can be seen below:

```
@RunWith(SpringJUnit4ClassRunner.class) ❶
@ContextConfiguration ❷
public class WithMockUserTests {
```

This is a basic example of how to setup Spring Security Test. The highlights are:

- ❶ `@RunWith` instructs the spring-test module that it should create an `ApplicationContext`. This is no different than using the existing Spring Test support. For additional information, refer to the [Spring Reference](#)
- ❷ `@ContextConfiguration` instructs the spring-test the configuration to use to create the `ApplicationContext`. Since no configuration is specified, the default configuration locations will be tried. This is no different than using the existing Spring Test support. For additional information, refer to the [Spring Reference](#)

Note

Spring Security hooks into Spring Test support using the `WithSecurityContextTestExecutionListener` which will ensure our tests are ran

with the correct user. It does this by populating the `SecurityContextHolder` prior to running our tests. If you are using reactive method security, you will also need `ReactorContextTestExecutionListener` which populates `ReactiveSecurityContextHolder`. After the test is done, it will clear out the `SecurityContextHolder`. If you only need Spring Security related support, you can replace `@ContextConfiguration` with `@SecurityTestExecutionListeners`.

Remember we added the `@PreAuthorize` annotation to our `HelloMessageService` and so it requires an authenticated user to invoke it. If we ran the following test, we would expect the following test will pass:

```
@Test(expected = AuthenticationCredentialsNotFoundException.class)
public void getMessageUnauthenticated() {
    messageService.getMessage();
}
```

@WithMockUser

The question is "How could we most easily run the test as a specific user?" The answer is to use `@WithMockUser`. The following test will be run as a user with the username "user", the password "password", and the roles "ROLE_USER".

```
@Test
@WithMockUser
public void getMessageWithMockUser() {
    String message = messageService.getMessage();
    ...
}
```

Specifically the following is true:

- The user with the username "user" does not have to exist since we are mocking the user
- The Authentication that is populated in the SecurityContext is of type `UsernamePasswordAuthenticationToken`
- The principal on the Authentication is Spring Security's `User` object
- The `User` will have the username of "user", the password "password", and a single `GrantedAuthority` named "ROLE_USER" is used.

Our example is nice because we are able to leverage a lot of defaults. What if we wanted to run the test with a different username? The following test would run with the username "customUser". Again, the user does not need to actually exist.

```
@Test
@WithMockUser("customUsername")
public void getMessageWithMockUserCustomUsername() {
    String message = messageService.getMessage();
    ...
}
```

We can also easily customize the roles. For example, this test will be invoked with the username "admin" and the roles "ROLE_USER" and "ROLE_ADMIN".

```

@Test
@WithMockUser(username="admin",roles={"USER","ADMIN"})
public void getMessageWithMockUserCustomUser() {
    String message = messageService.getMessage();
    ...
}

```

If we do not want the value to automatically be prefixed with `ROLE_` we can leverage the `authorities` attribute. For example, this test will be invoked with the username "admin" and the authorities "USER" and "ADMIN".

```

@Test
@WithMockUser(username = "admin", authorities = { "ADMIN", "USER" })
public void getMessageWithMockUserCustomAuthorities() {
    String message = messageService.getMessage();
    ...
}

```

Of course it can be a bit tedious placing the annotation on every test method. Instead, we can place the annotation at the class level and every test will use the specified user. For example, the following would run every test with a user with the username "admin", the password "password", and the roles "ROLE_USER" and "ROLE_ADMIN".

```

@RunWith(SpringJUnit4ClassRunner.class)
@ContextConfiguration
@WithMockUser(username="admin",roles={"USER","ADMIN"})
public class WithMockUserTests {

```

By default the `SecurityContext` is set during the `TestExecutionListener.beforeTestMethod` event. This is the equivalent of happening before JUnit's `@Before`. You can change this to happen during the `TestExecutionListener.beforeTestExecution` event which is after JUnit's `@Before` but before the test method is invoked.

```

@WithMockUser(setupBefore = TestExecutionEvent.TEST_EXECUTION)

```

@WithAnonymousUser

Using `@WithAnonymousUser` allows running as an anonymous user. This is especially convenient when you wish to run most of your tests with a specific user, but want to run a few tests as an anonymous user. For example, the following will run `withMockUser1` and `withMockUser2` using [@WithMockUser](#) and `anonymous` as an anonymous user.

```

@RunWith(SpringJUnit4ClassRunner.class)
@WithMockUser
public class WithUserClassLevelAuthenticationTests {

    @Test
    public void withMockUser1() {
    }

    @Test
    public void withMockUser2() {
    }

    @Test
    @WithAnonymousUser
    public void anonymous() throws Exception {
        // override default to run as anonymous user
    }
}

```

By default the `SecurityContext` is set during the `TestExecutionListener.beforeTestMethod` event. This is the equivalent of happening before JUnit's `@Before`. You can change this to happen during the `TestExecutionListener.beforeTestExecution` event which is after JUnit's `@Before` but before the test method is invoked.

```
@WithAnonymousUser(setupBefore = TestExecutionEvent.TEST_EXECUTION)
```

@WithUserDetails

While `@WithMockUser` is a very convenient way to get started, it may not work in all instances. For example, it is common for applications to expect that the `Authentication` principal be of a specific type. This is done so that the application can refer to the principal as the custom type and reduce coupling on Spring Security.

The custom principal is often times returned by a custom `UserDetailsService` that returns an object that implements both `UserDetails` and the custom type. For situations like this, it is useful to create the test user using the custom `UserDetailsService`. That is exactly what `@WithUserDetails` does.

Assuming we have a `UserDetailsService` exposed as a bean, the following test will be invoked with an `Authentication` of type `UsernamePasswordAuthenticationToken` and a principal that is returned from the `UserDetailsService` with the username of "user".

```
@Test
@WithUserDetails
public void getMessageWithUserDetails() {
    String message = messageService.getMessage();
    ...
}
```

We can also customize the username used to lookup the user from our `UserDetailsService`. For example, this test would be executed with a principal that is returned from the `UserDetailsService` with the username of "customUsername".

```
@Test
@WithUserDetails("customUsername")
public void getMessageWithUserDetailsCustomUsername() {
    String message = messageService.getMessage();
    ...
}
```

We can also provide an explicit bean name to look up the `UserDetailsService`. For example, this test would look up the username of "customUsername" using the `UserDetailsService` with the bean name "myUserDetailsService".

```
@Test
@WithUserDetails(value="customUsername", userDetailsServiceBeanName="myUserDetailsService")
public void getMessageWithUserDetailsServiceBeanName() {
    String message = messageService.getMessage();
    ...
}
```

Like `@WithMockUser` we can also place our annotation at the class level so that every test uses the same user. However unlike `@WithMockUser`, `@WithUserDetails` requires the user to exist.

By default the `SecurityContext` is set during the `TestExecutionListener.beforeTestMethod` event. This is the equivalent of happening before JUnit's `@Before`. You can change this to happen during the `TestExecutionListener.beforeTestExecution` event which is after JUnit's `@Before` but before the test method is invoked.

```
@WithUserDetails(setupBefore = TestExecutionEvent.TEST_EXECUTION)
```

@WithSecurityContext

We have seen that `@WithMockUser` is an excellent choice if we are not using a custom Authentication principal. Next we discovered that `@WithUserDetails` would allow us to use a custom `UserDetailsService` to create our Authentication principal but required the user to exist. We will now see an option that allows the most flexibility.

We can create our own annotation that uses the `@WithSecurityContext` to create any `SecurityContext` we want. For example, we might create an annotation named `@WithMockCustomUser` as shown below:

```
@Retention(RetentionPolicy.RUNTIME)
@WithSecurityContext(factory = WithMockCustomUserSecurityContextFactory.class)
public @interface WithMockCustomUser {

    String username() default "rob";

    String name() default "Rob Winch";
}
```

You can see that `@WithMockCustomUser` is annotated with the `@WithSecurityContext` annotation. This is what signals to Spring Security Test support that we intend to create a `SecurityContext` for the test. The `@WithSecurityContext` annotation requires we specify a `SecurityContextFactory` that will create a new `SecurityContext` given our `@WithMockCustomUser` annotation. You can find our `WithMockCustomUserSecurityContextFactory` implementation below:

```
public class WithMockCustomUserSecurityContextFactory
    implements WithSecurityContextFactory<WithMockCustomUser> {
    @Override
    public SecurityContext createSecurityContext(WithMockCustomUser customUser) {
        SecurityContext context = SecurityContextHolder.createEmptyContext();

        CustomUserDetails principal =
            new CustomUserDetails(customUser.name(), customUser.username());
        Authentication auth =
            new UsernamePasswordAuthenticationToken(principal, "password", principal.getAuthorities());
        context.setAuthentication(auth);
        return context;
    }
}
```

We can now annotate a test class or a test method with our new annotation and Spring Security's `WithSecurityContextTestExecutionListener` will ensure that our `SecurityContext` is populated appropriately.

When creating your own `WithSecurityContextFactory` implementations, it is nice to know that they can be annotated with standard Spring annotations. For example, the `WithUserDetailsSecurityContextFactory` uses the `@Autowired` annotation to acquire the `UserDetailsService`:

```

final class WithUserDetailsSecurityContextFactory
    implements WithSecurityContextFactory<WithUserDetails> {

    private UserDetailsService userDetailsService;

    @Autowired
    public WithUserDetailsSecurityContextFactory(UserDetailsService userDetailsService) {
        this.userDetailsService = userDetailsService;
    }

    public SecurityContext createSecurityContext(WithUserDetails withUser) {
        String username = withUser.value();
        Assert.hasLength(username, "value() must be non-empty String");
        UserDetails principal = userDetailsService.loadUserByUsername(username);
        Authentication authentication = new UsernamePasswordAuthenticationToken(principal,
        principal.getPassword(), principal.getAuthorities());
        SecurityContext context = SecurityContextHolder.createEmptyContext();
        context.setAuthentication(authentication);
        return context;
    }
}

```

By default the `SecurityContext` is set during the `TestExecutionListener.beforeTestMethod` event. This is the equivalent of happening before JUnit's `@Before`. You can change this to happen during the `TestExecutionListener.beforeTestExecution` event which is after JUnit's `@Before` but before the test method is invoked.

```
@WithSecurityContext(setupBefore = TestExecutionEvent.TEST_EXECUTION)
```

Test Meta Annotations

If you reuse the same user within your tests often, it is not ideal to have to repeatedly specify the attributes. For example, if there are many tests related to an administrative user with the username "admin" and the roles `ROLE_USER` and `ROLE_ADMIN` you would have to write:

```
@WithMockUser(username="admin",roles={"USER","ADMIN"})
```

Rather than repeating this everywhere, we can use a meta annotation. For example, we could create a meta annotation named `WithMockAdmin`:

```
@Retention(RetentionPolicy.RUNTIME)
@WithMockUser(value="rob",roles="ADMIN")
public @interface WithMockAdmin { }
```

Now we can use `@WithMockAdmin` in the same way as the more verbose `@WithMockUser`.

Meta annotations work with any of the testing annotations described above. For example, this means we could create a meta annotation for `@WithUserDetails("admin")` as well.

18.2 Spring MVC Test Integration

Spring Security provides comprehensive integration with [Spring MVC Test](#)

Setting Up MockMvc and Spring Security

In order to use Spring Security with Spring MVC Test it is necessary to add the Spring Security `FilterChainProxy` as a `Filter`. It is also necessary to add Spring Security's `TestSecurityContextHolderPostProcessor` to support [Running as a User in Spring MVC Test with Annotations](#). This can be done using Spring Security's `SecurityMockMvcConfigurers.springSecurity()`. For example:

Note

Spring Security's testing support requires `spring-test-4.1.3.RELEASE` or greater.

```
import static org.springframework.security.test.web.servlet.setup.SecurityMockMvcConfigurers.*;

@RunWith(SpringJUnit4ClassRunner.class)
@ContextConfiguration
@WebAppConfiguration
public class CsrfShowcaseTests {

    @Autowired
    private WebApplicationContext context;

    private MockMvc mvc;

    @Before
    public void setup() {
        mvc = MockMvcBuilders
            .webAppContextSetup(context)
            .apply(springSecurity()) ❶
            .build();
    }

    ...
}
```

- ❶ `SecurityMockMvcConfigurers.springSecurity()` will perform all of the initial setup we need to integrate Spring Security with Spring MVC Test

SecurityMockMvcRequestPostProcessors

Spring MVC Test provides a convenient interface called a `RequestPostProcessor` that can be used to modify a request. Spring Security provides a number of `RequestPostProcessor` implementations that make testing easier. In order to use Spring Security's `RequestPostProcessor` implementations ensure the following static import is used:

```
import static
    org.springframework.security.test.web.servlet.request.SecurityMockMvcRequestPostProcessors.*;
```

Testing with CSRF Protection

When testing any non-safe HTTP methods and using Spring Security's CSRF protection, you must be sure to include a valid CSRF Token in the request. To specify a valid CSRF token as a request parameter using the following:

```
mvc
    .perform(post("/").with(csrf()))
```

If you like you can include CSRF token in the header instead:

```
mvc
    .perform(post("/").with(csrf().asHeader()))
```

You can also test providing an invalid CSRF token using the following:

```
mvc
    .perform(post("/").with(csrf().useInvalidToken()))
```

Running a Test as a User in Spring MVC Test

It is often desirable to run tests as a specific user. There are two simple ways of populating the user:

- [Running as a User in Spring MVC Test with RequestPostProcessor](#)
- [Running as a User in Spring MVC Test with Annotations](#)

Running as a User in Spring MVC Test with RequestPostProcessor

There are a number of options available to associate a user to the current `HttpServletRequest`. For example, the following will run as a user (which does not need to exist) with the username "user", the password "password", and the role "ROLE_USER":

Note

The support works by associating the user to the `HttpServletRequest`. To associate the request to the `SecurityContextHolder` you need to ensure that the `SecurityContextPersistenceFilter` is associated with the `MockMvc` instance. A few ways to do this are:

- Invoking [apply\(springSecurity\(\)\)](#)
- Adding Spring Security's `FilterChainProxy` to `MockMvc`
- Manually adding `SecurityContextPersistenceFilter` to the `MockMvc` instance may make sense when using `MockMvcBuilders.standaloneSetup`

```
mvc
    .perform(get("/").with(user("user")))
```

You can easily make customizations. For example, the following will run as a user (which does not need to exist) with the username "admin", the password "pass", and the roles "ROLE_USER" and "ROLE_ADMIN".

```
mvc
    .perform(get("/admin").with(user("admin").password("pass").roles("USER", "ADMIN")))
```

If you have a custom `UserDetails` that you would like to use, you can easily specify that as well. For example, the following will use the specified `UserDetails` (which does not need to exist) to run with a `UsernamePasswordAuthenticationToken` that has a principal of the specified `UserDetails`:

```
mvc
    .perform(get("/").with(user(userDetails)))
```

You can run as an anonymous user using the following:

```
mvc
    .perform(get("/").with(anonymous()))
```

This is especially useful if you are running with a default user and wish to execute a few requests as an anonymous user.

If you want a custom `Authentication` (which does not need to exist) you can do so using the following:

```
mvc
    .perform(get("/").with(authentication(authentication)))
```

You can even customize the `SecurityContext` using the following:

```
mvc
    .perform(get("/").with(securityContext(securityContext)))
```

We can also ensure to run as a specific user for every request by using `MockMvcBuilders`'s default request. For example, the following will run as a user (which does not need to exist) with the username "admin", the password "password", and the role "ROLE_ADMIN":

```
mvc = MockMvcBuilders
    .webApplicationContextSetup(context)
    .defaultRequest(get("/").with(user("user").roles("ADMIN")))
    .apply(springSecurity())
    .build();
```

If you find you are using the same user in many of your tests, it is recommended to move the user to a method. For example, you can specify the following in your own class named `CustomSecurityMockMvcRequestPostProcessors`:

```
public static RequestPostProcessor rob() {
    return user("rob").roles("ADMIN");
}
```

Now you can perform a static import on `SecurityMockMvcRequestPostProcessors` and use that within your tests:

```
import static sample.CustomSecurityMockMvcRequestPostProcessors.*;

...

mvc
    .perform(get("/").with(rob()))
```

Running as a User in Spring MVC Test with Annotations

As an alternative to using a `RequestPostProcessor` to create your user, you can use annotations described in Section 18.1, "Testing Method Security". For example, the following will run the test with the user with username "user", password "password", and role "ROLE_USER":

```
@Test
@WithMockUser
public void requestProtectedUrlWithUser() throws Exception {
    mvc
        .perform(get("/"))
        ...
}
```

Alternatively, the following will run the test with the user with username "user", password "password", and role "ROLE_ADMIN":

```
@Test
@WithMockUser(roles="ADMIN")
public void requestProtectedUrlWithUser() throws Exception {
    mvc
        .perform(get("/"))
        ...
}
```

Testing HTTP Basic Authentication

While it has always been possible to authenticate with HTTP Basic, it was a bit tedious to remember the header name, format, and encode the values. Now this can be done using Spring Security's `HttpBasicRequestPostProcessor`. For example, the snippet below:


```
mvc
    .perform(get("/").with(httpBasic("user", "password")))
```

will attempt to use HTTP Basic to authenticate a user with the username "user" and the password "password" by ensuring the following header is populated on the HTTP Request:

```
Authorization: Basic dXNlcjpwYXNzd29yZA==
```

SecurityMockMvcRequestBuilders

Spring MVC Test also provides a `RequestBuilder` interface that can be used to create the `MockHttpServletRequest` used in your test. Spring Security provides a few `RequestBuilder` implementations that can be used to make testing easier. In order to use Spring Security's `RequestBuilder` implementations ensure the following static import is used:

```
import static org.springframework.security.test.web.servlet.request.SecurityMockMvcRequestBuilders.*;
```

Testing Form Based Authentication

You can easily create a request to test a form based authentication using Spring Security's testing support. For example, the following will submit a POST to `/login` with the username "user", the password "password", and a valid CSRF token:

```
mvc
    .perform(formLogin())
```

It is easy to customize the request. For example, the following will submit a POST to `/auth` with the username "admin", the password "pass", and a valid CSRF token:

```
mvc
    .perform(formLogin("/auth").user("admin").password("pass"))
```

We can also customize the parameters names that the username and password are included on. For example, this is the above request modified to include the username on the HTTP parameter "u" and the password on the HTTP parameter "p".

```
mvc
    .perform(formLogin("/auth").user("u", "admin").password("p", "pass"))
```

Testing Bearer Authentication

In order to make an authorized request on a resource server, you need a bearer token. If your resource server is configured for JWTs, then this would mean that the bearer token needs to be signed and then encoded according to the JWT specification. All of this can be quite daunting, especially when this isn't the focus of your test.

Fortunately, there are a number of simple ways that you can overcome this difficulty and allow your tests to focus on authorization and not on representing bearer tokens. We'll look at two of them now:

`jwt()` `RequestPostProcessor`

The first way is via a `RequestPostProcessor`. The simplest of these would look something like this:

```
mvc
    .perform(get("/endpoint").with(jwt()));
```

What this will do is create a mock `JWT`, passing it correctly through any authentication APIs so that it's available for your authorization mechanisms to verify.

By default, the `JWT` that it creates has the following characteristics:

```
{
  "headers" : { "alg" : "none" },
  "claims" : {
    "sub" : "user",
    "scope" : "read"
  }
}
```

And the resulting `JWT`, were it tested, would pass in the following way:

```
assertThat(jwt.getTokenValue()).isEqualTo("token");
assertThat(jwt.getHeaders().get("alg")).isEqualTo("none");
assertThat(jwt.getSubject()).isEqualTo("sub");
GrantedAuthority authority = jwt.getAuthorities().iterator().next();
assertThat(authority.getAuthority()).isEqualTo("read");
```

These values can, of course be configured.

Any headers or claims can be configured with their corresponding methods:

```
mvc
    .perform(get("/endpoint")
        .with(jwt(jwt -> jwt.header("kid", "one").claim("iss", "https://idp.example.org"))));
```

```
mvc
    .perform(get("/endpoint")
        .with(jwt(jwt -> jwt.claims(claims -> claims.remove("scope"))));
```

The `scope` and `scp` claims are processed the same way here as they are in a normal bearer token request. However, this can be overridden simply by providing the list of `GrantedAuthority` instances that you need for your test:

```
mvc
    .perform(get("/endpoint")
        .with(jwt().authorities(new SimpleGrantedAuthority("SCOPE_messages"))));
```

Or, if you have a custom `JWT` to `Collection<GrantedAuthority>` converter, you can also use that to derive the authorities:

```
mvc
    .perform(get("/endpoint")
        .with(jwt().authorities(new MyConverter())));
```

You can also specify a complete `JWT`, for which [Jwt.Builder](#) comes quite handy:

```
Jwt jwt = Jwt.withTokenValue("token")
    .header("alg", "none")
    .claim("sub", "user")
    .claim("scope", "read");

mvc
    .perform(get("/endpoint")
        .with(jwt(jwt)));
```

`authentication()` `RequestPostProcessor`

The second way is by using the `authentication()` `RequestPostProcessor`. Essentially, you can instantiate your own `JwtAuthenticationToken` and provide it in your test, like so:

```

Jwt jwt = Jwt.withTokenValue("token")
    .header("alg", "none")
    .claim("sub", "user")
    .build();
Collection<GrantedAuthority> authorities = AuthorityUtils.createAuthorityList("SCOPE_read");
JwtAuthenticationToken token = new JwtAuthenticationToken(jwt, authorities);

mvc
    .perform(get("/endpoint")
        .with(authentication(token)));

```

Note that as an alternative to these, you can also mock the `JwtDecoder` bean itself with a `@MockBean` annotation.

Testing Logout

While fairly trivial using standard Spring MVC Test, you can use Spring Security's testing support to make testing log out easier. For example, the following will submit a POST to `/logout` with a valid CSRF token:

```

mvc
    .perform(logout());

```

You can also customize the URL to post to. For example, the snippet below will submit a POST to `/signout` with a valid CSRF token:

```

mvc
    .perform(logout("/signout"));

```

SecurityMockMvcResultMatchers

At times it is desirable to make various security related assertions about a request. To accommodate this need, Spring Security Test support implements Spring MVC Test's `ResultMatcher` interface. In order to use Spring Security's `ResultMatcher` implementations ensure the following static import is used:

```

import static org.springframework.security.test.web.servlet.response.SecurityMockMvcResultMatchers.*;

```

Unauthenticated Assertion

At times it may be valuable to assert that there is no authenticated user associated with the result of a `MockMvc` invocation. For example, you might want to test submitting an invalid username and password and verify that no user is authenticated. You can easily do this with Spring Security's testing support using something like the following:

```

mvc
    .perform(formLogin().password("invalid"))
    .andExpect(unauthenticated());

```

Authenticated Assertion

It is often times that we must assert that an authenticated user exists. For example, we may want to verify that we authenticated successfully. We could verify that a form based login was successful with the following snippet of code:

```

mvc
    .perform(formLogin())
    .andExpect(authenticated());

```

If we wanted to assert the roles of the user, we could refine our previous code as shown below:

```
mvc
    .perform(formLogin().user("admin"))
    .andExpect(authenticated().withRoles("USER", "ADMIN"));
```

Alternatively, we could verify the username:

```
mvc
    .perform(formLogin().user("admin"))
    .andExpect(authenticated().withUsername("admin"));
```

We can also combine the assertions:

```
mvc
    .perform(formLogin().user("admin").roles("USER", "ADMIN"))
    .andExpect(authenticated().withUsername("admin"));
```

We can also make arbitrary assertions on the authentication

```
mvc
    .perform(formLogin())
    .andExpect(authenticated().withAuthentication(auth ->
        assertThat(auth).assertInstanceOf(UsernamePasswordAuthenticationToken.class)));
```

19. Spring Security Crypto Module

19.1 Introduction

The Spring Security Crypto module provides support for symmetric encryption, key generation, and password encoding. The code is distributed as part of the core module but has no dependencies on any other Spring Security (or Spring) code.

19.2 Encryptors

The Encryptors class provides factory methods for constructing symmetric encryptors. Using this class, you can create ByteEncryptors to encrypt data in raw byte[] form. You can also construct TextEncryptors to encrypt text strings. Encryptors are thread-safe.

BytesEncryptor

Use the Encryptors.standard factory method to construct a "standard" BytesEncryptor:

```
Encryptors.standard("password", "salt");
```

The "standard" encryption method is 256-bit AES using PKCS #5's PBKDF2 (Password-Based Key Derivation Function #2). This method requires Java 6. The password used to generate the SecretKey should be kept in a secure place and not be shared. The salt is used to prevent dictionary attacks against the key in the event your encrypted data is compromised. A 16-byte random initialization vector is also applied so each encrypted message is unique.

The provided salt should be in hex-encoded String form, be random, and be at least 8 bytes in length. Such a salt may be generated using a KeyGenerator:

```
String salt = KeyGenerators.string().generateKey(); // generates a random 8-byte salt that is then hex-encoded
```

TextEncryptor

Use the Encryptors.text factory method to construct a standard TextEncryptor:

```
Encryptors.text("password", "salt");
```

A TextEncryptor uses a standard BytesEncryptor to encrypt text data. Encrypted results are returned as hex-encoded strings for easy storage on the filesystem or in the database.

Use the Encryptors.queryableText factory method to construct a "queryable" TextEncryptor:

```
Encryptors.queryableText("password", "salt");
```

The difference between a queryable TextEncryptor and a standard TextEncryptor has to do with initialization vector (iv) handling. The iv used in a queryable TextEncryptor#encrypt operation is shared, or constant, and is not randomly generated. This means the same text encrypted multiple times will always produce the same encryption result. This is less secure, but necessary for encrypted data that needs to be queried against. An example of queryable encrypted text would be an OAuth apiKey.

19.3 Key Generators

The `KeyGenerators` class provides a number of convenience factory methods for constructing different types of key generators. Using this class, you can create a `BytesKeyGenerator` to generate `byte[]` keys. You can also construct a `StringKeyGenerator` to generate string keys. `KeyGenerators` are thread-safe.

BytesKeyGenerator

Use the `KeyGenerators.secureRandom` factory methods to generate a `BytesKeyGenerator` backed by a `SecureRandom` instance:

```
BytesKeyGenerator generator = KeyGenerators.secureRandom();
byte[] key = generator.generateKey();
```

The default key length is 8 bytes. There is also a `KeyGenerators.secureRandom` variant that provides control over the key length:

```
KeyGenerators.secureRandom(16);
```

Use the `KeyGenerators.shared` factory method to construct a `BytesKeyGenerator` that always returns the same key on every invocation:

```
KeyGenerators.shared(16);
```

StringKeyGenerator

Use the `KeyGenerators.string` factory method to construct a 8-byte, `SecureRandom` `KeyGenerator` that hex-encodes each key as a `String`:

```
KeyGenerators.string();
```

19.4 Password Encoding

The password package of the `spring-security-crypto` module provides support for encoding passwords. `PasswordEncoder` is the central service interface and has the following signature:

```
public interface PasswordEncoder {

    String encode(String rawPassword);

    boolean matches(String rawPassword, String encodedPassword);

}
```

The `matches` method returns `true` if the `rawPassword`, once encoded, equals the `encodedPassword`. This method is designed to support password-based authentication schemes.

The `BCryptPasswordEncoder` implementation uses the widely supported "bcrypt" algorithm to hash the passwords. Bcrypt uses a random 16 byte salt value and is a deliberately slow algorithm, in order to hinder password crackers. The amount of work it does can be tuned using the "strength" parameter which takes values from 4 to 31. The higher the value, the more work has to be done to calculate the hash. The default value is 10. You can change this value in your deployed system without affecting existing passwords, as the value is also stored in the encoded hash.

```
// Create an encoder with strength 16
BCryptPasswordEncoder encoder = new BCryptPasswordEncoder(16);
String result = encoder.encode("myPassword");
assertTrue(encoder.matches("myPassword", result));
```

The `Pbkdf2PasswordEncoder` implementation uses PBKDF2 algorithm to hash the passwords. In order to defeat password cracking PBKDF2 is a deliberately slow algorithm and should be tuned to take about .5 seconds to verify a password on your system.

```
// Create an encoder with all the defaults  
Pbkdf2PasswordEncoder encoder = new Pbkdf2PasswordEncoder();  
String result = encoder.encode("myPassword");  
assertTrue(encoder.matches("myPassword", result));
```

20. Appendix

20.1 Security Database Schema

There are various database schema used by the framework and this appendix provides a single reference point to them all. You only need to provide the tables for the areas of functionality you require.

DDL statements are given for the HSQLDB database. You can use these as a guideline for defining the schema for the database you are using.

User Schema

The standard JDBC implementation of the `UserDetailsService` (`JdbcDaoImpl`) requires tables to load the password, account status (enabled or disabled) and a list of authorities (roles) for the user. You will need to adjust this schema to match the database dialect you are using.

```
create table users(
    username varchar_ignorecase(50) not null primary key,
    password varchar_ignorecase(50) not null,
    enabled boolean not null
);

create table authorities (
    username varchar_ignorecase(50) not null,
    authority varchar_ignorecase(50) not null,
    constraint fk_authorities_users foreign key(username) references users(username)
);

create unique index ix_auth_username on authorities (username,authority);
```

For Oracle database

```
CREATE TABLE USERS (
    USERNAME NVARCHAR2(128) PRIMARY KEY,
    PASSWORD NVARCHAR2(128) NOT NULL,
    ENABLED CHAR(1) CHECK (ENABLED IN ('Y','N')) NOT NULL
);

CREATE TABLE AUTHORITIES (
    USERNAME NVARCHAR2(128) NOT NULL,
    AUTHORITY NVARCHAR2(128) NOT NULL
);

ALTER TABLE AUTHORITIES ADD CONSTRAINT AUTHORITIES_UNIQUE UNIQUE (USERNAME, AUTHORITY);
ALTER TABLE AUTHORITIES ADD CONSTRAINT AUTHORITIES_FK1 FOREIGN KEY (USERNAME) REFERENCES USERS
(USERNAME) ENABLE;
```

Group Authorities

Spring Security 2.0 introduced support for group authorities in `JdbcDaoImpl`. The table structure if groups are enabled is as follows. You will need to adjust this schema to match the database dialect you are using.


```

create table groups (
  id bigint generated by default as identity(start with 0) primary key,
  group_name varchar_ignorecase(50) not null
);

create table group_authorities (
  group_id bigint not null,
  authority varchar(50) not null,
  constraint fk_group_authorities_group foreign key(group_id) references groups(id)
);

create table group_members (
  id bigint generated by default as identity(start with 0) primary key,
  username varchar(50) not null,
  group_id bigint not null,
  constraint fk_group_members_group foreign key(group_id) references groups(id)
);

```

Remember that these tables are only required if you are using the provided JDBC `UserDetailsService` implementation. If you write your own or choose to implement `AuthenticationProvider` without a `UserDetailsService`, then you have complete freedom over how you store the data, as long as the interface contract is satisfied.

Persistent Login (Remember-Me) Schema

This table is used to store data used by the more secure [persistent token](#) remember-me implementation. If you are using `JdbcTokenRepositoryImpl` either directly or through the namespace, then you will need this table. Remember to adjust this schema to match the database dialect you are using.

```

create table persistent_logins (
  username varchar(64) not null,
  series varchar(64) primary key,
  token varchar(64) not null,
  last_used timestamp not null
);

```

ACL Schema

There are four tables used by the Spring Security [ACL](#) implementation.

1. `acl_sid` stores the security identities recognised by the ACL system. These can be unique principals or authorities which may apply to multiple principals.
2. `acl_class` defines the domain object types to which ACLs apply. The `class` column stores the Java class name of the object.
3. `acl_object_identity` stores the object identity definitions of specific domain objects.
4. `acl_entry` stores the ACL permissions which apply to a specific object identity and security identity.

It is assumed that the database will auto-generate the primary keys for each of the identities. The `JdbcMutableAclService` has to be able to retrieve these when it has created a new row in the `acl_sid` or `acl_class` tables. It has two properties which define the SQL needed to retrieve these values `classIdentityQuery` and `sidIdentityQuery`. Both of these default to call `identity()`

The ACL artifact JAR contains files for creating the ACL schema in HyperSQL (HSQLDB), PostgreSQL, MySQL/MariaDB, Microsoft SQL Server, and Oracle Database. These schemas are also demonstrated in the following sections.

HyperSQL

The default schema works with the embedded HSQLDB database that is used in unit tests within the framework.

```
create table acl_sid(
  id bigint generated by default as identity(start with 100) not null primary key,
  principal boolean not null,
  sid varchar_ignorecase(100) not null,
  constraint unique_uk_1 unique(sid,principal)
);

create table acl_class(
  id bigint generated by default as identity(start with 100) not null primary key,
  class varchar_ignorecase(100) not null,
  constraint unique_uk_2 unique(class)
);

create table acl_object_identity(
  id bigint generated by default as identity(start with 100) not null primary key,
  object_id_class bigint not null,
  object_id_identity varchar_ignorecase(36) not null,
  parent_object bigint,
  owner_sid bigint,
  entries_inheriting boolean not null,
  constraint unique_uk_3 unique(object_id_class,object_id_identity),
  constraint foreign_fk_1 foreign key(parent_object)references acl_object_identity(id),
  constraint foreign_fk_2 foreign key(object_id_class)references acl_class(id),
  constraint foreign_fk_3 foreign key(owner_sid)references acl_sid(id)
);

create table acl_entry(
  id bigint generated by default as identity(start with 100) not null primary key,
  acl_object_identity bigint not null,
  ace_order int not null,
  sid bigint not null,
  mask integer not null,
  granting boolean not null,
  audit_success boolean not null,
  audit_failure boolean not null,
  constraint unique_uk_4 unique(acl_object_identity,ace_order),
  constraint foreign_fk_4 foreign key(acl_object_identity) references acl_object_identity(id),
  constraint foreign_fk_5 foreign key(sid) references acl_sid(id)
);
```

PostgreSQL

```

create table acl_sid(
    id bigserial not null primary key,
    principal boolean not null,
    sid varchar(100) not null,
    constraint unique_uk_1 unique(sid,principal)
);

create table acl_class(
    id bigserial not null primary key,
    class varchar(100) not null,
    constraint unique_uk_2 unique(class)
);

create table acl_object_identity(
    id bigserial primary key,
    object_id_class bigint not null,
    object_id_identity varchar(36) not null,
    parent_object bigint,
    owner_sid bigint,
    entries_inheriting boolean not null,
    constraint unique_uk_3 unique(object_id_class,object_id_identity),
    constraint foreign_fk_1 foreign key(parent_object)references acl_object_identity(id),
    constraint foreign_fk_2 foreign key(object_id_class)references acl_class(id),
    constraint foreign_fk_3 foreign key(owner_sid)references acl_sid(id)
);

create table acl_entry(
    id bigserial primary key,
    acl_object_identity bigint not null,
    ace_order int not null,
    sid bigint not null,
    mask integer not null,
    granting boolean not null,
    audit_success boolean not null,
    audit_failure boolean not null,
    constraint unique_uk_4 unique(acl_object_identity,ace_order),
    constraint foreign_fk_4 foreign key(acl_object_identity) references acl_object_identity(id),
    constraint foreign_fk_5 foreign key(sid) references acl_sid(id)
);

```

You will have to set the `classIdentityQuery` and `sidIdentityQuery` properties of `JdbcMutableAclService` to the following values, respectively:

- `select currval(pg_get_serial_sequence('acl_class', 'id'))`
- `select currval(pg_get_serial_sequence('acl_sid', 'id'))`

MySQL and MariaDB

```
CREATE TABLE acl_sid (
  id BIGINT UNSIGNED NOT NULL AUTO_INCREMENT PRIMARY KEY,
  principal BOOLEAN NOT NULL,
  sid VARCHAR(100) NOT NULL,
  UNIQUE KEY unique_acl_sid (sid, principal)
) ENGINE=InnoDB;

CREATE TABLE acl_class (
  id BIGINT UNSIGNED NOT NULL AUTO_INCREMENT PRIMARY KEY,
  class VARCHAR(100) NOT NULL,
  UNIQUE KEY uk_acl_class (class)
) ENGINE=InnoDB;

CREATE TABLE acl_object_identity (
  id BIGINT UNSIGNED NOT NULL AUTO_INCREMENT PRIMARY KEY,
  object_id_class BIGINT UNSIGNED NOT NULL,
  object_id_identity VARCHAR(36) NOT NULL,
  parent_object BIGINT UNSIGNED,
  owner_sid BIGINT UNSIGNED,
  entries_inheriting BOOLEAN NOT NULL,
  UNIQUE KEY uk_acl_object_identity (object_id_class, object_id_identity),
  CONSTRAINT fk_acl_object_identity_parent FOREIGN KEY (parent_object) REFERENCES acl_object_identity
(id),
  CONSTRAINT fk_acl_object_identity_class FOREIGN KEY (object_id_class) REFERENCES acl_class (id),
  CONSTRAINT fk_acl_object_identity_owner FOREIGN KEY (owner_sid) REFERENCES acl_sid (id)
) ENGINE=InnoDB;

CREATE TABLE acl_entry (
  id BIGINT UNSIGNED NOT NULL AUTO_INCREMENT PRIMARY KEY,
  acl_object_identity BIGINT UNSIGNED NOT NULL,
  ace_order INTEGER NOT NULL,
  sid BIGINT UNSIGNED NOT NULL,
  mask INTEGER UNSIGNED NOT NULL,
  granting BOOLEAN NOT NULL,
  audit_success BOOLEAN NOT NULL,
  audit_failure BOOLEAN NOT NULL,
  UNIQUE KEY unique_acl_entry (acl_object_identity, ace_order),
  CONSTRAINT fk_acl_entry_object FOREIGN KEY (acl_object_identity) REFERENCES acl_object_identity
(id),
  CONSTRAINT fk_acl_entry_acl FOREIGN KEY (sid) REFERENCES acl_sid (id)
) ENGINE=InnoDB;
```

Microsoft SQL Server

```
CREATE TABLE acl_sid (
    id BIGINT NOT NULL IDENTITY PRIMARY KEY,
    principal BIT NOT NULL,
    sid VARCHAR(100) NOT NULL,
    CONSTRAINT unique_acl_sid UNIQUE (sid, principal)
);

CREATE TABLE acl_class (
    id BIGINT NOT NULL IDENTITY PRIMARY KEY,
    class VARCHAR(100) NOT NULL,
    CONSTRAINT uk_acl_class UNIQUE (class)
);

CREATE TABLE acl_object_identity (
    id BIGINT NOT NULL IDENTITY PRIMARY KEY,
    object_id_class BIGINT NOT NULL,
    object_id_identity VARCHAR(36) NOT NULL,
    parent_object BIGINT,
    owner_sid BIGINT,
    entries_inheriting BIT NOT NULL,
    CONSTRAINT uk_acl_object_identity UNIQUE (object_id_class, object_id_identity),
    CONSTRAINT fk_acl_object_identity_parent FOREIGN KEY (parent_object) REFERENCES acl_object_identity
(id),
    CONSTRAINT fk_acl_object_identity_class FOREIGN KEY (object_id_class) REFERENCES acl_class (id),
    CONSTRAINT fk_acl_object_identity_owner FOREIGN KEY (owner_sid) REFERENCES acl_sid (id)
);

CREATE TABLE acl_entry (
    id BIGINT NOT NULL IDENTITY PRIMARY KEY,
    acl_object_identity BIGINT NOT NULL,
    ace_order INTEGER NOT NULL,
    sid BIGINT NOT NULL,
    mask INTEGER NOT NULL,
    granting BIT NOT NULL,
    audit_success BIT NOT NULL,
    audit_failure BIT NOT NULL,
    CONSTRAINT unique_acl_entry UNIQUE (acl_object_identity, ace_order),
    CONSTRAINT fk_acl_entry_object FOREIGN KEY (acl_object_identity) REFERENCES acl_object_identity
(id),
    CONSTRAINT fk_acl_entry_acl FOREIGN KEY (sid) REFERENCES acl_sid (id)
);
```

Oracle Database

```

CREATE TABLE ACL_SID (
  ID NUMBER(18) PRIMARY KEY,
  PRINCIPAL NUMBER(1) NOT NULL CHECK (PRINCIPAL IN (0, 1)),
  SID NVARCHAR2(128) NOT NULL,
  CONSTRAINT ACL_SID_UNIQUE UNIQUE (SID, PRINCIPAL)
);
CREATE SEQUENCE ACL_SID_SQ START WITH 1 INCREMENT BY 1 NOMAXVALUE;
CREATE OR REPLACE TRIGGER ACL_SID_SQ_TR BEFORE INSERT ON ACL_SID FOR EACH ROW
BEGIN
  SELECT ACL_SID_SQ.NEXTVAL INTO :NEW.ID FROM DUAL;
END;

CREATE TABLE ACL_CLASS (
  ID NUMBER(18) PRIMARY KEY,
  CLASS NVARCHAR2(128) NOT NULL,
  CONSTRAINT ACL_CLASS_UNIQUE UNIQUE (CLASS)
);
CREATE SEQUENCE ACL_CLASS_SQ START WITH 1 INCREMENT BY 1 NOMAXVALUE;
CREATE OR REPLACE TRIGGER ACL_CLASS_ID_TR BEFORE INSERT ON ACL_CLASS FOR EACH ROW
BEGIN
  SELECT ACL_CLASS_SQ.NEXTVAL INTO :NEW.ID FROM DUAL;
END;

CREATE TABLE ACL_OBJECT_IDENTITY(
  ID NUMBER(18) PRIMARY KEY,
  OBJECT_ID_CLASS NUMBER(18) NOT NULL,
  OBJECT_ID_IDENTITY NVARCHAR2(64) NOT NULL,
  PARENT_OBJECT NUMBER(18),
  OWNER_SID NUMBER(18),
  ENTRIES_INHERITING NUMBER(1) NOT NULL CHECK (ENTRIES_INHERITING IN (0, 1)),
  CONSTRAINT ACL_OBJECT_IDENTITY_UNIQUE UNIQUE (OBJECT_ID_CLASS, OBJECT_ID_IDENTITY),
  CONSTRAINT ACL_OBJECT_IDENTITY_PARENT_FK FOREIGN KEY (PARENT_OBJECT) REFERENCES
  ACL_OBJECT_IDENTITY(ID),
  CONSTRAINT ACL_OBJECT_IDENTITY_CLASS_FK FOREIGN KEY (OBJECT_ID_CLASS) REFERENCES ACL_CLASS(ID),
  CONSTRAINT ACL_OBJECT_IDENTITY_OWNER_FK FOREIGN KEY (OWNER_SID) REFERENCES ACL_SID(ID)
);
CREATE SEQUENCE ACL_OBJECT_IDENTITY_SQ START WITH 1 INCREMENT BY 1 NOMAXVALUE;
CREATE OR REPLACE TRIGGER ACL_OBJECT_IDENTITY_ID_TR BEFORE INSERT ON ACL_OBJECT_IDENTITY FOR EACH ROW
BEGIN
  SELECT ACL_OBJECT_IDENTITY_SQ.NEXTVAL INTO :NEW.ID FROM DUAL;
END;

CREATE TABLE ACL_ENTRY (
  ID NUMBER(18) NOT NULL PRIMARY KEY,
  ACL_OBJECT_IDENTITY NUMBER(18) NOT NULL,
  ACE_ORDER INTEGER NOT NULL,
  SID NUMBER(18) NOT NULL,
  MASK INTEGER NOT NULL,
  GRANTING NUMBER(1) NOT NULL CHECK (GRANTING IN (0, 1)),
  AUDIT_SUCCESS NUMBER(1) NOT NULL CHECK (AUDIT_SUCCESS IN (0, 1)),
  AUDIT_FAILURE NUMBER(1) NOT NULL CHECK (AUDIT_FAILURE IN (0, 1)),
  CONSTRAINT ACL_ENTRY_UNIQUE UNIQUE (ACL_OBJECT_IDENTITY, ACE_ORDER),
  CONSTRAINT ACL_ENTRY_OBJECT_FK FOREIGN KEY (ACL_OBJECT_IDENTITY) REFERENCES ACL_OBJECT_IDENTITY
  (ID),
  CONSTRAINT ACL_ENTRY_ACL_FK FOREIGN KEY (SID) REFERENCES ACL_SID(ID)
);
CREATE SEQUENCE ACL_ENTRY_SQ START WITH 1 INCREMENT BY 1 NOMAXVALUE;
CREATE OR REPLACE TRIGGER ACL_ENTRY_ID_TRIGGER BEFORE INSERT ON ACL_ENTRY FOR EACH ROW
BEGIN
  SELECT ACL_ENTRY_SQ.NEXTVAL INTO :NEW.ID FROM DUAL;
END;

```

20.2 The Security Namespace

This appendix provides a reference to the elements available in the security namespace and information on the underlying beans they create (a knowledge of the individual classes and how they work together is assumed - you can find more information in the project Javadoc and elsewhere in this document). If you haven't used the namespace before, please read the [introductory chapter](#) on namespace configuration, as this is intended as a supplement to the information there. Using a good quality XML editor while editing a configuration based on the schema is recommended as this will provide contextual information on which elements and attributes are available as well as comments explaining their purpose. The namespace is written in [RELAX NG Compact](#) format and later converted into an XSD schema. If you are familiar with this format, you may wish to examine the [schema file](#) directly.

Web Application Security

<debug>

Enables Spring Security debugging infrastructure. This will provide human-readable (multi-line) debugging information to monitor requests coming into the security filters. This may include sensitive information, such as request parameters or headers, and should only be used in a development environment.

<http>

If you use an `<http>` element within your application, a `FilterChainProxy` bean named "springSecurityFilterChain" is created and the configuration within the element is used to build a filter chain within `FilterChainProxy`. As of Spring Security 3.1, additional `http` elements can be used to add extra filter chains³. Some core filters are always created in a filter chain and others will be added to the stack depending on the attributes and child elements which are present. The positions of the standard filters are fixed (see [the filter order table](#) in the namespace introduction), removing a common source of errors with previous versions of the framework when users had to configure the filter chain explicitly in the `FilterChainProxy` bean. You can, of course, still do this if you need full control of the configuration.

All filters which require a reference to the `AuthenticationManager` will be automatically injected with the internal instance created by the namespace configuration (see the [introductory chapter](#) for more on the `AuthenticationManager`).

Each `<http>` namespace block always creates an `SecurityContextPersistenceFilter`, an `ExceptionTranslationFilter` and a `FilterSecurityInterceptor`. These are fixed and cannot be replaced with alternatives.

<http> Attributes

The attributes on the `<http>` element control some of the properties on the core filters.

- **access-decision-manager-ref** Optional attribute specifying the ID of the `AccessDecisionManager` implementation which should be used for authorizing HTTP requests. By default an `AffirmativeBased` implementation is used for with a `RoleVoter` and an `AuthenticatedVoter`.
- **authentication-manager-ref** A reference to the `AuthenticationManager` used for the `FilterChain` created by this `http` element.

³See the [introductory chapter](#) for how to set up the mapping from your `web.xml`

- **auto-config** Automatically registers a login form, BASIC authentication, logout services. If set to "true", all of these capabilities are added (although you can still customize the configuration of each by providing the respective element). If unspecified, defaults to "false". Use of this attribute is not recommended. Use explicit configuration elements instead to avoid confusion.
- **create-session** Controls the eagerness with which an HTTP session is created by Spring Security classes. Options include:
 - `always` - Spring Security will proactively create a session if one does not exist.
 - `ifRequired` - Spring Security will only create a session only if one is required (default value).
 - `never` - Spring Security will never create a session, but will make use of one if the application does.
 - `stateless` - Spring Security will not create a session and ignore the session for obtaining a Spring Authentication.
- **disable-url-rewriting** Prevents session IDs from being appended to URLs in the application. Clients must use cookies if this attribute is set to `true`. The default is `true`.
- **entry-point-ref** Normally the `AuthenticationEntryPoint` used will be set depending on which authentication mechanisms have been configured. This attribute allows this behaviour to be overridden by defining a customized `AuthenticationEntryPoint` bean which will start the authentication process.
- **jaas-api-provision** If available, runs the request as the `Subject` acquired from the `JaasAuthenticationToken` which is implemented by adding a `JaasApiIntegrationFilter` bean to the stack. Defaults to `false`.
- **name** A bean identifier, used for referring to the bean elsewhere in the context.
- **once-per-request** Corresponds to the `observeOncePerRequest` property of `FilterSecurityInterceptor`. Defaults to `true`.
- **pattern** Defining a pattern for the [http](#) element controls the requests which will be filtered through the list of filters which it defines. The interpretation is dependent on the configured [request-matcher](#). If no pattern is defined, all requests will be matched, so the most specific patterns should be declared first.
- **realm** Sets the realm name used for basic authentication (if enabled). Corresponds to the `realmName` property on `BasicAuthenticationEntryPoint`.
- **request-matcher** Defines the `RequestMatcher` strategy used in the `FilterChainProxy` and the beans created by the `intercept-url` to match incoming requests. Options are currently `mvc`, `ant`, `regex` and `ciRegex`, for Spring MVC, ant, regular-expression and case-insensitive regular-expression respectively. A separate instance is created for each [intercept-url](#) element using its [pattern](#), [method](#) and [servlet-path](#) attributes. Ant paths are matched using an `AntPathRequestMatcher`, regular expressions are matched using a `RegexRequestMatcher` and for Spring MVC path matching the `MvcRequestMatcher` is used. See the Javadoc for these classes for more details on exactly how the matching is performed. Ant paths are the default strategy.
- **request-matcher-ref** A reference to a bean that implements `RequestMatcher` that will determine if this `FilterChain` should be used. This is a more powerful alternative to [pattern](#).
- **security** A request pattern can be mapped to an empty filter chain, by setting this attribute to `none`. No security will be applied and none of Spring Security's features will be available.

- **security-context-repository-ref** Allows injection of a custom `SecurityContextRepository` into the `SecurityContextPersistenceFilter`.
- **servlet-api-provision** Provides versions of `HttpServletRequest` security methods such as `isUserInRole()` and `getPrincipal()` which are implemented by adding a `SecurityContextHolderAwareRequestFilter` bean to the stack. Defaults to `true`.
- **use-expressions** Enables EL-expressions in the `access` attribute, as described in the chapter on [expression-based access-control](#). The default value is `true`.

Child Elements of `<http>`

- [access-denied-handler](#)
- [anonymous](#)
- [cors](#)
- [csrf](#)
- [custom-filter](#)
- [expression-handler](#)
- [form-login](#)
- [headers](#)
- [http-basic](#)
- [intercept-url](#)
- [jee](#)
- [logout](#)
- [openid-login](#)
- [port-mappings](#)
- [remember-me](#)
- [request-cache](#)
- [session-management](#)
- [x509](#)

`<access-denied-handler>`

This element allows you to set the `errorPage` property for the default `AccessDeniedHandler` used by the `ExceptionHandlerFilter`, using the [error-page](#) attribute, or to supply your own implementation using the [ref](#) attribute. This is discussed in more detail in the section on the [ExceptionHandlerFilter](#).

Parent Elements of `<access-denied-handler>`

- [http](#)

<access-denied-handler> Attributes

- **error-page** The access denied page that an authenticated user will be redirected to if they request a page which they don't have the authority to access.
- **ref** Defines a reference to a Spring bean of type `AccessDeniedHandler`.

<cors>

This element allows for configuring a `CorsFilter`. If no `CorsFilter` or `CorsConfigurationSource` is specified and Spring MVC is on the classpath, a `HandlerMappingIntrospector` is used as the `CorsConfigurationSource`.

<cors> Attributes

The attributes on the `<cors>` element control the headers element.

- **ref** Optional attribute that specifies the bean name of a `CorsFilter`.
- **cors-configuration-source-ref** Optional attribute that specifies the bean name of a `CorsConfigurationSource` to be injected into a `CorsFilter` created by the XML namespace.

Parent Elements of <cors>

- [http](#)

<headers>

This element allows for configuring additional (security) headers to be send with the response. It enables easy configuration for several headers and also allows for setting custom headers through the [header](#) element. Additional information, can be found in the [Security Headers](#) section of the reference.

- `Cache-Control`, `Pragma`, and `Expires` - Can be set using the [cache-control](#) element. This ensures that the browser does not cache your secured pages.
- `Strict-Transport-Security` - Can be set using the [hsts](#) element. This ensures that the browser automatically requests HTTPS for future requests.
- `X-Frame-Options` - Can be set using the [frame-options](#) element. The [X-Frame-Options](#) header can be used to prevent clickjacking attacks.
- `X-XSS-Protection` - Can be set using the [xss-protection](#) element. The [X-XSS-Protection](#) header can be used by browser to do basic control.
- `X-Content-Type-Options` - Can be set using the [content-type-options](#) element. The [X-Content-Type-Options](#) header prevents Internet Explorer from MIME-sniffing a response away from the declared content-type. This also applies to Google Chrome, when downloading extensions.
- `Public-Key-Pinning` or `Public-Key-Pinning-Report-Only` - Can be set using the [hpkp](#) element. This allows HTTPS websites to resist impersonation by attackers using mis-issued or otherwise fraudulent certificates.
- `Content-Security-Policy` or `Content-Security-Policy-Report-Only` - Can be set using the [content-security-policy](#) element. [Content Security Policy \(CSP\)](#) is a mechanism that web applications can leverage to mitigate content injection vulnerabilities, such as cross-site scripting (XSS).

- `Referrer-Policy` - Can be set using the [referrer-policy](#) element, [Referrer-Policy](#) is a mechanism that web applications can leverage to manage the referrer field, which contains the last page the user was on.
- `Feature-Policy` - Can be set using the [feature-policy](#) element, [Feature-Policy](#) is a mechanism that allows web developers to selectively enable, disable, and modify the behavior of certain APIs and web features in the browser.

<headers> Attributes

The attributes on the `<headers>` element control the headers element.

- **defaults-disabled** Optional attribute that specifies to disable the default Spring Security's HTTP response headers. The default is false (the default headers are included).
- **disabled** Optional attribute that specifies to disable Spring Security's HTTP response headers. The default is false (the headers are enabled).

Parent Elements of <headers>

- [http](#)

Child Elements of <headers>

- [cache-control](#)
- [content-security-policy](#)
- [content-type-options](#)
- [feature-policy](#)
- [frame-options](#)
- [header](#)
- [hpkp](#)
- [hsts](#)
- [referrer-policy](#)
- [xss-protection](#)

<cache-control>

Adds `Cache-Control`, `Pragma`, and `Expires` headers to ensure that the browser does not cache your secured pages.

<cache-control> Attributes

- **disabled** Specifies if Cache Control should be disabled. Default false.

Parent Elements of <cache-control>

- [headers](#)

<hsts>

When enabled adds the [Strict-Transport-Security](#) header to the response for any secure request. This allows the server to instruct browsers to automatically use HTTPS for future requests.

<hsts> Attributes

- **disabled** Specifies if Strict-Transport-Security should be disabled. Default false.
- **include-sub-domains** Specifies if subdomains should be included. Default true.
- **max-age-seconds** Specifies the maximum amount of time the host should be considered a Known HSTS Host. Default one year.
- **request-matcher-ref** The RequestMatcher instance to be used to determine if the header should be set. Default is if `HttpServletRequest.isSecure()` is true.
- **preload** Specifies if preload should be included. Default false.

Parent Elements of <hsts>

- [headers](#)

<hpkp>

When enabled adds the [Public Key Pinning Extension for HTTP](#) header to the response for any secure request. This allows HTTPS websites to resist impersonation by attackers using mis-issued or otherwise fraudulent certificates.

<hpkp> Attributes

- **disabled** Specifies if HTTP Public Key Pinning (HPKP) should be disabled. Default true.
- **include-sub-domains** Specifies if subdomains should be included. Default false.
- **max-age-seconds** Sets the value for the max-age directive of the Public-Key-Pins header. Default 60 days.
- **report-only** Specifies if the browser should only report pin validation failures. Default true.
- **report-uri** Specifies the URI to which the browser should report pin validation failures.

Parent Elements of <hpkp>

- [headers](#)

<pins>

The list of pins

Child Elements of <pins>

- [pin](#)

<pin>

A pin is specified using the base64-encoded SPKI fingerprint as value and the cryptographic hash algorithm as attribute

<pin> Attributes

- **algorithm** The cryptographic hash algorithm. Default is SHA256.

Parent Elements of <pin>

- [pins](#)

<content-security-policy>

When enabled adds the [Content Security Policy \(CSP\)](#) header to the response. CSP is a mechanism that web applications can leverage to mitigate content injection vulnerabilities, such as cross-site scripting (XSS).

<content-security-policy> Attributes

- **policy-directives** The security policy directive(s) for the Content-Security-Policy header or if report-only is set to true, then the Content-Security-Policy-Report-Only header is used.
- **report-only** Set to true, to enable the Content-Security-Policy-Report-Only header for reporting policy violations only. Defaults to false.

Parent Elements of <content-security-policy>

- [headers](#)

<referrer-policy>

When enabled adds the [Referrer Policy](#) header to the response.

<referrer-policy> Attributes

- **policy** The policy for the Referrer-Policy header. Default "no-referrer".

Parent Elements of <referrer-policy>

- [headers](#)

<feature-policy>

When enabled adds the [Feature Policy](#) header to the response.

<feature-policy> Attributes

- **policy-directives** The security policy directive(s) for the Feature-Policy header.

Parent Elements of <feature-policy>

- [headers](#)

<frame-options>

When enabled adds the [X-Frame-Options header](#) to the response, this allows newer browsers to do some security checks and prevent [clickjacking](#) attacks.

<frame-options> Attributes

- **disabled** If disabled, the X-Frame-Options header will not be included. Default false.

- **policy**

- **DENY** The page cannot be displayed in a frame, regardless of the site attempting to do so. This is the default when `frame-options-policy` is specified.
- **SAMEORIGIN** The page can only be displayed in a frame on the same origin as the page itself
- **ALLOW-FROM** `origin` The page can only be displayed in a frame on the specified origin.

In other words, if you specify DENY, not only will attempts to load the page in a frame fail when loaded from other sites, attempts to do so will fail when loaded from the same site. On the other hand, if you specify SAMEORIGIN, you can still use the page in a frame as long as the site including it in a frame it is the same as the one serving the page.

- **strategy** Select the `AllowFromStrategy` to use when using the ALLOW-FROM policy.
 - `static` Use a single static ALLOW-FROM value. The value can be set through the [value](#) attribute.
 - `regexp` Use a regular expression to validate incoming requests and if they are allowed. The regular expression can be set through the [value](#) attribute. The request parameter used to retrieve the value to validate can be specified using the [from-parameter](#).
 - `whitelist` A comma-separated list containing the allowed domains. The comma-separated list can be set through the [value](#) attribute. The request parameter used to retrieve the value to validate can be specified using the [from-parameter](#).
- **ref** Instead of using one of the predefined strategies it is also possible to use a custom `AllowFromStrategy`. The reference to this bean can be specified through this `ref` attribute.
- **value** The value to use when ALLOW-FROM is used a [strategy](#).
- **from-parameter** Specify the name of the request parameter to use when using `regexp` or `whitelist` for the ALLOW-FROM strategy.

Parent Elements of `<frame-options>`

- [headers](#)

`<xss-protection>`

Adds the [X-XSS-Protection header](#) to the response to assist in protecting against [reflected / Type-1 Cross-Site Scripting \(XSS\)](#) attacks. This is in no-way a full protection to XSS attacks!

`<xss-protection>` Attributes

- **xss-protection-disabled** Do not include the header for [reflected / Type-1 Cross-Site Scripting \(XSS\)](#) protection.
- **xss-protection-enabled** Explicitly enable or disable [reflected / Type-1 Cross-Site Scripting \(XSS\)](#) protection.
- **xss-protection-block** When true and `xss-protection-enabled` is true, adds `mode=block` to the header. This indicates to the browser that the page should not be loaded at all. When false and `xss-protection-enabled` is true, the page will still be rendered when an reflected attack is detected but the response will be modified to protect against the attack. Note that there are sometimes ways of bypassing this mode which can often times make blocking the page more desirable.

Parent Elements of <xss-protection>

- [headers](#)

<content-type-options>

Add the X-Content-Type-Options header with the value of nosniff to the response. This [disables MIME-sniffing](#) for IE8+ and Chrome extensions.

<content-type-options> Attributes

- **disabled** Specifies if Content Type Options should be disabled. Default false.

Parent Elements of <content-type-options>

- [headers](#)

<header>

Add additional headers to the response, both the name and value need to be specified.

<header-attributes> Attributes

- **header-name** The name of the header.
- **value** The value of the header to add.
- **ref** Reference to a custom implementation of the `HeaderWriter` interface.

Parent Elements of <header>

- [headers](#)

<anonymous>

Adds an `AnonymousAuthenticationFilter` to the stack and an `AnonymousAuthenticationProvider`. Required if you are using the `IS_AUTHENTICATED_ANONYMOUSLY` attribute.

Parent Elements of <anonymous>

- [http](#)

<anonymous> Attributes

- **enabled** With the default namespace setup, the anonymous "authentication" facility is automatically enabled. You can disable it using this property.
- **granted-authority** The granted authority that should be assigned to the anonymous request. Commonly this is used to assign the anonymous request particular roles, which can subsequently be used in authorization decisions. If unset, defaults to `ROLE_ANONYMOUS`.
- **key** The key shared between the provider and filter. This generally does not need to be set. If unset, it will default to a secure randomly generated value. This means setting this value can improve startup time when using the anonymous functionality since secure random values can take a while to be generated.

- **username** The username that should be assigned to the anonymous request. This allows the principal to be identified, which may be important for logging and auditing. If unset, defaults to `anonymousUser`.

<csrf>

This element will add [Cross Site Request Forger \(CSRF\)](#) protection to the application. It also updates the default RequestCache to only replay "GET" requests upon successful authentication. Additional information can be found in the [Cross Site Request Forgery \(CSRF\)](#) section of the reference.

Parent Elements of <csrf>

- [http](#)

<csrf> Attributes

- **disabled** Optional attribute that specifies to disable Spring Security's CSRF protection. The default is false (CSRF protection is enabled). It is highly recommended to leave CSRF protection enabled.
- **token-repository-ref** The `CsrfTokenRepository` to use. The default is `HttpSessionCsrfTokenRepository`.
- **request-matcher-ref** The `RequestMatcher` instance to be used to determine if CSRF should be applied. Default is any HTTP method except "GET", "TRACE", "HEAD", "OPTIONS".

<custom-filter>

This element is used to add a filter to the filter chain. It doesn't create any additional beans but is used to select a bean of type `javax.servlet.Filter` which is already defined in the application context and add that at a particular position in the filter chain maintained by Spring Security. Full details can be found in the [namespace chapter](#).

Parent Elements of <custom-filter>

- [http](#)

<custom-filter> Attributes

- **after** The filter immediately after which the custom-filter should be placed in the chain. This feature will only be needed by advanced users who wish to mix their own filters into the security filter chain and have some knowledge of the standard Spring Security filters. The filter names map to specific Spring Security implementation filters.
- **before** The filter immediately before which the custom-filter should be placed in the chain
- **position** The explicit position at which the custom-filter should be placed in the chain. Use if you are replacing a standard filter.
- **ref** Defines a reference to a Spring bean that implements `Filter`.

<expression-handler>

Defines the `SecurityExpressionHandler` instance which will be used if expression-based access-control is enabled. A default implementation (with no ACL support) will be used if not supplied.

Parent Elements of <expression-handler>

- [global-method-security](#)
- [http](#)
- [websocket-message-broker](#)

<expression-handler> Attributes

- **ref** Defines a reference to a Spring bean that implements `SecurityExpressionHandler`.

<form-login>

Used to add an `UsernamePasswordAuthenticationFilter` to the filter stack and an `LoginUrlAuthenticationEntryPoint` to the application context to provide authentication on demand. This will always take precedence over other namespace-created entry points. If no attributes are supplied, a login page will be generated automatically at the URL `/login`²³ The behaviour can be customized using the [<form-login> Attributes](#).

Parent Elements of <form-login>

- [http](#)

<form-login> Attributes

- **always-use-default-target** If set to `true`, the user will always start at the value given by [default-target-url](#), regardless of how they arrived at the login page. Maps to the `alwaysUseDefaultTargetUrl` property of `UsernamePasswordAuthenticationFilter`. Default value is `false`.
- **authentication-details-source-ref** Reference to an `AuthenticationDetailsSource` which will be used by the authentication filter
- **authentication-failure-handler-ref** Can be used as an alternative to [authentication-failure-url](#), giving you full control over the navigation flow after an authentication failure. The value should be the name of an `AuthenticationFailureHandler` bean in the application context.
- **authentication-failure-url** Maps to the `authenticationFailureUrl` property of `UsernamePasswordAuthenticationFilter`. Defines the URL the browser will be redirected to on login failure. Defaults to `/login?error`, which will be automatically handled by the automatic login page generator, re-rendering the login page with an error message.
- **authentication-success-handler-ref** This can be used as an alternative to [default-target-url](#) and [always-use-default-target](#), giving you full control over the navigation flow after a successful authentication. The value should be the name of an `AuthenticationSuccessHandler` bean in the application context. By default, an implementation of `SavedRequestAwareAuthenticationSuccessHandler` is used and injected with the [default-target-url](#).
- **default-target-url** Maps to the `defaultTargetUrl` property of `UsernamePasswordAuthenticationFilter`. If not set, the default value is `/` (the application

²³This feature is really just provided for convenience and is not intended for production (where a view technology will have been chosen and can be used to render a customized login page). The class `DefaultLoginPageGeneratingFilter` is responsible for rendering the login page and will provide login forms for both normal form login and/or OpenID if required.

root). A user will be taken to this URL after logging in, provided they were not asked to login while attempting to access a secured resource, when they will be taken to the originally requested URL.

- **login-page** The URL that should be used to render the login page. Maps to the `loginFormUrl` property of the `LoginUrlAuthenticationEntryPoint`. Defaults to `"/login"`.
- **login-processing-url** Maps to the `filterProcessesUrl` property of `UsernamePasswordAuthenticationFilter`. The default value is `"/login"`.
- **password-parameter** The name of the request parameter which contains the password. Defaults to `"password"`.
- **username-parameter** The name of the request parameter which contains the username. Defaults to `"username"`.
- **authentication-success-forward-url** Maps a `ForwardAuthenticationSuccessHandler` to `authenticationSuccessHandler` property of `UsernamePasswordAuthenticationFilter`.
- **authentication-failure-forward-url** Maps a `ForwardAuthenticationFailureHandler` to `authenticationFailureHandler` property of `UsernamePasswordAuthenticationFilter`.

<http-basic>

Adds a `BasicAuthenticationFilter` and `BasicAuthenticationEntryPoint` to the configuration. The latter will only be used as the configuration entry point if form-based login is not enabled.

Parent Elements of <http-basic>

- [http](#)

<http-basic> Attributes

- **authentication-details-source-ref** Reference to an `AuthenticationDetailsSource` which will be used by the authentication filter
- **entry-point-ref** Sets the `AuthenticationEntryPoint` which is used by the `BasicAuthenticationFilter`.

<http-firewall> Element

This is a top-level element which can be used to inject a custom implementation of `HttpFirewall` into the `FilterChainProxy` created by the namespace. The default implementation should be suitable for most applications.

<http-firewall> Attributes

- **ref** Defines a reference to a Spring bean that implements `HttpFirewall`.

<intercept-url>

This element is used to define the set of URL patterns that the application is interested in and to configure how they should be handled. It is used to construct the `FilterInvocationSecurityMetadataSource` used by the `FilterSecurityInterceptor`. It is also responsible for configuring a `ChannelProcessingFilter` if particular URLs need to be accessed by HTTPS, for example. When matching the specified patterns against an incoming request,

the matching is done in the order in which the elements are declared. So the most specific patterns should come first and the most general should come last.

Parent Elements of `<intercept-url>`

- [filter-security-metadata-source](#)
- [http](#)

`<intercept-url>` Attributes

- **access** Lists the access attributes which will be stored in the `FilterInvocationSecurityMetadataSource` for the defined URL pattern/method combination. This should be a comma-separated list of the security configuration attributes (such as role names).
- **method** The HTTP Method which will be used in combination with the pattern and servlet path (optional) to match an incoming request. If omitted, any method will match. If an identical pattern is specified with and without a method, the method-specific match will take precedence.
- **pattern** The pattern which defines the URL path. The content will depend on the `request-matcher` attribute from the containing `http` element, so will default to ant path syntax.
- **request-matcher-ref** A reference to a `RequestMatcher` that will be used to determine if this `<intercept-url>` is used.
- **requires-channel** Can be "http" or "https" depending on whether a particular URL pattern should be accessed over HTTP or HTTPS respectively. Alternatively the value "any" can be used when there is no preference. If this attribute is present on any `<intercept-url>` element, then a `ChannelProcessingFilter` will be added to the filter stack and its additional dependencies added to the application context.

If a `<port-mappings>` configuration is added, this will be used to by the `SecureChannelProcessor` and `InsecureChannelProcessor` beans to determine the ports used for redirecting to HTTP/HTTPS.

Note

This property is invalid for [filter-security-metadata-source](#)

- **servlet-path** The servlet path which will be used in combination with the pattern and HTTP method to match an incoming request. This attribute is only applicable when [request-matcher](#) is 'mvc'. In addition, the value is only required in the following 2 use cases: 1) There are 2 or more `HttpServlet`'s registered in the `ServletContext` that have mappings starting with `'/'` and are different; 2) The pattern starts with the same value of a registered `HttpServlet` path, excluding the default (root) `HttpServlet '/'`.

Note

This property is invalid for [filter-security-metadata-source](#)

`<jee>`

Adds a `J2eePreAuthenticatedProcessingFilter` to the filter chain to provide integration with container authentication.

Parent Elements of <jee>

- [http](#)

<jee> Attributes

- **mappable-roles** A comma-separated list of roles to look for in the incoming `HttpServletRequest`.
- **user-service-ref** A reference to a user-service (or `UserDetailsService` bean) Id

<logout>

Adds a `LogoutFilter` to the filter stack. This is configured with a `SecurityContextLogoutHandler`.

Parent Elements of <logout>

- [http](#)

<logout> Attributes

- **delete-cookies** A comma-separated list of the names of cookies which should be deleted when the user logs out.
- **invalidate-session** Maps to the `invalidateHttpSession` of the `SecurityContextLogoutHandler`. Defaults to "true", so the session will be invalidated on logout.
- **logout-success-url** The destination URL which the user will be taken to after logging out. Defaults to `<form-login-login-page>/?logout` (i.e. `/login?logout`)

Setting this attribute will inject the `SessionManagementFilter` with a `SimpleRedirectInvalidSessionStrategy` configured with the attribute value. When an invalid session ID is submitted, the strategy will be invoked, redirecting to the configured URL.

- **logout-url** The URL which will cause a logout (i.e. which will be processed by the filter). Defaults to `/logout`.
- **success-handler-ref** May be used to supply an instance of `LogoutSuccessHandler` which will be invoked to control the navigation after logging out.

<openid-login>

Similar to `<form-login>` and has the same attributes. The default value for `login-processing-url` is `/login/openid`. An `OpenIDAuthenticationFilter` and `OpenIDAuthenticationProvider` will be registered. The latter requires a reference to a `UserDetailsService`. Again, this can be specified by `id`, using the `user-service-ref` attribute, or will be located automatically in the application context.

Parent Elements of <openid-login>

- [http](#)

<openid-login> Attributes

- **always-use-default-target** Whether the user should always be redirected to the `default-target-url` after login.

- **authentication-details-source-ref** Reference to an `AuthenticationDetailsSource` which will be used by the authentication filter
- **authentication-failure-handler-ref** Reference to an `AuthenticationFailureHandler` bean which should be used to handle a failed authentication request. Should not be used in combination with `authentication-failure-url` as the implementation should always deal with navigation to the subsequent destination
- **authentication-failure-url** The URL for the login failure page. If no login failure URL is specified, Spring Security will automatically create a failure login URL at `/login?login_error` and a corresponding filter to render that login failure URL when requested.
- **authentication-success-forward-url** Maps a `ForwardAuthenticationSuccessHandler` to `authenticationSuccessHandler` property of `UsernamePasswordAuthenticationFilter`.
- **authentication-failure-forward-url** Maps a `ForwardAuthenticationFailureHandler` to `authenticationFailureHandler` property of `UsernamePasswordAuthenticationFilter`.
- **authentication-success-handler-ref** Reference to an `AuthenticationSuccessHandler` bean which should be used to handle a successful authentication request. Should not be used in combination with [default-target-url](#) (or [always-use-default-target](#)) as the implementation should always deal with navigation to the subsequent destination
- **default-target-url** The URL that will be redirected to after successful authentication, if the user's previous action could not be resumed. This generally happens if the user visits a login page without having first requested a secured operation that triggers authentication. If unspecified, defaults to the root of the application.
- **login-page** The URL for the login page. If no login URL is specified, Spring Security will automatically create a login URL at `/login` and a corresponding filter to render that login URL when requested.
- **login-processing-url** The URL that the login form is posted to. If unspecified, it defaults to `/login`.
- **password-parameter** The name of the request parameter which contains the password. Defaults to "password".
- **user-service-ref** A reference to a user-service (or `UserDetailsService` bean) Id
- **username-parameter** The name of the request parameter which contains the username. Defaults to "username".

Child Elements of `<openid-login>`

- [attribute-exchange](#)

`<attribute-exchange>`

The `attribute-exchange` element defines the list of attributes which should be requested from the identity provider. An example can be found in the [OpenID Support](#) section of the namespace configuration chapter. More than one can be used, in which case each must have an `identifier-match` attribute, containing a regular expression which is matched against the supplied OpenID identifier. This allows different attribute lists to be fetched from different providers (Google, Yahoo etc).

Parent Elements of `<attribute-exchange>`

- [openid-login](#)

<attribute-exchange> Attributes

- **identifier-match** A regular expression which will be compared against the claimed identity, when deciding which attribute-exchange configuration to use during authentication.

Child Elements of <attribute-exchange>

- [openid-attribute](#)

<openid-attribute>

Attributes used when making an OpenID AX [Fetch Request](#)

Parent Elements of <openid-attribute>

- [attribute-exchange](#)

<openid-attribute> Attributes

- **count** Specifies the number of attributes that you wish to get back. For example, return 3 emails. The default value is 1.
- **name** Specifies the name of the attribute that you wish to get back. For example, email.
- **required** Specifies if this attribute is required to the OP, but does not error out if the OP does not return the attribute. Default is false.
- **type** Specifies the attribute type. For example, <https://axschema.org/contact/email>. See your OP's documentation for valid attribute types.

<port-mappings>

By default, an instance of `PortMapperImpl` will be added to the configuration for use in redirecting to secure and insecure URLs. This element can optionally be used to override the default mappings which that class defines. Each child `<port-mapping>` element defines a pair of HTTP:HTTPS ports. The default mappings are 80:443 and 8080:8443. An example of overriding these can be found in the [namespace introduction](#).

Parent Elements of <port-mappings>

- [http](#)

Child Elements of <port-mappings>

- [port-mapping](#)

<port-mapping>

Provides a method to map http ports to https ports when forcing a redirect.

Parent Elements of <port-mapping>

- [port-mappings](#)

<port-mapping> Attributes

- **http** The http port to use.

- **https** The https port to use.

<remember-me>

Adds the `RememberMeAuthenticationFilter` to the stack. This in turn will be configured with either a `TokenBasedRememberMeServices`, a `PersistentTokenBasedRememberMeServices` or a user-specified bean implementing `RememberMeServices` depending on the attribute settings.

Parent Elements of <remember-me>

- [http](#)

<remember-me> Attributes

- **authentication-success-handler-ref** Sets the `authenticationSuccessHandler` property on the `RememberMeAuthenticationFilter` if custom navigation is required. The value should be the name of a `AuthenticationSuccessHandler` bean in the application context.
- **data-source-ref** A reference to a `DataSource` bean. If this is set, `PersistentTokenBasedRememberMeServices` will be used and configured with a `JdbcTokenRepositoryImpl` instance.
- **remember-me-parameter** The name of the request parameter which toggles remember-me authentication. Defaults to "remember-me". Maps to the "parameter" property of `AbstractRememberMeServices`.
- **remember-me-cookie** The name of cookie which store the token for remember-me authentication. Defaults to "remember-me". Maps to the "cookieName" property of `AbstractRememberMeServices`.
- **key** Maps to the "key" property of `AbstractRememberMeServices`. Should be set to a unique value to ensure that remember-me cookies are only valid within the one application²⁵. If this is not set a secure random value will be generated. Since generating secure random values can take a while, setting this value explicitly can help improve startup times when using the remember-me functionality.
- **services-alias** Exports the internally defined `RememberMeServices` as a bean alias, allowing it to be used by other beans in the application context.
- **services-ref** Allows complete control of the `RememberMeServices` implementation that will be used by the filter. The value should be the id of a bean in the application context which implements this interface. Should also implement `LogoutHandler` if a logout filter is in use.
- **token-repository-ref** Configures a `PersistentTokenBasedRememberMeServices` but allows the use of a custom `PersistentTokenRepository` bean.
- **token-validity-seconds** Maps to the `tokenValiditySeconds` property of `AbstractRememberMeServices`. Specifies the period in seconds for which the remember-me cookie should be valid. By default it will be valid for 14 days.
- **use-secure-cookie** It is recommended that remember-me cookies are only submitted over HTTPS and thus should be flagged as "secure". By default, a secure cookie will be used if the connection over which the login request is made is secure (as it should be). If you set this property to `false`, secure cookies will not be used. Setting it to `true` will always set the secure flag on the cookie. This attribute maps to the `useSecureCookie` property of `AbstractRememberMeServices`.

- **user-service-ref** The remember-me services implementations require access to a `UserDetailsService`, so there has to be one defined in the application context. If there is only one, it will be selected and used automatically by the namespace configuration. If there are multiple instances, you can specify a bean `id` explicitly using this attribute.

<request-cache> Element

Sets the `RequestCache` instance which will be used by the `ExceptionTranslationFilter` to store request information before invoking an `AuthenticationEntryPoint`.

Parent Elements of <request-cache>

- [http](#)

<request-cache> Attributes

- **ref** Defines a reference to a Spring bean that is a `RequestCache`.

<session-management>

Session-management related functionality is implemented by the addition of a `SessionManagementFilter` to the filter stack.

Parent Elements of <session-management>

- [http](#)

<session-management> Attributes

- **invalid-session-url** Setting this attribute will inject the `SessionManagementFilter` with a `SimpleRedirectInvalidSessionStrategy` configured with the attribute value. When an invalid session ID is submitted, the strategy will be invoked, redirecting to the configured URL.
- **invalid-session-url** Allows injection of the `InvalidSessionStrategy` instance used by the `SessionManagementFilter`. Use either this or the `invalid-session-url` attribute but not both.
- **session-authentication-error-url** Defines the URL of the error page which should be shown when the `SessionAuthenticationStrategy` raises an exception. If not set, an unauthorized (401) error code will be returned to the client. Note that this attribute doesn't apply if the error occurs during a form-based login, where the URL for authentication failure will take precedence.
- **session-authentication-strategy-ref** Allows injection of the `SessionAuthenticationStrategy` instance used by the `SessionManagementFilter`
- **session-fixation-protection** Indicates how session fixation protection will be applied when a user authenticates. If set to "none", no protection will be applied. "newSession" will create a new empty session, with only Spring Security-related attributes migrated. "migrateSession" will create a new session and copy all session attributes to the new session. In Servlet 3.1 (Java EE 7) and newer containers, specifying "changeSessionId" will keep the existing session and use the container-supplied session fixation protection (`HttpServletRequest#changeSessionId()`). Defaults to "changeSessionId" in Servlet 3.1 and newer containers, "migrateSession" in older containers. Throws an exception if "changeSessionId" is used in older containers.

If session fixation protection is enabled, the `SessionManagementFilter` is injected with an appropriately configured `DefaultSessionAuthenticationStrategy`. See the Javadoc for this class for more details.

Child Elements of <session-management>

- [concurrency-control](#)

<concurrency-control>

Adds support for concurrent session control, allowing limits to be placed on the number of active sessions a user can have. A `ConcurrentSessionFilter` will be created, and a `ConcurrentSessionControlAuthenticationStrategy` will be used with the `SessionManagementFilter`. If a `form-login` element has been declared, the strategy object will also be injected into the created authentication filter. An instance of `SessionRegistry` (a `SessionRegistryImpl` instance unless the user wishes to use a custom bean) will be created for use by the strategy.

Parent Elements of <concurrency-control>

- [session-management](#)

<concurrency-control> Attributes

- **error-if-maximum-exceeded** If set to "true" a `SessionAuthenticationException` will be raised when a user attempts to exceed the maximum allowed number of sessions. The default behaviour is to expire the original session.
- **expired-url** The URL a user will be redirected to if they attempt to use a session which has been "expired" by the concurrent session controller because the user has exceeded the number of allowed sessions and has logged in again elsewhere. Should be set unless `exception-if-maximum-exceeded` is set. If no value is supplied, an expiry message will just be written directly back to the response.
- **expired-url** Allows injection of the `ExpiredSessionStrategy` instance used by the `ConcurrentSessionFilter`
- **max-sessions** Maps to the `maximumSessions` property of `ConcurrentSessionControlAuthenticationStrategy`. Specify `-1` as the value to support unlimited sessions.
- **session-registry-alias** It can also be useful to have a reference to the internal session registry for use in your own beans or an admin interface. You can expose the internal bean using the `session-registry-alias` attribute, giving it a name that you can use elsewhere in your configuration.
- **session-registry-ref** The user can supply their own `SessionRegistry` implementation using the `session-registry-ref` attribute. The other concurrent session control beans will be wired up to use it.

<x509>

Adds support for X.509 authentication. An `X509AuthenticationFilter` will be added to the stack and an `Http403ForbiddenEntryPoint` bean will be created. The latter will only be used if no other authentication mechanisms are in use (its only functionality is to return an HTTP 403 error code). A `PreAuthenticatedAuthenticationProvider` will also be created which delegates the loading of user authorities to a `UserDetailsService`.

Parent Elements of <x509>

- [http](#)

<x509> Attributes

- **authentication-details-source-ref** A reference to an `AuthenticationDetailsSource`
- **subject-principal-regex** Defines a regular expression which will be used to extract the username from the certificate (for use with the `UserDetailsService`).
- **user-service-ref** Allows a specific `UserDetailsService` to be used with X.509 in the case where multiple instances are configured. If not set, an attempt will be made to locate a suitable instance automatically and use that.

<filter-chain-map>

Used to explicitly configure a `FilterChainProxy` instance with a `FilterChainMap`

<filter-chain-map> Attributes

- **request-matcher** Defines the strategy to use for matching incoming requests. Currently the options are 'ant' (for ant path patterns), 'regex' for regular expressions and 'ciRegex' for case-insensitive regular expressions.

Child Elements of <filter-chain-map>

- [filter-chain](#)

<filter-chain>

Used within to define a specific URL pattern and the list of filters which apply to the URLs matching that pattern. When multiple filter-chain elements are assembled in a list in order to configure a `FilterChainProxy`, the most specific patterns must be placed at the top of the list, with most general ones at the bottom.

Parent Elements of <filter-chain>

- [filter-chain-map](#)

<filter-chain> Attributes

- **filters** A comma separated list of references to Spring beans that implement `Filter`. The value "none" means that no `Filter` should be used for this `FilterChain`.
- **pattern** A pattern that creates `RequestMatcher` in combination with the [request-matcher](#)
- **request-matcher-ref** A reference to a `RequestMatcher` that will be used to determine if any `Filter` from the `filters` attribute should be invoked.

<filter-security-metadata-source>

Used to explicitly configure a `FilterSecurityMetadataSource` bean for use with a `FilterSecurityInterceptor`. Usually only needed if you are configuring a `FilterChainProxy` explicitly, rather than using the `<http>` element. The `intercept-url` elements used should only contain `pattern`, `method` and `access` attributes. Any others will result in a configuration error.

<filter-security-metadata-source> Attributes

- **id** A bean identifier, used for referring to the bean elsewhere in the context.

- **request-matcher** Defines the strategy use for matching incoming requests. Currently the options are 'ant' (for ant path patterns), 'regex' for regular expressions and 'ciRegex' for case-insensitive regular expressions.
- **use-expressions** Enables the use of expressions in the 'access' attributes in <intercept-url> elements rather than the traditional list of configuration attributes. Defaults to 'true'. If enabled, each attribute should contain a single Boolean expression. If the expression evaluates to 'true', access will be granted.

Child Elements of <filter-security-metadata-source>

- [intercept-url](#)

WebSocket Security

Spring Security 4.0+ provides support for authorizing messages. One concrete example of where this is useful is to provide authorization in WebSocket based applications.

<websocket-message-broker>

The websocket-message-broker element has two different modes. If the [websocket-message-broker@id](#) is not specified, then it will do the following things:

- Ensure that any `SimpAnnotationMethodMessageHandler` has the `AuthenticationPrincipalArgumentResolver` registered as a custom argument resolver. This allows the use of `@AuthenticationPrincipal` to resolve the principal of the current `Authentication`
- Ensures that the `SecurityContextChannelInterceptor` is automatically registered for the `clientInboundChannel`. This populates the `SecurityContextHolder` with the user that is found in the `Message`
- Ensures that a `ChannelSecurityInterceptor` is registered with the `clientInboundChannel`. This allows authorization rules to be specified for a message.
- Ensures that a `CsrfChannelInterceptor` is registered with the `clientInboundChannel`. This ensures that only requests from the original domain are enabled.
- Ensures that a `CsrfTokenHandshakeInterceptor` is registered with `WebSocketHttpRequestHandler`, `TransportHandlingSockJsService`, or `DefaultSockJsService`. This ensures that the expected `CsrfToken` from the `HttpServletRequest` is copied into the `WebSocket Session` attributes.

If additional control is necessary, the `id` can be specified and a `ChannelSecurityInterceptor` will be assigned to the specified `id`. All the wiring with Spring's messaging infrastructure can then be done manually. This is more cumbersome, but provides greater control over the configuration.

<websocket-message-broker> Attributes

- **id** A bean identifier, used for referring to the `ChannelSecurityInterceptor` bean elsewhere in the context. If specified, Spring Security requires explicit configuration within Spring Messaging. If not specified, Spring Security will automatically integrate with the messaging infrastructure as described in the section called "<websocket-message-broker>"
- **same-origin-disabled** Disables the requirement for CSRF token to be present in the Stomp headers (default false). Changing the default is useful if it is necessary to allow other origins to make SockJS connections.

Child Elements of <websocket-message-broker>

- [expression-handler](#)
- [intercept-message](#)

<intercept-message>

Defines an authorization rule for a message.

Parent Elements of <intercept-message>

- [websocket-message-broker](#)

<intercept-message> Attributes

- **pattern** An ant based pattern that matches on the Message destination. For example, "/" **matches any Message with a destination**; "/admin/" matches any Message that has a destination that starts with "/admin/**".
- **type** The type of message to match on. Valid values are defined in `SimpMessageType` (i.e. `CONNECT`, `CONNECT_ACK`, `HEARTBEAT`, `MESSAGE`, `SUBSCRIBE`, `UNSUBSCRIBE`, `DISCONNECT`, `DISCONNECT_ACK`, `OTHER`).
- **access** The expression used to secure the Message. For example, "denyAll" will deny access to all of the matching Messages; "permitAll" will grant access to all of the matching Messages; "hasRole('ADMIN')" requires the current user to have the role 'ROLE_ADMIN' for the matching Messages.

Authentication Services

Before Spring Security 3.0, an `AuthenticationManager` was automatically registered internally. Now you must register one explicitly using the `<authentication-manager>` element. This creates an instance of Spring Security's `ProviderManager` class, which needs to be configured with a list of one or more `AuthenticationProvider` instances. These can either be created using syntax elements provided by the namespace, or they can be standard bean definitions, marked for addition to the list using the `authentication-provider` element.

<authentication-manager>

Every Spring Security application which uses the namespace must have include this element somewhere. It is responsible for registering the `AuthenticationManager` which provides authentication services to the application. All elements which create `AuthenticationProvider` instances should be children of this element.

<authentication-manager> Attributes

- **alias** This attribute allows you to define an alias name for the internal instance for use in your own configuration. Its use is described in the [namespace introduction](#).
- **erase-credentials** If set to true, the `AuthenticationManager` will attempt to clear any credentials data in the returned `Authentication` object, once the user has been authenticated. Literally it maps to the `eraseCredentialsAfterAuthentication` property of the `ProviderManager`. This is discussed in the [Core Services](#) chapter.

- **id** This attribute allows you to define an id for the internal instance for use in your own configuration. It is the same as the alias element, but provides a more consistent experience with elements that use the id attribute.

Child Elements of <authentication-manager>

- [authentication-provider](#)
- [ldap-authentication-provider](#)

<authentication-provider>

Unless used with a `ref` attribute, this element is shorthand for configuring a [DaoAuthenticationProvider](#). `DaoAuthenticationProvider` loads user information from a `UserDetailsService` and compares the username/password combination with the values supplied at login. The `UserDetailsService` instance can be defined either by using an available namespace element (`jdbc-user-service` or by using the `user-service-ref` attribute to point to a bean defined elsewhere in the application context). You can find examples of these variations in the [namespace introduction](#).

Parent Elements of <authentication-provider>

- [authentication-manager](#)

<authentication-provider> Attributes

- **ref** Defines a reference to a Spring bean that implements `AuthenticationProvider`.

If you have written your own `AuthenticationProvider` implementation (or want to configure one of Spring Security's own implementations as a traditional bean for some reason, then you can use the following syntax to add it to the internal list of `ProviderManager`:

```
<security:authentication-manager>
<security:authentication-provider ref="myAuthenticationProvider" />
</security:authentication-manager>
<bean id="myAuthenticationProvider" class="com.something.MyAuthenticationProvider"/>
```

- **user-service-ref** A reference to a bean that implements `UserDetailsService` that may be created using the standard bean element or the custom `user-service` element.

Child Elements of <authentication-provider>

- [jdbc-user-service](#)
- [ldap-user-service](#)
- [password-encoder](#)
- [user-service](#)

<jdbc-user-service>

Causes creation of a JDBC-based `UserDetailsService`.

<jdbc-user-service> Attributes

- **authorities-by-username-query** An SQL statement to query for a user's granted authorities given a username.

The default is

```
select username, authority from authorities where username = ?
```

- **cache-ref** Defines a reference to a cache for use with a `UserDetailsService`.
- **data-source-ref** The bean ID of the `DataSource` which provides the required tables.
- **group-authorities-by-username-query** An SQL statement to query user's group authorities given a username. The default is

```
select
  g.id, g.group_name, ga.authority
from
  groups g, group_members gm, group_authorities ga
where
  gm.username = ? and g.id = ga.group_id and g.id = gm.group_id
```

- **id** A bean identifier, used for referring to the bean elsewhere in the context.
- **role-prefix** A non-empty string prefix that will be added to role strings loaded from persistent storage (default is "ROLE_"). Use the value "none" for no prefix in cases where the default is non-empty.
- **users-by-username-query** An SQL statement to query a username, password, and enabled status given a username. The default is

```
select username, password, enabled from users where username = ?
```

<password-encoder>

Authentication providers can optionally be configured to use a password encoder as described in the [namespace introduction](#). This will result in the bean being injected with the appropriate `PasswordEncoder` instance.

Parent Elements of <password-encoder>

- [authentication-provider](#)
- [password-compare](#)

<password-encoder> Attributes

- **hash** Defines the hashing algorithm used on user passwords. We recommend strongly against using MD4, as it is a very weak hashing algorithm.
- **ref** Defines a reference to a Spring bean that implements `PasswordEncoder`.

<user-service>

Creates an in-memory `UserDetailsService` from a properties file or a list of "user" child elements. Usernames are converted to lower-case internally to allow for case-insensitive lookups, so this should not be used if case-sensitivity is required.

<user-service> Attributes

- **id** A bean identifier, used for referring to the bean elsewhere in the context.
- **properties** The location of a Properties file where each line is in the format of

```
username=password,grantedAuthority[,grantedAuthority][,enabled|disabled]
```

Child Elements of <user-service>

- [user](#)

<user>

Represents a user in the application.

Parent Elements of <user>

- [user-service](#)

<user> Attributes

- **authorities** One or more authorities granted to the user. Separate authorities with a comma (but no space). For example, "ROLE_USER,ROLE_ADMINISTRATOR"
- **disabled** Can be set to "true" to mark an account as disabled and unusable.
- **locked** Can be set to "true" to mark an account as locked and unusable.
- **name** The username assigned to the user.
- **password** The password assigned to the user. This may be hashed if the corresponding authentication provider supports hashing (remember to set the "hash" attribute of the "user-service" element). This attribute be omitted in the case where the data will not be used for authentication, but only for accessing authorities. If omitted, the namespace will generate a random value, preventing its accidental use for authentication. Cannot be empty.

Method Security

<global-method-security>

This element is the primary means of adding support for securing methods on Spring Security beans. Methods can be secured by the use of annotations (defined at the interface or class level) or by defining a set of pointcuts as child elements, using AspectJ syntax.

<global-method-security> Attributes

- **access-decision-manager-ref** Method security uses the same `AccessDecisionManager` configuration as web security, but this can be overridden using this attribute. By default an `AffirmativeBased` implementation is used for with a `RoleVoter` and an `AuthenticatedVoter`.
- **authentication-manager-ref** A reference to an `AuthenticationManager` that should be used for method security.
- **jsr250-annotations** Specifies whether JSR-250 style attributes are to be used (for example "RolesAllowed"). This will require the `javax.annotation.security` classes on the classpath. Setting this to true also adds a `Jsr250Voter` to the `AccessDecisionManager`, so you need to make sure you do this if you are using a custom implementation and want to use these annotations.
- **metadata-source-ref** An external `MethodSecurityMetadataSource` instance can be supplied which will take priority over other sources (such as the default annotations).
- **mode** This attribute can be set to "aspectj" to specify that AspectJ should be used instead of the default Spring AOP. Secured methods must be woven with the `AnnotationSecurityAspect` from the `spring-security-aspects` module.

It is important to note that AspectJ follows Java's rule that annotations on interfaces are not inherited. This means that methods that define the Security annotations on the interface will not be secured. Instead, you must place the Security annotation on the class when using AspectJ.

- **order** Allows the advice "order" to be set for the method security interceptor.
- **pre-post-annotations** Specifies whether the use of Spring Security's pre and post invocation annotations (`@PreFilter`, `@PreAuthorize`, `@PostFilter`, `@PostAuthorize`) should be enabled for this application context. Defaults to "disabled".
- **proxy-target-class** If true, class based proxying will be used instead of interface based proxying.
- **run-as-manager-ref** A reference to an optional `RunAsManager` implementation which will be used by the configured `MethodSecurityInterceptor`
- **secured-annotations** Specifies whether the use of Spring Security's `@Secured` annotations should be enabled for this application context. Defaults to "disabled".

Child Elements of `<global-method-security>`

- [after-invocation-provider](#)
- [expression-handler](#)
- [pre-post-annotation-handling](#)
- [protect-pointcut](#)

`<after-invocation-provider>`

This element can be used to decorate an `AfterInvocationProvider` for use by the security interceptor maintained by the `<global-method-security>` namespace. You can define zero or more of these within the `global-method-security` element, each with a `ref` attribute pointing to an `AfterInvocationProvider` bean instance within your application context.

Parent Elements of `<after-invocation-provider>`

- [global-method-security](#)

`<after-invocation-provider>` Attributes

- **ref** Defines a reference to a Spring bean that implements `AfterInvocationProvider`.

`<pre-post-annotation-handling>`

Allows the default expression-based mechanism for handling Spring Security's pre and post invocation annotations (`@PreFilter`, `@PreAuthorize`, `@PostFilter`, `@PostAuthorize`) to be replaced entirely. Only applies if these annotations are enabled.

Parent Elements of `<pre-post-annotation-handling>`

- [global-method-security](#)

Child Elements of `<pre-post-annotation-handling>`

- [invocation-attribute-factory](#)

- [post-invocation-advice](#)
- [pre-invocation-advice](#)

<invocation-attribute-factory>

Defines the `PrePostInvocationAttributeFactory` instance which is used to generate pre and post invocation metadata from the annotated methods.

Parent Elements of <invocation-attribute-factory>

- [pre-post-annotation-handling](#)

<invocation-attribute-factory> Attributes

- **ref** Defines a reference to a Spring bean Id.

<post-invocation-advice>

Customizes the `PostInvocationAdviceProvider` with the `ref` as the `PostInvocationAuthorizationAdvice` for the <pre-post-annotation-handling> element.

Parent Elements of <post-invocation-advice>

- [pre-post-annotation-handling](#)

<post-invocation-advice> Attributes

- **ref** Defines a reference to a Spring bean Id.

<pre-invocation-advice>

Customizes the `PreInvocationAuthorizationAdviceVoter` with the `ref` as the `PreInvocationAuthorizationAdviceVoter` for the <pre-post-annotation-handling> element.

Parent Elements of <pre-invocation-advice>

- [pre-post-annotation-handling](#)

<pre-invocation-advice> Attributes

- **ref** Defines a reference to a Spring bean Id.

Securing Methods using

<protect-pointcut> Rather than defining security attributes on an individual method or class basis using the `@Secured` annotation, you can define cross-cutting security constraints across whole sets of methods and interfaces in your service layer using the <protect-pointcut> element. You can find an example in the [namespace introduction](#).

Parent Elements of <protect-pointcut>

- [global-method-security](#)

<protect-pointcut> Attributes

- **access** Access configuration attributes list that applies to all methods matching the pointcut, e.g. "ROLE_A,ROLE_B"

- **expression** An AspectJ expression, including the 'execution' keyword. For example, 'execution(int com.foo.TargetObject.countLength(String))' (without the quotes).

<intercept-methods>

Can be used inside a bean definition to add a security interceptor to the bean and set up access configuration attributes for the bean's methods

<intercept-methods> Attributes

- **access-decision-manager-ref** Optional AccessDecisionManager bean ID to be used by the created method security interceptor.

Child Elements of <intercept-methods>

- [protect](#)

<method-security-metadata-source>

Creates a MethodSecurityMetadataSource instance

<method-security-metadata-source> Attributes

- **id** A bean identifier, used for referring to the bean elsewhere in the context.
- **use-expressions** Enables the use of expressions in the 'access' attributes in <intercept-url> elements rather than the traditional list of configuration attributes. Defaults to 'false'. If enabled, each attribute should contain a single Boolean expression. If the expression evaluates to 'true', access will be granted.

Child Elements of <method-security-metadata-source>

- [protect](#)

<protect>

Defines a protected method and the access control configuration attributes that apply to it. We strongly advise you NOT to mix "protect" declarations with any services provided "global-method-security".

Parent Elements of <protect>

- [intercept-methods](#)
- [method-security-metadata-source](#)

<protect> Attributes

- **access** Access configuration attributes list that applies to the method, e.g. "ROLE_A,ROLE_B".
- **method** A method name

LDAP Namespace Options

LDAP is covered in some details in [its own chapter](#). We will expand on that here with some explanation of how the namespace options map to Spring beans. The LDAP implementation uses Spring LDAP extensively, so some familiarity with that project's API may be useful.

Defining the LDAP Server using the

`<ldap-server>` Element This element sets up a Spring LDAP `ContextSource` for use by the other LDAP beans, defining the location of the LDAP server and other information (such as a username and password, if it doesn't allow anonymous access) for connecting to it. It can also be used to create an embedded server for testing. Details of the syntax for both options are covered in the [LDAP chapter](#). The actual `ContextSource` implementation is `DefaultSpringSecurityContextSource` which extends Spring LDAP's `LdapContextSource` class. The `manager-dn` and `manager-password` attributes map to the latter's `userDn` and `password` properties respectively.

If you only have one server defined in your application context, the other LDAP namespace-defined beans will use it automatically. Otherwise, you can give the element an "id" attribute and refer to it from other namespace beans using the `server-ref` attribute. This is actually the bean id of the `ContextSource` instance, if you want to use it in other traditional Spring beans.

`<ldap-server>` Attributes

- **mode** Explicitly specifies which embedded ldap server should use. Values are `apacheds` and `unboundid`. By default, it will depends if the library is available in the classpath.
- **id** A bean identifier, used for referring to the bean elsewhere in the context.
- **ldif** Explicitly specifies an ldif file resource to load into an embedded LDAP server. The ldif is should be a Spring resource pattern (i.e. `classpath:init.ldif`). The default is `classpath*:*.ldif`
- **manager-dn** Username (DN) of the "manager" user identity which will be used to authenticate to a (non-embedded) LDAP server. If omitted, anonymous access will be used.
- **manager-password** The password for the manager DN. This is required if the `manager-dn` is specified.
- **port** Specifies an IP port number. Used to configure an embedded LDAP server, for example. The default value is 33389.
- **root** Optional root suffix for the embedded LDAP server. Default is `"dc=springframework,dc=org"`
- **url** Specifies the ldap server URL when not using the embedded LDAP server.

`<ldap-authentication-provider>`

This element is shorthand for the creation of an `LdapAuthenticationProvider` instance. By default this will be configured with a `BindAuthenticator` instance and a `DefaultAuthoritiesPopulator`. As with all namespace authentication providers, it must be included as a child of the `authentication-provider` element.

Parent Elements of `<ldap-authentication-provider>`

- [authentication-manager](#)

`<ldap-authentication-provider>` Attributes

- **group-role-attribute** The LDAP attribute name which contains the role name which will be used within Spring Security. Maps to the `DefaultLdapAuthoritiesPopulator`'s `groupRoleAttribute` property. Defaults to `"cn"`.

- **group-search-base** Search base for group membership searches. Maps to the `DefaultLdapAuthoritiesPopulator`'s `groupSearchBase` constructor argument. Defaults to "" (searching from the root).
- **group-search-filter** Group search filter. Maps to the `DefaultLdapAuthoritiesPopulator`'s `groupSearchFilter` property. Defaults to `(uniqueMember={0})`. The substituted parameter is the DN of the user.
- **role-prefix** A non-empty string prefix that will be added to role strings loaded from persistent. Maps to the `DefaultLdapAuthoritiesPopulator`'s `rolePrefix` property. Defaults to "ROLE_". Use the value "none" for no prefix in cases where the default is non-empty.
- **server-ref** The optional server to use. If omitted, and a default LDAP server is registered (using `<ldap-server>` with no `ld`), that server will be used.
- **user-context-mapper-ref** Allows explicit customization of the loaded user object by specifying a `UserDetailsContextMapper` bean which will be called with the context information from the user's directory entry
- **user-details-class** Allows the `objectClass` of the user entry to be specified. If set, the framework will attempt to load standard attributes for the defined class into the returned `UserDetails` object
- **user-dn-pattern** If your users are at a fixed location in the directory (i.e. you can work out the DN directly from the username without doing a directory search), you can use this attribute to map directly to the DN. It maps directly to the `userDnPatterns` property of `AbstractLdapAuthenticator`. The value is a specific pattern used to build the user's DN, for example `uid={0},ou=people`. The key `{0}` must be present and will be substituted with the username.
- **user-search-base** Search base for user searches. Defaults to "". Only used with a 'user-search-filter'.

If you need to perform a search to locate the user in the directory, then you can set these attributes to control the search. The `BindAuthenticator` will be configured with a `FilterBasedLdapUserSearch` and the attribute values map directly to the first two arguments of that bean's constructor. If these attributes aren't set and no `user-dn-pattern` has been supplied as an alternative, then the default search values of `user-search-filter="(uid={0})"` and `user-search-base=""` will be used.

- **user-search-filter** The LDAP filter used to search for users (optional). For example `"(uid={0})"`. The substituted parameter is the user's login name.

If you need to perform a search to locate the user in the directory, then you can set these attributes to control the search. The `BindAuthenticator` will be configured with a `FilterBasedLdapUserSearch` and the attribute values map directly to the first two arguments of that bean's constructor. If these attributes aren't set and no `user-dn-pattern` has been supplied as an alternative, then the default search values of `user-search-filter="(uid={0})"` and `user-search-base=""` will be used.

Child Elements of `<ldap-authentication-provider>`

- [password-compare](#)

`<password-compare>`

This is used as child element to `<ldap-provider>` and switches the authentication strategy from `BindAuthenticator` to `PasswordComparisonAuthenticator`.

Parent Elements of <password-compare>

- [ldap-authentication-provider](#)

<password-compare> Attributes

- **hash** Defines the hashing algorithm used on user passwords. We recommend strongly against using MD4, as it is a very weak hashing algorithm.
- **password-attribute** The attribute in the directory which contains the user password. Defaults to "userPassword".

Child Elements of <password-compare>

- [password-encoder](#)

<ldap-user-service>

This element configures an LDAP `UserDetailsService`. The class used is `LdapUserDetailsService` which is a combination of a `FilterBasedLdapUserSearch` and a `DefaultLdapAuthoritiesPopulator`. The attributes it supports have the same usage as in `<ldap-provider>`.

<ldap-user-service> Attributes

- **cache-ref** Defines a reference to a cache for use with a `UserDetailsService`.
- **group-role-attribute** The LDAP attribute name which contains the role name which will be used within Spring Security. Defaults to "cn".
- **group-search-base** Search base for group membership searches. Defaults to "" (searching from the root).
- **group-search-filter** Group search filter. Defaults to `(uniqueMember={0})`. The substituted parameter is the DN of the user.
- **id** A bean identifier, used for referring to the bean elsewhere in the context.
- **role-prefix** A non-empty string prefix that will be added to role strings loaded from persistent storage (e.g. "ROLE_"). Use the value "none" for no prefix in cases where the default is non-empty.
- **server-ref** The optional server to use. If omitted, and a default LDAP server is registered (using `<ldap-server>` with no `id`), that server will be used.
- **user-context-mapper-ref** Allows explicit customization of the loaded user object by specifying a `UserDetailsContextMapper` bean which will be called with the context information from the user's directory entry
- **user-details-class** Allows the `objectClass` of the user entry to be specified. If set, the framework will attempt to load standard attributes for the defined class into the returned `UserDetails` object
- **user-search-base** Search base for user searches. Defaults to "". Only used with a 'user-search-filter'.
- **user-search-filter** The LDAP filter used to search for users (optional). For example `(uid={0})`. The substituted parameter is the user's login name.

20.3 Spring Security Dependencies

This appendix provides a reference of the modules in Spring Security and the additional dependencies that they require in order to function in a running application. We don't include dependencies that are only used when building or testing Spring Security itself. Nor do we include transitive dependencies which are required by external dependencies.

The version of Spring required is listed on the project website, so the specific versions are omitted for Spring dependencies below. Note that some of the dependencies listed as "optional" below may still be required for other non-security functionality in a Spring application. Also dependencies listed as "optional" may not actually be marked as such in the project's Maven POM files if they are used in most applications. They are "optional" only in the sense that you don't need them unless you are using the specified functionality.

Where a module depends on another Spring Security module, the non-optional dependencies of the module it depends on are also assumed to be required and are not listed separately.

spring-security-core

The core module must be included in any project using Spring Security.

Table 20.1. Core Dependencies

Dependency	Version	Description
ehcache	1.6.2	Required if the Ehcache-based user cache implementation is used (optional).
spring-aop		Method security is based on Spring AOP
spring-beans		Required for Spring configuration
spring-expression		Required for expression-based method security (optional)
spring-jdbc		Required if using a database to store user data (optional).
spring-tx		Required if using a database to store user data (optional).
aspectjrt	1.6.10	Required if using AspectJ support (optional).
jsr250-api	1.0	Required if you are using JSR-250 method-security annotations (optional).

spring-security-remoting

This module is typically required in web applications which use the Servlet API.

Table 20.2. Remoting Dependencies

Dependency	Version	Description
spring-security-core		
spring-web		Required for clients which use HTTP remoting support.

spring-security-web

This module is typically required in web applications which use the Servlet API.

Table 20.3. Web Dependencies

Dependency	Version	Description
spring-security-core		
spring-web		Spring web support classes are used extensively.
spring-jdbc		Required for JDBC-based persistent remember-me token repository (optional).
spring-tx		Required by remember-me persistent token repository implementations (optional).

spring-security-ldap

This module is only required if you are using LDAP authentication.

Table 20.4. LDAP Dependencies

Dependency	Version	Description
spring-security-core		
spring-ldap-core	1.3.0	LDAP support is based on Spring LDAP.
spring-tx		Data exception classes are required.
apache-ds ¹	1.5.5	Required if you are using an embedded LDAP server (optional).
shared-ldap	0.9.15	Required if you are using an embedded LDAP server (optional).
ldapsdk	4.1	Mozilla LdapSDK. Used for decoding LDAP password

Dependency	Version	Description
		policy controls if you are using password-policy functionality with OpenLDAP, for example.

¹The modules `apacheds-core`, `apacheds-core-entry`, `apacheds-protocol-shared`, `apacheds-protocol-ldap` and `apacheds-server-jndi` are required.

spring-security-config

This module is required if you are using Spring Security namespace configuration.

Table 20.5. Config Dependencies

Dependency	Version	Description
spring-security-core		
spring-security-web		Required if you are using any web-related namespace configuration (optional).
spring-security-ldap		Required if you are using the LDAP namespace options (optional).
spring-security-openid		Required if you are using OpenID authentication (optional).
aspectjweaver	1.6.10	Required if using the protect-pointcut namespace syntax (optional).

spring-security-acl

The ACL module.

Table 20.6. ACL Dependencies

Dependency	Version	Description
spring-security-core		
ehcache	1.6.2	Required if the Ehcache-based ACL cache implementation is used (optional if you are using your own implementation).
spring-jdbc		Required if you are using the default JDBC-based <code>AclService</code> (optional if you implement your own).
spring-tx		Required if you are using the default JDBC-based <code>AclService</code>

Dependency	Version	Description
		(optional if you implement your own).

spring-security-cas

The CAS module provides integration with JA-SIG CAS.

Table 20.7. CAS Dependencies

Dependency	Version	Description
spring-security-core		
spring-security-web		
cas-client-core	3.1.12	The JA-SIG CAS Client. This is the basis of the Spring Security integration.
ehcache	1.6.2	Required if you are using the Ehcache-based ticket cache (optional).

spring-security-openid

The OpenID module.

Table 20.8. OpenID Dependencies

Dependency	Version	Description
spring-security-core		
spring-security-web		
openid4java-nodeps	0.9.6	Spring Security's OpenID integration uses OpenID4Java.
httpClient	4.1.1	openid4java-nodeps depends on HttpClient 4.
guice	2.0	openid4java-nodeps depends on Guice 2.

spring-security-taglibs

Provides Spring Security's JSP tag implementations.

Table 20.9. Taglib Dependencies

Dependency	Version	Description
spring-security-core		

Dependency	Version	Description
spring-security-web		
spring-security-acl		Required if you are using the <code>accesscontrollist</code> tag or <code>hasPermission()</code> expressions with ACLs (optional).
spring-expression		Required if you are using SPEL expressions in your tag access constraints.

20.4 Proxy Server Configuration

When using a proxy server it is important to ensure that you have configured your application properly. For example, many applications will have a load balancer that responds to request for <https://example.com/> by forwarding the request to an application server at <https://192.168.1:8080> Without proper configuration, the application server will not know that the load balancer exists and treat the request as though <https://192.168.1:8080> was requested by the client.

To fix this you can use [RFC 7239](#) to specify that a load balancer is being used. To make the application aware of this, you need to either configure your application server aware of the X-Forwarded headers. For example Tomcat uses the [RemotepValve](#) and Jetty uses [ForwardedRequestCustomizer](#). Alternatively, Spring 4.3+ users can leverage [ForwardedHeaderFilter](#).

Spring Boot users may use the `server.use-forward-headers` property to configure the application. See the [Spring Boot documentation](#) for further details.

20.5 Spring Security FAQ

- the section called “General Questions”
- the section called “Common Problems”
- the section called “Spring Security Architecture Questions”
- the section called “Common "Howto" Requests”

General Questions

1. the section called “Will Spring Security take care of all my application security requirements?”
2. the section called “Why not just use web.xml security?”
3. the section called “What Java and Spring Framework versions are required?”
4. the section called “I’m new to Spring Security and I need to build an application that supports CAS single sign-on over HTTPS, while allowing Basic authentication locally for certain URLs, authenticating against multiple back end user information sources (LDAP and JDBC). I’ve copied some configuration files I found but it doesn’t work.”

Will Spring Security take care of all my application security requirements?

Spring Security provides you with a very flexible framework for your authentication and authorization requirements, but there are many other considerations for building a secure application that are outside its scope. Web applications are vulnerable to all kinds of attacks which you should be familiar with, preferably before you start development so you can design and code with them in mind from the beginning. Check out the <http://www.owasp.org/> [OWASP web site] for information on the major issues facing web application developers and the countermeasures you can use against them.

Why not just use web.xml security?

Let's assume you're developing an enterprise application based on Spring. There are four security concerns you typically need to address: authentication, web request security, service layer security (i.e. your methods that implement business logic), and domain object instance security (i.e. different domain objects have different permissions). With these typical requirements in mind:

1. *Authentication*: The servlet specification provides an approach to authentication. However, you will need to configure the container to perform authentication which typically requires editing of container-specific "realm" settings. This makes a non-portable configuration, and if you need to write an actual Java class to implement the container's authentication interface, it becomes even more non-portable. With Spring Security you achieve complete portability - right down to the WAR level. Also, Spring Security offers a choice of production-proven authentication providers and mechanisms, meaning you can switch your authentication approaches at deployment time. This is particularly valuable for software vendors writing products that need to work in an unknown target environment.
2. *Web request security*: The servlet specification provides an approach to secure your request URIs. However, these URIs can only be expressed in the servlet specification's own limited URI path format. Spring Security provides a far more comprehensive approach. For instance, you can use Ant paths or regular expressions, you can consider parts of the URI other than simply the requested page (e.g. you can consider HTTP GET parameters) and you can implement your own runtime source of configuration data. This means your web request security can be dynamically changed during the actual execution of your webapp.
3. *Service layer and domain object security*: The absence of support in the servlet specification for services layer security or domain object instance security represent serious limitations for multi-tiered applications. Typically developers either ignore these requirements, or implement security logic within their MVC controller code (or even worse, inside the views). There are serious disadvantages with this approach:
 - a. *Separation of concerns*: Authorization is a crosscutting concern and should be implemented as such. MVC controllers or views implementing authorization code makes it more difficult to test both the controller and authorization logic, more difficult to debug, and will often lead to code duplication.
 - b. *Support for rich clients and web services*: If an additional client type must ultimately be supported, any authorization code embedded within the web layer is non-reusable. It should be considered that Spring remoting exporters only export service layer beans (not MVC controllers). As such authorization logic needs to be located in the services layer to support a multitude of client types.
 - c. *Layering issues*: An MVC controller or view is simply the incorrect architectural layer to implement authorization decisions concerning services layer methods or domain object instances. Whilst the Principal may be passed to the services layer to enable it to make the authorization decision, doing so would introduce an additional argument on every services layer method. A more elegant approach is to use a ThreadLocal to hold the Principal, although this would likely increase

development time to a point where it would become more economical (on a cost-benefit basis) to simply use a dedicated security framework.

- d. *Authorisation code quality*: It is often said of web frameworks that they "make it easier to do the right things, and harder to do the wrong things". Security frameworks are the same, because they are designed in an abstract manner for a wide range of purposes. Writing your own authorization code from scratch does not provide the "design check" a framework would offer, and in-house authorization code will typically lack the improvements that emerge from widespread deployment, peer review and new versions.

For simple applications, servlet specification security may just be enough. Although when considered within the context of web container portability, configuration requirements, limited web request security flexibility, and non-existent services layer and domain object instance security, it becomes clear why developers often look to alternative solutions.

What Java and Spring Framework versions are required?

Spring Security 3.0 and 3.1 require at least JDK 1.5 and also require Spring 3.0.3 as a minimum. Ideally you should be using the latest release versions to avoid problems.

Spring Security 2.0.x requires a minimum JDK version of 1.4 and is built against Spring 2.0.x. It should also be compatible with applications using Spring 2.5.x.

I'm new to Spring Security and I need to build an application that supports CAS single sign-on over HTTPS, while allowing Basic authentication locally for certain URLs, authenticating against multiple back end user information sources (LDAP and JDBC). I've copied some configuration files I found but it doesn't work.

What could be wrong?

Or substitute an alternative complex scenario...

Realistically, you need an understanding of the technologies you are intending to use before you can successfully build applications with them. Security is complicated. Setting up a simple configuration using a login form and some hard-coded users using Spring Security's namespace is reasonably straightforward. Moving to using a backed JDBC database is also easy enough. But if you try and jump straight to a complicated deployment scenario like this you will almost certainly be frustrated. There is a big jump in the learning curve required to set up systems like CAS, configure LDAP servers and install SSL certificates properly. So you need to take things one step at a time.

From a Spring Security perspective, the first thing you should do is follow the "Getting Started" guide on the web site. This will take you through a series of steps to get up and running and get some idea of how the framework operates. If you are using other technologies which you aren't familiar with then you should do some research and try to make sure you can use them in isolation before combining them in a complex system.

Common Problems

1. Authentication

- a. the section called "When I try to log in, I get an error message that says "Bad Credentials". What's wrong?"
- b. the section called "My application goes into an "endless loop" when I try to login, what's going on?"

- c. the section called "I get an exception with the message "Access is denied (user is anonymous);". What's wrong?"
- d. the section called "Why can I still see a secured page even after I've logged out of my application?"
- e. the section called "I get an exception with the message "An Authentication object was not found in the SecurityContext". What's wrong?"
- f. the section called "I can't get LDAP authentication to work."

2. Session Management

- a. the section called "I'm using Spring Security's concurrent session control to prevent users from logging in more than once at a time."
- b. the section called "Why does the session Id change when I authenticate through Spring Security?"
- c. the section called "I'm using Tomcat (or some other servlet container) and have enabled HTTPS for my login page, switching back to HTTP afterwards."
- d. the section called "I'm trying to use the concurrent session-control support but it won't let me log back in, even if I'm sure I've logged out and haven't exceeded the allowed sessions."
- e. the section called "Spring Security is creating a session somewhere, even though I've configured it not to, by setting the create-session attribute to never."

3. Miscellaneous

- a. the section called "I get a 403 Forbidden when performing a POST"
- b. the section called "I'm forwarding a request to another URL using the RequestDispatcher, but my security constraints aren't being applied."
- c. the section called "I have added Spring Security's <global-method-security> element to my application context but if I add security annotations to my Spring MVC controller beans (Struts actions etc.) then they don't seem to have an effect."
- d. the section called "I have a user who has definitely been authenticated, but when I try to access the SecurityContextHolder during some requests, the Authentication is null."
- e. the section called "The authorize JSP Tag doesn't respect my method security annotations when using the URL attribute."

When I try to log in, I get an error message that says "Bad Credentials". What's wrong?

This means that authentication has failed. It doesn't say why, as it is good practice to avoid giving details which might help an attacker guess account names or passwords.

This also means that if you ask this question in the forum, you will not get an answer unless you provide additional information. As with any issue you should check the output from the debug log, note any exception stacktraces and related messages. Step through the code in a debugger to see where the authentication fails and why. Write a test case which exercises your authentication configuration outside of the application. More often than not, the failure is due to a difference in the password data stored in a database and that entered by the user. If you are using hashed passwords, make sure the value stored

in your database is *exactly* the same as the value produced by the `PasswordEncoder` configured in your application.

My application goes into an "endless loop" when I try to login, what's going on?

A common user problem with infinite loop and redirecting to the login page is caused by accidentally configuring the login page as a "secured" resource. Make sure your configuration allows anonymous access to the login page, either by excluding it from the security filter chain or marking it as requiring `ROLE_ANONYMOUS`.

If your `AccessDecisionManager` includes an `AuthenticatedVoter`, you can use the attribute `"IS_AUTHENTICATED_ANONYMOUSLY"`. This is automatically available if you are using the standard namespace configuration setup.

From Spring Security 2.0.1 onwards, when you are using namespace-based configuration, a check will be made on loading the application context and a warning message logged if your login page appears to be protected.

I get an exception with the message "Access is denied (user is anonymous);". What's wrong?

This is a debug level message which occurs the first time an anonymous user attempts to access a protected resource.

```
DEBUG [ExceptionTranslationFilter] - Access is denied (user is anonymous); redirecting to authentication
entry point
org.springframework.security.AccessDeniedException: Access is denied
at org.springframework.security.vote.AffirmativeBased.decide(AffirmativeBased.java:68)
at
org.springframework.security.intercept.AbstractSecurityInterceptor.beforeInvocation(AbstractSecurityInterceptor.java:262)
```

It is normal and shouldn't be anything to worry about.

Why can I still see a secured page even after I've logged out of my application?

The most common reason for this is that your browser has cached the page and you are seeing a copy which is being retrieved from the browsers cache. Verify this by checking whether the browser is actually sending the request (check your server access logs, the debug log or use a suitable browser debugging plugin such as "Tamper Data" for Firefox). This has nothing to do with Spring Security and you should configure your application or server to set the appropriate `Cache-Control` response headers. Note that SSL requests are never cached.

I get an exception with the message "An Authentication object was not found in the SecurityContext". What's wrong?

This is a another debug level message which occurs the first time an anonymous user attempts to access a protected resource, but when you do not have an `AnonymousAuthenticationFilter` in your filter chain configuration.

```
DEBUG [ExceptionTranslationFilter] - Authentication exception occurred; redirecting to authentication
entry point
org.springframework.security.AuthenticationCredentialsNotFoundException:
An Authentication object was not found in the SecurityContext
at
org.springframework.security.intercept.AbstractSecurityInterceptor.credentialsNotFound(AbstractSecurityInterceptor.java:34)
at
org.springframework.security.intercept.AbstractSecurityInterceptor.beforeInvocation(AbstractSecurityInterceptor.java:254)
```

It is normal and shouldn't be anything to worry about.

I can't get LDAP authentication to work.

What's wrong with my configuration?

Note that the permissions for an LDAP directory often do not allow you to read the password for a user. Hence it is often not possible to use the the section called "What is a UserDetailsService and do I need one?" where Spring Security compares the stored password with the one submitted by the user. The most common approach is to use LDAP "bind", which is one of the operations supported by [the LDAP protocol](#). With this approach, Spring Security validates the password by attempting to authenticate to the directory as the user.

The most common problem with LDAP authentication is a lack of knowledge of the directory server tree structure and configuration. This will be different in different companies, so you have to find it out yourself. Before adding a Spring Security LDAP configuration to an application, it's a good idea to write a simple test using standard Java LDAP code (without Spring Security involved), and make sure you can get that to work first. For example, to authenticate a user, you could use the following code:

```
@Test
public void ldapAuthenticationIsSuccessful() throws Exception {
    Hashtable<String,String> env = new Hashtable<String,String>();
    env.put(Context.SECURITY_AUTHENTICATION, "simple");
    env.put(Context.SECURITY_PRINCIPAL, "cn=joe,ou=users,dc=mycompany,dc=com");
    env.put(Context.PROVIDER_URL, "ldap://mycompany.com:389/dc=mycompany,dc=com");
    env.put(Context.SECURITY_CREDENTIALS, "joespassword");
    env.put(Context.INITIAL_CONTEXT_FACTORY, "com.sun.jndi.ldap.LdapCtxFactory");

    InitialLdapContext ctx = new InitialLdapContext(env, null);
}
```

Session Management

Session management issues are a common source of forum questions. If you are developing Java web applications, you should understand how the session is maintained between the servlet container and the user's browser. You should also understand the difference between secure and non-secure cookies and the implications of using HTTP/HTTPS and switching between the two. Spring Security has nothing to do with maintaining the session or providing session identifiers. This is entirely handled by the servlet container.

I'm using Spring Security's concurrent session control to prevent users from logging in more than once at a time.

When I open another browser window after logging in, it doesn't stop me from logging in again. Why can I log in more than once?

Browsers generally maintain a single session per browser instance. You cannot have two separate sessions at once. So if you log in again in another window or tab you are just reauthenticating in the same session. The server doesn't know anything about tabs, windows or browser instances. All it sees are HTTP requests and it ties those to a particular session according to the value of the JSESSIONID cookie that they contain. When a user authenticates during a session, Spring Security's concurrent session control checks the number of *other authenticated sessions* that they have. If they are already authenticated with the same session, then re-authenticating will have no effect.

Why does the session Id change when I authenticate through Spring Security?

With the default configuration, Spring Security changes the session ID when the user authenticates. If you're using a Servlet 3.1 or newer container, the session ID is simply changed. If you're using an

older container, Spring Security invalidates the existing session, creates a new session, and transfers the session data to the new session. Changing the session identifier in this manner prevents "session-fixation" attacks. You can find more about this online and in the reference manual.

I'm using Tomcat (or some other servlet container) and have enabled HTTPS for my login page, switching back to HTTP afterwards.

It doesn't work - I just end up back at the login page after authenticating.

This happens because sessions created under HTTPS, for which the session cookie is marked as "secure", cannot subsequently be used under HTTP. The browser will not send the cookie back to the server and any session state will be lost (including the security context information). Starting a session in HTTP first should work as the session cookie won't be marked as secure. However, Spring Security's [Session Fixation Protection](#) can interfere with this because it results in a new session ID cookie being sent back to the user's browser, usually with the secure flag. To get around this, you can disable session fixation protection, but in newer Servlet containers you can also configure session cookies to never use the secure flag. Note that switching between HTTP and HTTPS is not a good idea in general, as any application which uses HTTP at all is vulnerable to man-in-the-middle attacks. To be truly secure, the user should begin accessing your site in HTTPS and continue using it until they log out. Even clicking on an HTTPS link from a page accessed over HTTP is potentially risky. If you need more convincing, check out a tool like [sslstrip](#).

I'm not switching between HTTP and HTTPS but my session is still getting lost

Sessions are maintained either by exchanging a session cookie or by adding a `jsessionid` parameter to URLs (this happens automatically if you are using JSTL to output URLs, or if you call `HttpServletResponse.encodeUrl` on URLs (before a redirect, for example). If clients have cookies disabled, and you are not rewriting URLs to include the `jsessionid`, then the session will be lost. Note that the use of cookies is preferred for security reasons, as it does not expose the session information in the URL.

I'm trying to use the concurrent session-control support but it won't let me log back in, even if I'm sure I've logged out and haven't exceeded the allowed sessions.

Make sure you have added the listener to your `web.xml` file. It is essential to make sure that the Spring Security session registry is notified when a session is destroyed. Without it, the session information will not be removed from the registry.

```
<listener>
  <listener-class>org.springframework.security.web.session.HttpSessionEventPublisher</listener-
class>
</listener>
```

Spring Security is creating a session somewhere, even though I've configured it not to, by setting the create-session attribute to never.

This usually means that the user's application is creating a session somewhere, but that they aren't aware of it. The most common culprit is a JSP. Many people aren't aware that JSPs create sessions by default. To prevent a JSP from creating a session, add the directive `<%@ page session="false" %>` to the top of the page.

If you are having trouble working out where a session is being created, you can add some debugging code to track down the location(s). One way to do this would be to add a `javax.servlet.http.HttpSessionListener` to your application, which calls `Thread.dumpStack()` in the `sessionCreated` method.

I get a 403 Forbidden when performing a POST

If an HTTP 403 Forbidden is returned for HTTP POST, but works for HTTP GET then the issue is most likely related to [CSRF](#). Either provide the CSRF Token or disable CSRF protection (not recommended).

I'm forwarding a request to another URL using the RequestDispatcher, but my security constraints aren't being applied.

Filters are not applied by default to forwards or includes. If you really want the security filters to be applied to forwards and/or includes, then you have to configure these explicitly in your `web.xml` using the `<dispatcher>` element, a child element of `<filter-mapping>`.

I have added Spring Security's <global-method-security> element to my application context but if I add security annotations to my Spring MVC controller beans (Struts actions etc.) then they don't seem to have an effect.

In a Spring web application, the application context which holds the Spring MVC beans for the dispatcher servlet is often separate from the main application context. It is often defined in a file called `myapp-servlet.xml`, where "myapp" is the name assigned to the `SpringDispatcherServlet` in `web.xml`. An application can have multiple `DispatcherServlets`, each with its own isolated application context. The beans in these "child" contexts are not visible to the rest of the application. The "parent" application context is loaded by the `ContextLoaderListener` you define in your `web.xml` and is visible to all the child contexts. This parent context is usually where you define your security configuration, including the `<global-method-security>` element). As a result any security constraints applied to methods in these web beans will not be enforced, since the beans cannot be seen from the `DispatcherServlet` context. You need to either move the `<global-method-security>` declaration to the web context or moved the beans you want secured into the main application context.

Generally we would recommend applying method security at the service layer rather than on individual web controllers.

I have a user who has definitely been authenticated, but when I try to access the SecurityContextHolder during some requests, the Authentication is null.

Why can't I see the user information?

If you have excluded the request from the security filter chain using the attribute `filters='none'` in the `<intercept-url>` element that matches the URL pattern, then the `SecurityContextHolder` will not be populated for that request. Check the debug log to see whether the request is passing through the filter chain. (You are reading the debug log, right?).

The authorize JSP Tag doesn't respect my method security annotations when using the URL attribute.

Method security will not hide links when using the `url` attribute in `<sec:authorize>` because we cannot readily reverse engineer what URL is mapped to what controller endpoint as controllers can rely on headers, current user, etc to determine what method to invoke.

Spring Security Architecture Questions

1. the section called "How do I know which package class X is in?"
2. the section called "How do the namespace elements map to conventional bean configurations?"

3. the section called "What does "ROLE_" mean and why do I need it on my role names?"
4. the section called "How do I know which dependencies to add to my application to work with Spring Security?"
5. the section called "What dependencies are needed to run an embedded ApacheDS LDAP server?"
6. the section called "What is a UserDetailsService and do I need one?"

How do I know which package class X is in?

The best way of locating classes is by installing the Spring Security source in your IDE. The distribution includes source jars for each of the modules the project is divided up into. Add these to your project source path and you can navigate directly to Spring Security classes (`Ctrl-Shift-T` in Eclipse). This also makes debugging easier and allows you to troubleshoot exceptions by looking directly at the code where they occur to see what's going on there.

How do the namespace elements map to conventional bean configurations?

There is a general overview of what beans are created by the namespace in the namespace appendix of the reference guide. There is also a detailed blog article called "Behind the Spring Security Namespace" on blog.springsource.com. If you want to know the full details then the code is in the `spring-security-config` module within the Spring Security 3.0 distribution. You should probably read the chapters on namespace parsing in the standard Spring Framework reference documentation first.

What does "ROLE_" mean and why do I need it on my role names?

Spring Security has a voter-based architecture which means that an access decision is made by a series of `AccessDecisionVoters`. The voters act on the "configuration attributes" which are specified for a secured resource (such as a method invocation). With this approach, not all attributes may be relevant to all voters and a voter needs to know when it should ignore an attribute (abstain) and when it should vote to grant or deny access based on the attribute value. The most common voter is the `RoleVoter` which by default votes whenever it finds an attribute with the "ROLE_" prefix. It makes a simple comparison of the attribute (such as "ROLE_USER") with the names of the authorities which the current user has been assigned. If it finds a match (they have an authority called "ROLE_USER"), it votes to grant access, otherwise it votes to deny access.

The prefix can be changed by setting the `rolePrefix` property of `RoleVoter`. If you only need to use roles in your application and have no need for other custom voters, then you can set the prefix to a blank string, in which case the `RoleVoter` will treat all attributes as roles.

How do I know which dependencies to add to my application to work with Spring Security?

It will depend on what features you are using and what type of application you are developing. With Spring Security 3.0, the project jars are divided into clearly distinct areas of functionality, so it is straightforward to work out which Spring Security jars you need from your application requirements. All applications will need the `spring-security-core` jar. If you're developing a web application, you need the `spring-security-web` jar. If you're using security namespace configuration you need the `spring-security-config` jar, for LDAP support you need the `spring-security-ldap` jar and so on.

For third-party jars the situation isn't always quite so obvious. A good starting point is to copy those from one of the pre-built sample applications WEB-INF/lib directories. For a basic application, you can start with the tutorial sample. If you want to use LDAP,

with an embedded test server, then use the LDAP sample as a starting point. The reference manual also includes <http://static.springsource.org/spring-security/site/docs/3.1.x/reference/springsecurity-single.html#appendix-dependencies> [an appendix] listing the first-level dependencies for each Spring Security module with some information on whether they are optional and what they are required for.

If you are building your project with maven, then adding the appropriate Spring Security modules as dependencies to your pom.xml will automatically pull in the core jars that the framework requires. Any which are marked as "optional" in the Spring Security POM files will have to be added to your own pom.xml file if you need them.

What dependencies are needed to run an embedded ApacheDS LDAP server?

If you are using Maven, you need to add the following to your pom dependencies:

```
<dependency>
  <groupId>org.apache.directory.server</groupId>
  <artifactId>apacheds-core</artifactId>
  <version>1.5.5</version>
  <scope>runtime</scope>
</dependency>
<dependency>
  <groupId>org.apache.directory.server</groupId>
  <artifactId>apacheds-server-jndi</artifactId>
  <version>1.5.5</version>
  <scope>runtime</scope>
</dependency>
```

The other required jars should be pulled in transitively.

What is a UserDetailsService and do I need one?

`UserDetailsService` is a DAO interface for loading data that is specific to a user account. It has no other function other to load that data for use by other components within the framework. It is not responsible for authenticating the user. Authenticating a user with a username/password combination is most commonly performed by the `DaoAuthenticationProvider`, which is injected with a `UserDetailsService` to allow it to load the password (and other data) for a user in order to compare it with the submitted value. Note that if you are using LDAP, [this approach may not work](#).

If you want to customize the authentication process then you should implement `AuthenticationProvider` yourself. See this [blog article](#) for an example integrating Spring Security authentication with Google App Engine.

Common "Howto" Requests

1. the section called "I need to login in with more information than just the username."
2. the section called "How do I apply different intercept-url constraints where only the fragment value of the requested URLs differs (e.g./foo#bar and /foo#blah?)"
3. the section called "How do I access the user's IP Address (or other web-request data) in a UserDetailsService?"
4. the section called "How do I access the HttpSession from a UserDetailsService?"
5. the section called "How do I access the user's password in a UserDetailsService?"
6. the section called "How do I define the secured URLs within an application dynamically?"

7. the section called “How do I authenticate against LDAP but load user roles from a database?”
8. the section called “I want to modify the property of a bean that is created by the namespace, but there is nothing in the schema to support it.”

I need to login in with more information than just the username.

How do I add support for extra login fields (e.g. a company name)?

This question comes up repeatedly in the Spring Security forum so you will find more information there by searching the archives (or through google).

The submitted login information is processed by an instance of `UsernamePasswordAuthenticationFilter`. You will need to customize this class to handle the extra data field(s). One option is to use your own customized authentication token class (rather than the standard `UsernamePasswordAuthenticationToken`), another is simply to concatenate the extra fields with the username (for example, using a ":" as the separator) and pass them in the username property of `UsernamePasswordAuthenticationToken`.

You will also need to customize the actual authentication process. If you are using a custom authentication token class, for example, you will have to write an `AuthenticationProvider` to handle it (or extend the standard `DaoAuthenticationProvider`). If you have concatenated the fields, you can implement your own `UserDetailsService` which splits them up and loads the appropriate user data for authentication.

How do I apply different intercept-url constraints where only the fragment value of the requested URLs differs (e.g./foo#bar and /foo#blah)?

You can't do this, since the fragment is not transmitted from the browser to the server. The URLs above are identical from the server's perspective. This is a common question from GWT users.

How do I access the user's IP Address (or other web-request data) in a UserDetailsService?

Obviously you can't (without resorting to something like thread-local variables) since the only information supplied to the interface is the username. Instead of implementing `UserDetailsService`, you should implement `AuthenticationProvider` directly and extract the information from the supplied `Authentication` token.

In a standard web setup, the `getDetails()` method on the `Authentication` object will return an instance of `WebAuthenticationDetails`. If you need additional information, you can inject a custom `AuthenticationDetailsSource` into the authentication filter you are using. If you are using the namespace, for example with the `<form-login>` element, then you should remove this element and replace it with a `<custom-filter>` declaration pointing to an explicitly configured `UsernamePasswordAuthenticationFilter`.

How do I access the HttpSession from a UserDetailsService?

You can't, since the `UserDetailsService` has no awareness of the servlet API. If you want to store custom user data, then you should customize the `UserDetails` object which is returned. This can then be accessed at any point, via the thread-local `SecurityContextHolder`. A call to `SecurityContextHolder.getContext().getAuthentication().getPrincipal()` will return this custom object.

If you really need to access the session, then it must be done by customizing the web tier.

How do I access the user's password in a UserDetailsService?

You can't (and shouldn't). You are probably misunderstanding its purpose. See "[What is a UserDetailsService?](#)" above.

How do I define the secured URLs within an application dynamically?

People often ask about how to store the mapping between secured URLs and security metadata attributes in a database, rather than in the application context.

The first thing you should ask yourself is if you really need to do this. If an application requires securing, then it also requires that the security be tested thoroughly based on a defined policy. It may require auditing and acceptance testing before being rolled out into a production environment. A security-conscious organization should be aware that the benefits of their diligent testing process could be wiped out instantly by allowing the security settings to be modified at runtime by changing a row or two in a configuration database. If you have taken this into account (perhaps using multiple layers of security within your application) then Spring Security allows you to fully customize the source of security metadata. You can make it fully dynamic if you choose.

Both method and web security are protected by subclasses of `AbstractSecurityInterceptor` which is configured with a `SecurityMetadataSource` from which it obtains the metadata for a particular method or filter invocation. For web security, the interceptor class is `FilterSecurityInterceptor` and it uses the marker interface `FilterInvocationSecurityMetadataSource`. The "secured object" type it operates on is a `FilterInvocation`. The default implementation which is used (both in the namespace `<http>` and when configuring the interceptor explicitly, stores the list of URL patterns and their corresponding list of "configuration attributes" (instances of `ConfigAttribute`) in an in-memory map.

To load the data from an alternative source, you must be using an explicitly declared security filter chain (typically Spring Security's `FilterChainProxy`) in order to customize the `FilterSecurityInterceptor` bean. You can't use the namespace. You would then implement `FilterInvocationSecurityMetadataSource` to load the data as you please for a particular `FilterInvocation`³⁷. A very basic outline would look something like this:

```
public class MyFilterSecurityMetadataSource implements FilterInvocationSecurityMetadataSource {

    public List<ConfigAttribute> getAttributes(Object object) {
        FilterInvocation fi = (FilterInvocation) object;
        String url = fi.getRequestUrl();
        String httpMethod = fi.getRequest().getMethod();
        List<ConfigAttribute> attributes = new ArrayList<ConfigAttribute>();

        // Lookup your database (or other source) using this information and populate the
        // list of attributes

        return attributes;
    }

    public Collection<ConfigAttribute> getAllConfigAttributes() {
        return null;
    }

    public boolean supports(Class<?> clazz) {
        return FilterInvocation.class.isAssignableFrom(clazz);
    }
}
```

³⁷The `FilterInvocation` object contains the `HttpServletRequest`, so you can obtain the URL or any other relevant information on which to base your decision on what the list of returned attributes will contain.

For more information, look at the code for `DefaultFilterInvocationSecurityMetadataSource`.

How do I authenticate against LDAP but load user roles from a database?

The `LdapAuthenticationProvider` bean (which handles normal LDAP authentication in Spring Security) is configured with two separate strategy interfaces, one which performs the authentication and one which loads the user authorities, called `LdapAuthenticator` and `LdapAuthoritiesPopulator` respectively. The `DefaultLdapAuthoritiesPopulator` loads the user authorities from the LDAP directory and has various configuration parameters to allow you to specify how these should be retrieved.

To use JDBC instead, you can implement the interface yourself, using whatever SQL is appropriate for your schema:

```
public class MyAuthoritiesPopulator implements LdapAuthoritiesPopulator {
    @Autowired
    JdbcTemplate template;

    List<GrantedAuthority> getGrantedAuthorities(DirContextOperations userData, String username) {
        List<GrantedAuthority> = template.query("select role from roles where username = ?",
            new
String[] {username},
            new
RowMapper<GrantedAuthority>() {
                /**
                 * We're assuming here that you're using the standard convention of using the role
                 * prefix "ROLE_" to mark attributes which are supported by Spring Security's RoleVoter.
                 */
                public GrantedAuthority mapRow(ResultSet rs, int rowNum) throws SQLException {
                    return new SimpleGrantedAuthority("ROLE_" + rs.getString(1));
                }
            }
        );
    }
}
```

You would then add a bean of this type to your application context and inject it into the `LdapAuthenticationProvider`. This is covered in the section on configuring LDAP using explicit Spring beans in the LDAP chapter of the reference manual. Note that you can't use the namespace for configuration in this case. You should also consult the Javadoc for the relevant classes and interfaces.

I want to modify the property of a bean that is created by the namespace, but there is nothing in the schema to support it.

What can I do short of abandoning namespace use?

The namespace functionality is intentionally limited, so it doesn't cover everything that you can do with plain beans. If you want to do something simple, like modify a bean, or inject a different dependency, you can do this by adding a `BeanPostProcessor` to your configuration. More information can be found in the [Spring Reference Manual](#). In order to do this, you need to know a bit about which beans are created, so you should also read the blog article in the above question on [how the namespace maps to Spring beans](#).

Normally, you would add the functionality you require to the `postProcessBeforeInitialization` method of `BeanPostProcessor`. Let's say that you want to customize the `AuthenticationDetailsSource` used by the `UsernamePasswordAuthenticationFilter`, (created by the `form-login` element). You want to extract a particular header called `CUSTOM_HEADER` from the request and make use of it while authenticating the user. The processor class would look like this:

```
public class BeanPostProcessor implements BeanPostProcessor {

    public Object postProcessAfterInitialization(Object bean, String name) {
        if (bean instanceof UsernamePasswordAuthenticationFilter) {
            System.out.println("***** Post-processing " + name);
            ((UsernamePasswordAuthenticationFilter)bean).setAuthenticationDetailsSource(
                new AuthenticationDetailsSource() {
                    public Object buildDetails(Object context) {
                        return
                            ((HttpServletRequest)context).getHeader("CUSTOM_HEADER");
                    }
                });
        }
        return bean;
    }

    public Object postProcessBeforeInitialization(Object bean, String name) {
        return bean;
    }
}
```

You would then register this bean in your application context. Spring will automatically invoke it on the beans defined in the application context.

Part III. Reactive Applications

21. WebFlux Security

Spring Security's WebFlux support relies on a `WebFilter` and works the same for Spring WebFlux and Spring WebFlux.Fn. You can find a few sample applications that demonstrate the code below:

- Hello WebFlux [hellowebflux](#)
- Hello WebFlux.Fn [hellowebfluxfn](#)
- Hello WebFlux Method [hellowebflux-method](#)

21.1 Minimal WebFlux Security Configuration

You can find a minimal WebFlux Security configuration below:

```
@EnableWebFluxSecurity
public class HelloWebfluxSecurityConfig {

    @Bean
    public MapReactiveUserDetailsService userDetailsService() {
        UserDetails user = User.withDefaultPasswordEncoder()
            .username("user")
            .password("user")
            .roles("USER")
            .build();
        return new MapReactiveUserDetailsService(user);
    }
}
```

This configuration provides form and http basic authentication, sets up authorization to require an authenticated user for accessing any page, sets up a default log in page and a default log out page, sets up security related HTTP headers, CSRF protection, and more.

21.2 Explicit WebFlux Security Configuration

You can find an explicit version of the minimal WebFlux Security configuration below:

```
@Configuration
@EnableWebFluxSecurity
public class HelloWebfluxSecurityConfig {

    @Bean
    public MapReactiveUserDetailsService userDetailsService() {
        UserDetails user = User.withDefaultPasswordEncoder()
            .username("user")
            .password("user")
            .roles("USER")
            .build();
        return new MapReactiveUserDetailsService(user);
    }

    @Bean
    public SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
        http
            .authorizeExchange(exchanges ->
                exchanges
                    .anyExchange().authenticated()
            )
            .httpBasic(withDefaults())
            .formLogin(withDefaults());
        return http.build();
    }
}
```

This configuration explicitly sets up all the same things as our minimal configuration. From here you can easily make the changes to the defaults.

22. Protection Against Exploits

22.1 Cross Site Request Forgery (CSRF) for WebFlux Environments

This section discusses Spring Security's [Cross Site Request Forgery \(CSRF\)](#) support for WebFlux environments.

Using Spring Security CSRF Protection

The steps to using Spring Security's CSRF protection are outlined below:

- [Use proper HTTP verbs](#)
- [Configure CSRF Protection](#)
- [Include the CSRF Token](#)

Use proper HTTP verbs

The first step to protecting against CSRF attacks is to ensure your website uses proper HTTP verbs. This is covered in detail in [Safe Methods Must be Idempotent](#).

Configure CSRF Protection

The next step is to configure Spring Security's CSRF protection within your application. Spring Security's CSRF protection is enabled by default, but you may need to customize the configuration. Below are a few common customizations.

Custom CsrfTokenRepository

By default Spring Security stores the expected CSRF token in the `WebSession` using `WebSessionServerCsrfTokenRepository`. There can be cases where users will want to configure a custom `ServerCsrfTokenRepository`. For example, it might be desirable to persist the `CsrfToken` in a cookie to [support a JavaScript based application](#).

By default the `CookieServerCsrfTokenRepository` will write to a cookie named `XSRF-TOKEN` and read it from a header named `X-XSRF-TOKEN` or the HTTP parameter `_csrf`. These defaults come from [AngularJS](#)

You can configure `CookieCsrfTokenRepository` in Java Configuration using:

```
@Bean
public SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .csrf(csrf -> csrf.csrfTokenRepository(CookieServerCsrfTokenRepository.withHttpOnlyFalse()))
    return http.build();
}
```

Example 22.1 Store CSRF Token in a Cookie with Java Configuration

Note

The sample explicitly sets `cookieHttpOnly=false`. This is necessary to allow JavaScript (i.e. AngularJS) to read it. If you do not need the ability to read the cookie with

JavaScript directly, it is recommended to omit `cookieHttpOnly=false` (by using `new CookieServerCsrfTokenRepository()` instead) to improve security.

Disable CSRF Protection

CSRF protection is enabled by default. However, it is simple to disable CSRF protection if it [makes sense for your application](#).

The Java configuration below will disable CSRF protection.

```
@Bean
public SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .csrf(csrf -> csrf.disable())
    return http.build();
}
```

Example 22.2 Disable CSRF Java Configuration

Include the CSRF Token

In order for the [synchronizer token pattern](#) to protect against CSRF attacks, we must include the actual CSRF token in the HTTP request. This must be included in a part of the request (i.e. form parameter, HTTP header, etc) that is not automatically included in the HTTP request by the browser.

Spring Security's [CsrfWebFilter](#) exposes a [Mono<CsrfToken>](#) as a `ServerWebExchange` attribute named `org.springframework.security.web.server.csrf.CsrfToken`. This means that any view technology can access the `Mono<CsrfToken>` to expose the expected token as either a [form](#) or [meta tag](#).

If your view technology does not provide a simple way to subscribe to the `Mono<CsrfToken>`, a common pattern is to use Spring's `@ControllerAdvice` to expose the `CsrfToken` directly. For example, the following code will place the `CsrfToken` on the default attribute name (`_csrf`) used by Spring Security's [CsrfRequestDataValueProcessor](#) to automatically include the CSRF token as a hidden input.

```
@ControllerAdvice
public class SecurityControllerAdvice {
    @ModelAttribute
    Mono<CsrfToken> csrfToken(ServerWebExchange exchange) {
        Mono<CsrfToken> csrfToken = exchange.getAttribute(CsrfToken.class.getName());
        return csrfToken.doOnSuccess(token -> exchange.getAttributes()
            .put(CsrfRequestDataValueProcessor.DEFAULT_CSRF_ATTR_NAME, token));
    }
}
```

Example 22.3 CsrfToken as @ModelAttribute

Fortunately, Thymeleaf provides [integration](#) that works without any additional work.

Form URL Encoded

In order to post an HTML form the CSRF token must be included in the form as a hidden input. For example, the rendered HTML might look like:

```
<input type="hidden"
    name="_csrf"
    value="4bfd1575-3ad1-4d21-96c7-4ef2d9f86721"/>
```

Example 22.4 CSRF Token HTML

Next we will discuss various ways of including the CSRF token in a form as a hidden input.

Automatic CSRF Token Inclusion

Spring Security's CSRF support provides integration with Spring's [RequestDataValueProcessor](#) via its [CsrfRequestDataValueProcessor](#). In order for `CsrfRequestDataValueProcessor` to work, the `Mono<CsrfToken>` must be subscribed to and the `CsrfToken` must be [exposed as an attribute](#) that matches [DEFAULT_CSRF_ATTR_NAME](#).

Fortunately, Thymeleaf [provides support](#) to take care of all the boilerplate for you by integrating with `RequestDataValueProcessor` to ensure that forms that have an unsafe HTTP method (i.e. post) will automatically include the actual CSRF token.

CsrfToken Request Attribute

If the [other options](#) for including the actual CSRF token in the request do not work, you can take advantage of the fact that the `Mono<CsrfToken>` [is exposed](#) as a `ServerWebExchange` attribute named `org.springframework.security.web.server.csrf.CsrfToken`.

The Thymeleaf sample below assumes that you [expose](#) the `CsrfToken` on an attribute named `_csrf`.

```
<form th:action="@{/logout}"
      method="post">
  <input type="submit"
        value="Log out" />
  <input type="hidden"
        th:name="${_csrf.parameterName}"
        th:value="${_csrf.token}"/>
</form>
```

Example 22.5 CSRF Token in Form with Request Attribute

Ajax and JSON Requests

If you are using JSON, then it is not possible to submit the CSRF token within an HTTP parameter. Instead you can submit the token within a HTTP header.

In the following sections we will discuss various ways of including the CSRF token as an HTTP request header in JavaScript based applications.

Automatic Inclusion

Spring Security can easily be [configured](#) to store the expected CSRF token in a cookie. By storing the expected CSRF in a cookie, JavaScript frameworks like [AngularJS](#) will automatically include the actual CSRF token in the HTTP request headers.

Meta tags

An alternative pattern to [exposing the CSRF in a cookie](#) is to include the CSRF token within your meta tags. The HTML might look something like this:

```
<html>
<head>
  <meta name="_csrf" content="4bfd1575-3ad1-4d21-96c7-4ef2d9f86721"/>
  <meta name="_csrf_header" content="X-CSRF-TOKEN"/>
  <!-- ... -->
</head>
<!-- ... -->
```

Example 22.6 CSRF meta tag HTML

Once the meta tags contained the CSRF token, the JavaScript code would read the meta tags and include the CSRF token as a header. If you were using jQuery, this could be done with the following:

```
$(function () {
    var token = $("meta[name='_csrf']").attr("content");
    var header = $("meta[name='_csrf_header']").attr("content");
    $(document).ajaxSend(function(e, xhr, options) {
        xhr.setRequestHeader(header, token);
    });
});
```

Example 22.7 AJAX send CSRF Token

The sample below assumes that you [expose](#) the `CsrfToken` on an attribute named `_csrf`. An example of doing this with Thymeleaf is shown below:

```
<html>
<head>
  <meta name="_csrf" th:content="${_csrf.token}"/>
  <!-- default header name is X-CSRF-TOKEN -->
  <meta name="_csrf_header" th:content="${_csrf.headerName}"/>
  <!-- ... -->
</head>
<!-- ... -->
```

Example 22.8 CSRF meta tag JSP

CSRF Considerations

There are a few special considerations to consider when implementing protection against CSRF attacks. This section discusses those considerations as it pertains to WebFlux environments. Refer to the section called “CSRF Considerations” for a more general discussion.

Logging In

It is important to [require CSRF for log in](#) requests to protect against forging log in attempts. Spring Security’s WebFlux support does this out of the box.

Logging Out

It is important to [require CSRF for log out](#) requests to protect against forging log out attempts. By default Spring Security’s `LogoutWebFilter` only processes HTTP post requests. This ensures that log out requires a CSRF token and that a malicious user cannot forcibly log out your users.

The easiest approach is to use a form to log out. If you really want a link, you can use JavaScript to have the link perform a POST (i.e. maybe on a hidden form). For browsers with JavaScript that is disabled, you can optionally have the link take the user to a log out confirmation page that will perform the POST.

If you really want to use HTTP GET with logout you can do so, but remember this is generally not recommended. For example, the following Java Configuration will perform logout with the URL `/logout` is requested with any HTTP method:

```
@Bean
public SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .logout(logout -> logout.requiresLogout(new PathPatternParserServerWebExchangeMatcher("/
logout")))
        return http.build();
}
```

Example 22.9 Log out with HTTP GET

CSRF and Session Timeouts

By default Spring Security stores the CSRF token in the `WebSession`. This can lead to a situation where the session expires which means there is not an expected CSRF token to validate against.

We've already discussed [general solutions](#) to session timeouts. This section discusses the specifics of CSRF timeouts as it pertains to the WebFlux support.

It is simple to change storage of the expected CSRF token to be in a cookie. For details, refer to the the section called "Custom `CsrfTokenRepository`" section.

Multipart (file upload)

We have [already discussed](#) how protecting multipart requests (file uploads) from CSRF attacks causes a [chicken and the egg](#) problem. This section discusses how to implement placing the CSRF token in the [body](#) and [url](#) within a WebFlux application.

Note

More information about using multipart forms with Spring can be found within the [Multipart Data](#) section of the Spring reference.

Place CSRF Token in the Body

We have [already discussed](#) the trade-offs of placing the CSRF token in the body.

In a WebFlux application, this can be configured with the following configuration:

```
@Bean
public SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .csrf(csrf -> csrf.tokenFromMultipartDataEnabled(true))
        return http.build();
}
```

Example 22.10 Enable obtaining CSRF token from multipart/form-data

Include CSRF Token in URL

We have [already discussed](#) the trade-offs of placing the CSRF token in the URL. Since the `CsrfToken` is exposed as an `ServerHttpRequest` [request attribute](#), we can use that to create an action with the CSRF token in it. An example with Thymeleaf is shown below:

```
<form method="post"
    th:action="@{/upload(${_csrf.parameterName}=${_csrf.token})}"
    enctype="multipart/form-data">
```

Example 22.11 CSRF Token in Action

HiddenHttpMethodFilter

We have [already discussed](#) overriding the HTTP method.

In a Spring WebFlux application, overriding the HTTP method is done using [HiddenHttpMethodFilter](#).

22.2 Security HTTP Response Headers

This section discusses Spring Security's support for adding various security headers to the response of WebFlux.

Default Security Headers

Spring Security allows users to easily inject the default security headers to assist in protecting their application. The default for Spring Security is to include the following headers:

```
Cache-Control: no-cache, no-store, max-age=0, must-revalidate
Pragma: no-cache
Expires: 0
X-Content-Type-Options: nosniff
Strict-Transport-Security: max-age=31536000 ; includeSubDomains
X-Frame-Options: DENY
X-XSS-Protection: 1; mode=block
```

Note

Strict-Transport-Security is only added on HTTPS requests

For additional details on each of these headers, refer to the corresponding sections:

- [Cache Control](#)
- [Content Type Options](#)
- [HTTP Strict Transport Security](#)
- [X-Frame-Options](#)
- [X-XSS-Protection](#)

While each of these headers are considered best practice, it should be noted that not all clients utilize the headers, so additional testing is encouraged.

You can customize specific headers. For example, assume that want your HTTP response headers to look like the following:

```
Cache-Control: no-cache, no-store, max-age=0, must-revalidate
Pragma: no-cache
Expires: 0
X-Content-Type-Options: nosniff
X-Frame-Options: SAMEORIGIN
X-XSS-Protection: 1; mode=block
```

Specifically, you want all of the default headers with the following customizations:

- [X-Frame-Options](#) to allow any request from same domain
- [HTTP Strict Transport Security \(HSTS\)](#) will not be added to the response

You can easily do this with the following Java Configuration:


```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .hsts(hsts ->
                    hsts
                        .disable()
                )
                .frameOptions(frameOptions ->
                    frameOptions
                        .mode(Mode.SAMEORIGIN)
                )
        );
    return http.build();
}

```

If you do not want the defaults to be added and want explicit control over what should be used, you can disable the defaults. An example for both Java and XML based configuration is provided below:

If necessary, you can disable all of the HTTP Security response headers with the following Java Configuration:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .disable()
        );
    return http.build();
}

```

Cache Control

In the past Spring Security required you to provide your own cache control for your web application. This seemed reasonable at the time, but browser caches have evolved to include caches for secure connections as well. This means that a user may view an authenticated page, log out, and then a malicious user can use the browser history to view the cached page. To help mitigate this Spring Security has added cache control support which will insert the following headers into you response by default.

```

Cache-Control: no-cache, no-store, max-age=0, must-revalidate
Pragma: no-cache
Expires: 0

```

If you actually want to cache specific responses, your application can selectively set the cache control headers to override the header set by Spring Security. This is useful to ensure things like CSS, JavaScript, and images are properly cached.

You can also disable cache control using the following Java Configuration:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .cache(cache -> cache.disable())
        );
    return http.build();
}

```

Content Type Options

Historically browsers, including Internet Explorer, would try to guess the content type of a request using [content sniffing](#). This allowed browsers to improve the user experience by guessing the content type on resources that had not specified the content type. For example, if a browser encountered a JavaScript file that did not have the content type specified, it would be able to guess the content type and then execute it.

Note

== There are many additional things one should do (i.e. only display the document in a distinct domain, ensure Content-Type header is set, sanitize the document, etc) when allowing content to be uploaded. However, these measures are out of the scope of what Spring Security provides. It is also important to point out when disabling content sniffing, you must specify the content type in order for things to work properly. ==

The problem with content sniffing is that this allowed malicious users to use polyglots (i.e. a file that is valid as multiple content types) to execute XSS attacks. For example, some sites may allow users to submit a valid postscript document to a website and view it. A malicious user might create a [postscript document that is also a valid JavaScript file](#) and execute a XSS attack with it.

Content sniffing can be disabled by adding the following header to our response:

```
X-Content-Type-Options: nosniff
```

Just as with the cache control element, the nosniff directive is added by default. However, if need to disable the header, the following may be used:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .contentTypeOptions(contentTypeOptions -> contentTypeOptions.disable())
        );
    return http.build();
}
```

HTTP Strict Transport Security (HSTS)

When you type in your bank's website, do you enter `mybank.example.com` or do you enter <https://mybank.example.com>? If you omit the https protocol, you are potentially vulnerable to [Man in the Middle attacks](#). Even if the website performs a redirect to <https://mybank.example.com> a malicious user could intercept the initial HTTP request and manipulate the response (i.e. redirect to <https://mibank.example.com> and steal their credentials).

Many users omit the https protocol and this is why [HTTP Strict Transport Security \(HSTS\)](#) was created. Once `mybank.example.com` is added as a [HSTS host](#), a browser can know ahead of time that any request to `mybank.example.com` should be interpreted as <https://mybank.example.com>. This greatly reduces the possibility of a Man in the Middle attack occurring.

Note

== In accordance with [RFC6797](#), the HSTS header is only injected into HTTPS responses. In order for the browser to acknowledge the header, the browser must first trust the CA that signed the SSL certificate used to make the connection (not just the SSL certificate). ==

One way for a site to be marked as a HSTS host is to have the host preloaded into the browser. Another is to add the "Strict-Transport-Security" header to the response. For example the following would instruct the browser to treat the domain as an HSTS host for a year (there are approximately 31536000 seconds in a year):

```
Strict-Transport-Security: max-age=31536000 ; includeSubDomains ; preload
```

The optional `includeSubDomains` directive instructs Spring Security that subdomains (i.e. `secure.mybank.example.com`) should also be treated as an HSTS domain.

The optional `preload` directive instructs Spring Security that domain should be preloaded in browser as HSTS domain. For more details on HSTS preload please see <https://hstspreload.org>.

As with the other headers, Spring Security adds HSTS by default. You can customize HSTS headers with Java Configuration:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .hsts(hsts ->
                    hsts
                        .includeSubdomains(true)
                        .preload(true)
                        .maxAge(Duration.ofDays(365))
                    )
                )
        );
    return http.build();
}
```

X-Frame-Options

Allowing your website to be added to a frame can be a security issue. For example, using clever CSS styling users could be tricked into clicking on something that they were not intending ([video demo](#)). For example, a user that is logged into their bank might click a button that grants access to other users. This sort of attack is known as [Clickjacking](#).

Note

== Another modern approach to dealing with clickjacking is to use the section called "Content Security Policy (CSP)". ==

There are a number ways to mitigate clickjacking attacks. For example, to protect legacy browsers from clickjacking attacks you can use [frame breaking code](#). While not perfect, the frame breaking code is the best you can do for the legacy browsers.

A more modern approach to address clickjacking is to use [X-Frame-Options](#) header:

```
X-Frame-Options: DENY
```

The X-Frame-Options response header instructs the browser to prevent any site with this header in the response from being rendered within a frame. By default, Spring Security disables rendering within an iframe.

You can customize X-Frame-Options with Java Configuration using the following:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .frameOptions(frameOptions ->
                    frameOptions
                        .mode(SAMEORIGIN)
                )
        );
    return http.build();
}
```

X-XSS-Protection

Some browsers have built in support for filtering out [reflected XSS attacks](#). This is by no means foolproof, but does assist in XSS protection.

The filtering is typically enabled by default, so adding the header typically just ensures it is enabled and instructs the browser what to do when a XSS attack is detected. For example, the filter might try to change the content in the least invasive way to still render everything. At times, this type of replacement can become a [XSS vulnerability in itself](#). Instead, it is best to block the content rather than attempt to fix it. To do this we can add the following header:

```
X-XSS-Protection: 1; mode=block
```

This header is included by default. However, we can customize with Java Configuration with the following:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .xssProtection(xssProtection -> xssProtection.disable())
        );
    return http.build();
}
```

Content Security Policy (CSP)

[Content Security Policy \(CSP\)](#) is a mechanism that web applications can leverage to mitigate content injection vulnerabilities, such as cross-site scripting (XSS). CSP is a declarative policy that provides a facility for web application authors to declare and ultimately inform the client (user-agent) about the sources from which the web application expects to load resources.

Note

== Content Security Policy is not intended to solve all content injection vulnerabilities. Instead, CSP can be leveraged to help reduce the harm caused by content injection attacks. As a first line of defense, web application authors should validate their input and encode their output. ==

A web application may employ the use of CSP by including one of the following HTTP headers in the response:

- **Content-Security-Policy**
- **Content-Security-Policy-Report-Only**

Each of these headers are used as a mechanism to deliver a **security policy** to the client. A security policy contains a set of **security policy directives** (for example, *script-src* and *object-src*), each responsible for declaring the restrictions for a particular resource representation.

For example, a web application can declare that it expects to load scripts from specific, trusted sources, by including the following header in the response:

```
Content-Security-Policy: script-src https://trustedscripts.example.com
```

An attempt to load a script from another source other than what is declared in the *script-src* directive will be blocked by the user-agent. Additionally, if the [report-uri](#) directive is declared in the security policy, then the violation will be reported by the user-agent to the declared URL.

For example, if a web application violates the declared security policy, the following response header will instruct the user-agent to send violation reports to the URL specified in the policy's *report-uri* directive.

```
Content-Security-Policy: script-src https://trustedscripts.example.com; report-uri /csp-report-endpoint/
```

[Violation reports](#) are standard JSON structures that can be captured either by the web application's own API or by a publicly hosted CSP violation reporting service, such as, [REPORT-URI](#).

The **Content-Security-Policy-Report-Only** header provides the capability for web application authors and administrators to monitor security policies, rather than enforce them. This header is typically used when experimenting and/or developing security policies for a site. When a policy is deemed effective, it can be enforced by using the *Content-Security-Policy* header field instead.

Given the following response header, the policy declares that scripts may be loaded from one of two possible sources.

```
Content-Security-Policy-Report-Only: script-src 'self' https://trustedscripts.example.com; report-uri /csp-report-endpoint/
```

If the site violates this policy, by attempting to load a script from *evil.com*, the user-agent will send a violation report to the declared URL specified by the *report-uri* directive, but still allow the violating resource to load nevertheless.

Configuring Content Security Policy

It's important to note that Spring Security **does not add** Content Security Policy by default. The web application author must declare the security policy(s) to enforce and/or monitor for the protected resources.

For example, given the following security policy:

```
script-src 'self' https://trustedscripts.example.com; object-src https://trustedplugins.example.com;
report-uri /csp-report-endpoint/
```

You can enable the CSP header using Java configuration as shown below:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .contentSecurityPolicy(contentSecurityPolicy ->
                    contentSecurityPolicy
                        .policyDirectives("script-src 'self' https://trustedscripts.example.com; object-
src https://trustedplugins.example.com; report-uri /csp-report-endpoint/")
                )
        );
    return http.build();
}

```

To enable the CSP *'report-only'* header, provide the following Java configuration:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .contentSecurityPolicy(contentSecurityPolicy ->
                    contentSecurityPolicy
                        .policyDirectives("script-src 'self' https://trustedscripts.example.com; object-
src https://trustedplugins.example.com; report-uri /csp-report-endpoint/")
                        .reportOnly()
                )
        );
    return http.build();
}

```

Additional Resources

Applying Content Security Policy to a web application is often a non-trivial undertaking. The following resources may provide further assistance in developing effective security policies for your site.

[An Introduction to Content Security Policy](#)

[CSP Guide - Mozilla Developer Network](#)

[W3C Candidate Recommendation](#)

Referrer Policy

[Referrer Policy](#) is a mechanism that web applications can leverage to manage the referrer field, which contains the last page the user was on.

Spring Security's approach is to use [Referrer Policy](#) header, which provides different [policies](#):

```
Referrer-Policy: same-origin
```

The Referrer-Policy response header instructs the browser to let the destination know the source where the user was previously.

Configuring Referrer Policy

Spring Security **doesn't add** Referrer Policy header by default.

You can enable the Referrer-Policy header using Java configuration as shown below:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .referrerPolicy(referrerPolicy ->
                    referrerPolicy
                        .policy(ReferrerPolicy.SAME_ORIGIN)
                )
        );
    return http.build();
}
```

Feature Policy

[Feature Policy](#) is a mechanism that allows web developers to selectively enable, disable, and modify the behavior of certain APIs and web features in the browser.

```
Feature-Policy: geolocation 'self'
```

With Feature Policy, developers can opt-in to a set of "policies" for the browser to enforce on specific features used throughout your site. These policies restrict what APIs the site can access or modify the browser's default behavior for certain features.

Configuring Feature Policy

Spring Security **doesn't add** Feature Policy header by default.

You can enable the Feature-Policy header using Java configuration as shown below:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .headers(headers ->
            headers
                .featurePolicy("geolocation 'self'")
        );
    return http.build();
}
```

Clear Site Data

[Clear Site Data](#) is a mechanism by which any browser-side data - cookies, local storage, and the like - can be removed when an HTTP response contains this header:

```
Clear-Site-Data: "cache", "cookies", "storage", "executionContexts"
```

This is a nice clean-up action to perform on logout.

Configuring Clear Site Data

Spring Security **doesn't add** the Clear Site Data header by default.

You can configure your application to send down this header on logout like so:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    ServerLogoutHandler securityContext = new SecurityContextServerLogoutHandler();
    ServerLogoutHandler clearSiteData = new HeaderWriterServerLogoutHandler(new
ClearSiteDataServerHttpHeadersWriter());
    DelegatingServerLogoutHandler logoutHandler = new DelegatingServerLogoutHandler(securityContext,
clearSiteData);

    http
        // ...
        .logout()
            .logoutHandler(logoutHandler);
    return http.build();
}

```

Note

It's not recommended that you configure this header writer via the `headers()` directive. The reason for this is that any session state, say the `JSESSIONID` cookie, would be removed, effectively logging the user out.

22.3 Redirect to HTTPS

HTTPS is required to provide a secure application. Spring Security can be configured to perform a redirect to https using the following Java Configuration:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .redirectToHttps(withDefaults());
    return http.build();
}

```

The configuration can easily be wrapped around an if statement to only be turned on in production. Alternatively, it can be enabled by looking for a property about the request that only happens in production. For example, if the production environment adds a header named `X-Forwarded-Proto` the following Java Configuration could be used:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .redirectToHttps(redirectToHttps ->
            redirectToHttps
                .httpsRedirectWhen(e -> e.getRequest().getHeaders().containsKey("X-Forwarded-Proto"))
        );
    return http.build();
}

```


23. OAuth2 WebFlux

Spring Security provides OAuth2 and WebFlux integration for reactive applications.

23.1 OAuth 2.0 Login

The OAuth 2.0 Login feature provides an application with the capability to have users log in to the application by using their existing account at an OAuth 2.0 Provider (e.g. GitHub) or OpenID Connect 1.0 Provider (such as Google). OAuth 2.0 Login implements the use cases: "Login with Google" or "Login with GitHub".

Note

OAuth 2.0 Login is implemented by using the **Authorization Code Grant**, as specified in the [OAuth 2.0 Authorization Framework](#) and [OpenID Connect Core 1.0](#).

Spring Boot 2.0 Sample

Spring Boot 2.0 brings full auto-configuration capabilities for OAuth 2.0 Login.

This section shows how to configure the [OAuth 2.0 Login WebFlux sample](#) using *Google* as the *Authentication Provider* and covers the following topics:

- [Initial setup](#)
- [Setting the redirect URI](#)
- [Configure application.yml](#)
- [Boot up the application](#)

Initial setup

To use Google's OAuth 2.0 authentication system for login, you must set up a project in the Google API Console to obtain OAuth 2.0 credentials.

Note

[Google's OAuth 2.0 implementation](#) for authentication conforms to the [OpenID Connect 1.0](#) specification and is [OpenID Certified](#).

Follow the instructions on the [OpenID Connect](#) page, starting in the section, "Setting up OAuth 2.0".

After completing the "Obtain OAuth 2.0 credentials" instructions, you should have a new OAuth Client with credentials consisting of a Client ID and a Client Secret.

Setting the redirect URI

The redirect URI is the path in the application that the end-user's user-agent is redirected back to after they have authenticated with Google and have granted access to the OAuth Client ([created in the previous step](#)) on the Consent page.

In the "Set a redirect URI" sub-section, ensure that the **Authorized redirect URIs** field is set to <http://localhost:8080/login/oauth2/code/google>.

Tip

The default redirect URI template is `{baseUrl}/login/oauth2/code/{registrationId}`. The **registrationId** is a unique identifier for the [ClientRegistration](#). For our example, the `registrationId` is `google`.

Important

If the OAuth Client is running behind a proxy server, it is recommended to check [Proxy Server Configuration](#) to ensure the application is correctly configured. Also, see the supported [URI template variables](#) for `redirect-uri`.

Configure `application.yml`

Now that you have a new OAuth Client with Google, you need to configure the application to use the OAuth Client for the *authentication flow*. To do so:

1. Go to `application.yml` and set the following configuration:

```
spring:
  security:
    oauth2:
      client:
        registration: ❶
        google: ❷
          client-id: google-client-id
          client-secret: google-client-secret
```

- ❶ `spring.security.oauth2.client.registration` is the base property prefix for OAuth Client properties.
- ❷ Following the base property prefix is the ID for the [ClientRegistration](#), such as `google`.

Example 23.1 OAuth Client properties

2. Replace the values in the `client-id` and `client-secret` property with the OAuth 2.0 credentials you created earlier.

Boot up the application

Launch the Spring Boot 2.0 sample and go to <http://localhost:8080>. You are then redirected to the default *auto-generated* login page, which displays a link for Google.

Click on the Google link, and you are then redirected to Google for authentication.

After authenticating with your Google account credentials, the next page presented to you is the Consent screen. The Consent screen asks you to either allow or deny access to the OAuth Client you created earlier. Click **Allow** to authorize the OAuth Client to access your email address and basic profile information.

At this point, the OAuth Client retrieves your email address and basic profile information from the [UserInfo Endpoint](#) and establishes an authenticated session.

Using OpenID Provider Configuration

For well known providers, Spring Security provides the necessary defaults for the OAuth Authorization Provider's configuration. If you are working with your own Authorization Provider that supports [OpenID Provider Configuration](#) or [Authorization Server Metadata](#), the [OpenID Provider Configuration Response](#)'s `issuer-uri` can be used to configure the application.

```

spring:
  security:
    oauth2:
      client:
        provider:
          keycloak:
            issuer-uri: https://idp.example.com/auth/realms/demo
        registration:
          keycloak:
            client-id: spring-security
            client-secret: 6cea952f-10d0-4d00-ac79-cc865820dc2c

```

The `issuer-uri` instructs Spring Security to query in series the endpoints <https://idp.example.com/auth/realms/demo/.well-known/openid-configuration>, <https://idp.example.com/.well-known/openid-configuration/auth/realms/demo>, or <https://idp.example.com/.well-known/oauth-authorization-server/auth/realms/demo> to discover the configuration.

Note

Spring Security will query the endpoints one at a time, stopping at the first that gives a 200 response.

The `client-id` and `client-secret` are linked to the provider because `keycloak` is used for both the provider and the registration.

Explicit OAuth2 Login Configuration

A minimal OAuth2 Login configuration is shown below:

```

@Bean
ReactiveClientRegistrationRepository clientRegistrations() {
    ClientRegistration clientRegistration = ClientRegistrations
        .fromIssuerLocation("https://idp.example.com/auth/realms/demo")
        .clientId("spring-security")
        .clientSecret("6cea952f-10d0-4d00-ac79-cc865820dc2c")
        .build();
    return new InMemoryReactiveClientRegistrationRepository(clientRegistration);
}

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .oauth2Login(withDefaults());
    return http.build();
}

```

Additional configuration options can be seen below:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        // ...
        .oauth2Login(oauth2Login ->
            oauth2Login
                .authenticationConverter(converter)
                .authenticationManager(manager)
                .authorizedClientRepository(authorizedClients)
                .clientRegistrationRepository(clientRegistrations)
            );
    return http.build();
}

```

23.2 OAuth2 Client

Spring Security's OAuth Support allows obtaining an access token without authenticating. A basic configuration with Spring Boot can be seen below:

```
spring:
  security:
    oauth2:
      client:
        registration:
          github:
            client-id: replace-with-client-id
            client-secret: replace-with-client-secret
            scope: read:user,public_repo
```

You will need to replace the `client-id` and `client-secret` with values registered with GitHub.

The next step is to instruct Spring Security that you wish to act as an OAuth2 Client so that you can obtain an access token.

```
@Bean
SecurityWebFilterChain configure(ServerHttpSecurity http) throws Exception {
    http
        // ...
        .oauth2Client(withDefaults());
    return http.build();
}
```

You can now leverage Spring Security's Chapter 26, *WebClient* or [@RegisteredOAuth2AuthorizedClient](#) support to obtain and use the access token.

23.3 OAuth 2.0 Resource Server

Spring Security supports protecting endpoints using two forms of OAuth 2.0 [Bearer Tokens](#):

- [JWT](#)
- Opaque Tokens

This is handy in circumstances where an application has delegated its authority management to an [authorization server](#) (for example, Okta or Ping Identity). This authorization server can be consulted by resource servers to authorize requests.

Note

A complete working example for [JWTs](#) is available in the [Spring Security repository](#).

Dependencies

Most Resource Server support is collected into `spring-security-oauth2-resource-server`. However, the support for decoding and verifying JWTs is in `spring-security-oauth2-jose`, meaning that both are necessary in order to have a working resource server that supports JWT-encoded Bearer Tokens.

Minimal Configuration for JWTs

When using [Spring Boot](#), configuring an application as a resource server consists of two basic steps. First, include the needed dependencies and second, indicate the location of the authorization server.

Specifying the Authorization Server

In a Spring Boot application, to specify which authorization server to use, simply do:

```
spring:
  security:
    oauth2:
      resourceserver:
        jwt:
          issuer-uri: https://idp.example.com/issuer
```

Where <https://idp.example.com/issuer> is the value contained in the `iss` claim for JWT tokens that the authorization server will issue. Resource Server will use this property to further self-configure, discover the authorization server's public keys, and subsequently validate incoming JWTs.

Note

To use the `issuer-uri` property, it must also be true that one of <https://idp.example.com/issuer/.well-known/openid-configuration>, <https://idp.example.com/.well-known/openid-configuration/issuer>, or <https://idp.example.com/.well-known/oauth-authorization-server/issuer> is a supported endpoint for the authorization server. This endpoint is referred to as a [Provider Configuration](#) endpoint or a [Authorization Server Metadata](#) endpoint.

And that's it!

Startup Expectations

When this property and these dependencies are used, Resource Server will automatically configure itself to validate JWT-encoded Bearer Tokens.

It achieves this through a deterministic startup process:

1. Hit the Provider Configuration or Authorization Server Metadata endpoint, processing the response for the `jwt: issuer-uri` property
2. Configure the validation strategy to query `jwt: issuer-uri` for valid public keys
3. Configure the validation strategy to validate each JWTs `iss` claim against <https://idp.example.com>.

A consequence of this process is that the authorization server must be up and receiving requests in order for Resource Server to successfully start up.

Note

If the authorization server is down when Resource Server queries it (given appropriate timeouts), then startup will fail.

Runtime Expectations

Once the application is started up, Resource Server will attempt to process any request containing an `Authorization: Bearer` header:

```
GET / HTTP/1.1
Authorization: Bearer some-token-value # Resource Server will process this
```

So long as this scheme is indicated, Resource Server will attempt to process the request according to the Bearer Token specification.

Given a well-formed JWT, Resource Server will:

1. Validate its signature against a public key obtained from the `jwks_url` endpoint during startup and matched against the JWTs header
2. Validate the JWTs `exp` and `nbf` timestamps and the JWTs `iss` claim, and
3. Map each scope to an authority with the prefix `SCOPE_`.

Note

As the authorization server makes available new keys, Spring Security will automatically rotate the keys used to validate the JWT tokens.

The resulting `Authentication#getPrincipal`, by default, is a Spring Security `Jwt` object, and `Authentication#getName` maps to the JWT's `sub` property, if one is present.

From here, consider jumping to:

[How to Configure without Tying Resource Server startup to an authorization server's availability](#)

[How to Configure without Spring Boot](#)

Specifying the Authorization Server JWK Set Uri Directly

If the authorization server doesn't support any configuration endpoints, or if Resource Server must be able to start up independently from the authorization server, then the `jwk-set-uri` can be supplied as well:

```
spring:
  security:
    oauth2:
      resourceserver:
        jwt:
          issuer-uri: https://idp.example.com
          jwk-set-uri: https://idp.example.com/.well-known/jwks.json
```

Note

The JWK Set uri is not standardized, but can typically be found in the authorization server's documentation

Consequently, Resource Server will not ping the authorization server at startup. We still specify the `issuer-uri` so that Resource Server still validates the `iss` claim on incoming JWTs.

Note

This property can also be supplied directly on the [DSL](#).

Overriding or Replacing Boot Auto Configuration

There are two `@Beans` that Spring Boot generates on Resource Server's behalf.

The first is a `SecurityWebFilterChain` that configures the app as a resource server. When including `spring-security-oauth2-jose`, this `WebSecurityConfigurerAdapter` looks like:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        .authorizeExchange(exchanges ->
            exchanges
                .anyExchange().authenticated()
            )
        .oauth2ResourceServer(OAuth2ResourceServerSpec::jwt)
    return http.build();
}
```

If the application doesn't expose a `SecurityWebFilterChain` bean, then Spring Boot will expose the above default one.

Replacing this is as simple as exposing the bean within the application:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        .authorizeExchange(exchanges ->
            exchanges
                .pathMatchers("/message/**").hasAuthority("SCOPE_message:read")
                .anyExchange().authenticated()
            )
        .oauth2ResourceServer(oauth2ResourceServer ->
            oauth2ResourceServer
                .jwt(withDefaults())
            );
    return http.build();
}
```

The above requires the scope of `message:read` for any URL that starts with `/messages/`.

Methods on the `oauth2ResourceServer` DSL will also override or replace auto configuration.

For example, the second `@Bean` Spring Boot creates is a `ReactiveJwtDecoder`, which decodes `String` tokens into validated instances of `Jwt`:

```
@Bean
public ReactiveJwtDecoder jwtDecoder() {
    return ReactiveJwtDecoders.fromIssuerLocation(issuerUri);
}
```

Note

Calling [ReactiveJwtDecoders#fromIssuerLocation](#) is what invokes the Provider Configuration or Authorization Server Metadata endpoint in order to derive the JWK Set Uri. If the application doesn't expose a `ReactiveJwtDecoder` bean, then Spring Boot will expose the above default one.

And its configuration can be overridden using `jwtSetUri()` or replaced using `decoder()`.

Using `jwtSetUri()`

An authorization server's JWK Set Uri can be configured [as a configuration property](#) or it can be supplied in the DSL:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        .authorizeExchange(exchanges ->
            exchanges
                .anyExchange().authenticated()
            )
        .oauth2ResourceServer(oauth2ResourceServer ->
            oauth2ResourceServer
                .jwt(jwt ->
                    jwt
                        .jwkSetUri("https://idp.example.com/.well-known/jwks.json")
                    )
            );
    return http.build();
}

```

Using `jwkSetUri()` takes precedence over any configuration property.

Using `decoder()`

More powerful than `jwkSetUri()` is `decoder()`, which will completely replace any Boot auto configuration of `JwtDecoder`:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        .authorizeExchange()
            .anyExchange().authenticated()
            .and()
        .oauth2ResourceServer()
            .jwt()
                .decoder(myCustomDecoder());
    return http.build();
}

```

This is handy when deeper configuration, like [validation](#), is necessary.

Exposing a `ReactiveJwtDecoder` @Bean

Or, exposing a `ReactiveJwtDecoder` @Bean has the same effect as `decoder()`:

```

@Bean
public ReactiveJwtDecoder jwtDecoder() {
    return NimbusReactiveJwtDecoder.withJwkSetUri(jwkSetUri).build();
}

```

Configuring Trusted Algorithms

By default, `NimbusReactiveJwtDecoder`, and hence `Resource Server`, will only trust and verify tokens using RS256.

You can customize this via [Spring Boot](#) or [the NimbusJwtDecoder builder](#).

Via Spring Boot

The simplest way to set the algorithm is as a property:

```

spring:
  security:
    oauth2:
      resourceserver:
        jwt:
          jws-algorithm: RS512
          jwk-set-uri: https://idp.example.org/.well-known/jwks.json

```


Using a Builder

For greater power, though, we can use a builder that ships with `NimbusReactiveJwtDecoder`:

```
@Bean
ReactiveJwtDecoder jwtDecoder() {
    return NimbusReactiveJwtDecoder.fromJwkSetUri(this.jwkSetUri)
        .jwsAlgorithm(RS512).build();
}
```

Calling `jwsAlgorithm` more than once will configure `NimbusReactiveJwtDecoder` to trust more than one algorithm, like so:

```
@Bean
ReactiveJwtDecoder jwtDecoder() {
    return NimbusReactiveJwtDecoder.fromJwkSetUri(this.jwkSetUri)
        .jwsAlgorithm(RS512).jwsAlgorithm(EC512).build();
}
```

Or, you can call `jwsAlgorithms`:

```
@Bean
ReactiveJwtDecoder jwtDecoder() {
    return NimbusReactiveJwtDecoder.fromJwkSetUri(this.jwkSetUri)
        .jwsAlgorithms(algorithms -> {
            algorithms.add(RS512);
            algorithms.add(EC512);
        }).build();
}
```

Trusting a Single Asymmetric Key

Simpler than backing a Resource Server with a JWK Set endpoint is to hard-code an RSA public key. The public key can be provided via [Spring Boot](#) or by [Using a Builder](#).

Via Spring Boot

Specifying a key via Spring Boot is quite simple. The key's location can be specified like so:

```
spring:
  security:
    oauth2:
      resource-server:
        jwt:
          public-key-location: classpath:my-key.pub
```

Or, to allow for a more sophisticated lookup, you can post-process the `RsaKeyConversionServicePostProcessor`:

```
@Bean
BeanFactoryPostProcessor conversionServiceCustomizer() {
    return beanFactory ->
        beanFactory.getBean(RsaKeyConversionServicePostProcessor.class)
            .setResourceLoader(new CustomResourceLoader());
}
```

Specify your key's location:

```
key.location: hdfs://my-key.pub
```

And then autowire the value:

```
@Value("${key.location}")
RSAPublicKey key;
```

Using a Builder

To wire an `RSAPublicKey` directly, you can simply use the appropriate `NimbusReactiveJwtDecoder` builder, like so:

```
@Bean
public ReactiveJwtDecoder jwtDecoder() {
    return NimbusReactiveJwtDecoder.withPublicKey(this.key).build();
}
```

Trusting a Single Symmetric Key

Using a single symmetric key is also simple. You can simply load in your `SecretKey` and use the appropriate `NimbusReactiveJwtDecoder` builder, like so:

```
@Bean
public ReactiveJwtDecoder jwtDecoder() {
    return NimbusReactiveJwtDecoder.withSecretKey(this.key).build();
}
```

Configuring Authorization

A JWT that is issued from an OAuth 2.0 Authorization Server will typically either have a `scope` or `scp` attribute, indicating the scopes (or authorities) it's been granted, for example:

```
{ ..., "scope" : "messages contacts" }
```

When this is the case, Resource Server will attempt to coerce these scopes into a list of granted authorities, prefixing each scope with the string `"SCOPE_"`.

This means that to protect an endpoint or method with a scope derived from a JWT, the corresponding expressions should include this prefix:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        .authorizeExchange(exchanges ->exchanges
            .mvcMatchers("/contacts/**").hasAuthority("SCOPE_contacts")
            .mvcMatchers("/messages/**").hasAuthority("SCOPE_messages")
            .anyExchange().authenticated()
        )
        .oauth2ResourceServer(OAuth2ResourceServerSpec::jwt);
    return http.build();
}
```

Or similarly with method security:

```
@PreAuthorize("hasAuthority('SCOPE_messages')")
public Flux<Message> getMessages(...) {}
```

Extracting Authorities Manually

However, there are a number of circumstances where this default is insufficient. For example, some authorization servers don't use the `scope` attribute, but instead have their own custom attribute. Or, at other times, the resource server may need to adapt the attribute or a composition of attributes into internalized authorities.

To this end, the DSL exposes `jwtAuthenticationConverter()`:

```

@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        .authorizeExchange()
            .anyExchange().authenticated()
            .and()
        .oauth2ResourceServer()
            .jwt()
                .jwtAuthenticationConverter(grantedAuthoritiesExtractor());
    return http.build();
}

Converter<Jwt, Mono<AbstractAuthenticationToken>> grantedAuthoritiesExtractor() {
    JwtAuthenticationConverter jwtAuthenticationConverter =
        new JwtAuthenticationConverter();
    jwtAuthenticationConverter.setJwtGrantedAuthoritiesConverter
        (new GrantedAuthoritiesExtractor());
    return new ReactiveJwtAuthenticationConverterAdapter(jwtAuthenticationConverter);
}

```

which is responsible for converting a `Jwt` into an `Authentication`. As part of its configuration, we can supply a subsidiary converter to go from `Jwt` to a `Collection` of granted authorities.

That final converter might be something like `GrantedAuthoritiesExtractor` below:

```

static class GrantedAuthoritiesExtractor
    implements Converter<Jwt, Collection<GrantedAuthority>> {

    public Collection<GrantedAuthority> convert(Jwt jwt) {
        Collection<String> authorities = (Collection<String>)
            jwt.getClaims().get("mycustomclaim");

        return authorities.stream()
            .map(SimpleGrantedAuthority::new)
            .collect(Collectors.toList());
    }
}

```

For more flexibility, the DSL supports entirely replacing the converter with any class that implements `Converter<Jwt, Mono<AbstractAuthenticationToken>>`:

```

static class CustomAuthenticationConverter implements Converter<Jwt, Mono<AbstractAuthenticationToken>>
{
    public AbstractAuthenticationToken convert(Jwt jwt) {
        return Mono.just(jwt).map(this::doConversion);
    }
}

```

Configuring Validation

Using [minimal Spring Boot configuration](#), indicating the authorization server's issuer uri, Resource Server will default to verifying the `iss` claim as well as the `exp` and `nbf` timestamp claims.

In circumstances where validation needs to be customized, Resource Server ships with two standard validators and also accepts custom `OAuth2TokenValidator` instances.

Customizing Timestamp Validation

JWT's typically have a window of validity, with the start of the window indicated in the `nbf` claim and the end indicated in the `exp` claim.

However, every server can experience clock drift, which can cause tokens to appear expired to one server, but not to another. This can cause some implementation heartburn as the number of collaborating servers increases in a distributed system.

Resource Server uses `JwtTimestampValidator` to verify a token's validity window, and it can be configured with a `clockSkew` to alleviate the above problem:

```
@Bean
ReactiveJwtDecoder jwtDecoder() {
    NimbusReactiveJwtDecoder jwtDecoder = (NimbusReactiveJwtDecoder)
        ReactiveJwtDecoders.fromIssuerLocation(issuerUri);

    OAuth2TokenValidator<Jwt> withClockSkew = new DelegatingOAuth2TokenValidator<>(
        new JwtTimestampValidator(Duration.ofSeconds(60)),
        new IssuerValidator(issuerUri));

    jwtDecoder.setJwtValidator(withClockSkew);

    return jwtDecoder;
}
```

Note

By default, Resource Server configures a clock skew of 30 seconds.

Configuring a Custom Validator

Adding a check for the `aud` claim is simple with the `OAuth2TokenValidator` API:

```
public class AudienceValidator implements OAuth2TokenValidator<Jwt> {
    OAuth2Error error = new OAuth2Error("invalid_token", "The required audience is missing", null);

    public OAuth2TokenValidatorResult validate(Jwt jwt) {
        if (jwt.getAudience().contains("messaging")) {
            return OAuth2TokenValidatorResult.success();
        } else {
            return OAuth2TokenValidatorResult.failure(error);
        }
    }
}
```

Then, to add into a resource server, it's a matter of specifying the `ReactiveJwtDecoder` instance:

```
@Bean
ReactiveJwtDecoder jwtDecoder() {
    NimbusReactiveJwtDecoder jwtDecoder = (NimbusReactiveJwtDecoder)
        ReactiveJwtDecoders.fromIssuerLocation(issuerUri);

    OAuth2TokenValidator<Jwt> audienceValidator = new AudienceValidator();
    OAuth2TokenValidator<Jwt> withIssuer = JwtValidators.createDefaultWithIssuer(issuerUri);
    OAuth2TokenValidator<Jwt> withAudience = new DelegatingOAuth2TokenValidator<>(withIssuer,
        audienceValidator);

    jwtDecoder.setJwtValidator(withAudience);

    return jwtDecoder;
}
```

Minimal Configuration for Introspection

Typically, an opaque token can be verified via an [OAuth 2.0 Introspection Endpoint](#), hosted by the authorization server. This can be handy when revocation is a requirement.

When using [Spring Boot](#), configuring an application as a resource server that uses introspection consists of two basic steps. First, include the needed dependencies and second, indicate the introspection endpoint details.

Specifying the Authorization Server

To specify where the introspection endpoint is, simply do:

```
security:
  oauth2:
    resourceserver:
      opaque-token:
        introspection-uri: https://idp.example.com/introspect
        client-id: client
        client-secret: secret
```

Where <https://idp.example.com/introspect> is the introspection endpoint hosted by your authorization server and `client-id` and `client-secret` are the credentials needed to hit that endpoint.

Resource Server will use these properties to further self-configure and subsequently validate incoming JWTs.

Note

When using introspection, the authorization server's word is the law. If the authorization server responds that the token is valid, then it is.

And that's it!

Startup Expectations

When this property and these dependencies are used, Resource Server will automatically configure itself to validate Opaque Bearer Tokens.

This startup process is quite a bit simpler than for JWTs since no endpoints need to be discovered and no additional validation rules get added.

Runtime Expectations

Once the application is started up, Resource Server will attempt to process any request containing an `Authorization: Bearer` header:

```
GET / HTTP/1.1
Authorization: Bearer some-token-value # Resource Server will process this
```

So long as this scheme is indicated, Resource Server will attempt to process the request according to the Bearer Token specification.

Given an Opaque Token, Resource Server will

1. Query the provided introspection endpoint using the provided credentials and the token
2. Inspect the response for an `{ 'active' : true }` attribute
3. Map each scope to an authority with the prefix `SCOPE_`

The resulting `Authentication#getPrincipal`, by default, is a Spring Security [OAuth2AuthenticatedPrincipal](#) object, and `Authentication#getName` maps to the token's `sub` property, if one is present.

From here, you may want to jump to:

- [Looking Up Attributes Post-Authentication](#)
- [Extracting Authorities Manually](#)
- [Using Introspection with JWTs](#)

Looking Up Attributes Post-Authentication

Once a token is authenticated, an instance of `BearerTokenAuthentication` is set in the `SecurityContext`.

This means that it's available in `@Controller` methods when using `@EnableWebFlux` in your configuration:

```
@GetMapping("/foo")
public Mono<String> foo(BearerTokenAuthentication authentication) {
    return Mono.just(authentication.getTokenAttributes().get("sub") + " is the subject");
}
```

Since `BearerTokenAuthentication` holds an `OAuth2AuthenticatedPrincipal`, that also means that it's available to controller methods, too:

```
@GetMapping("/foo")
public Mono<String> foo(@AuthenticationPrincipal OAuth2AuthenticatedPrincipal principal) {
    return Mono.just(principal.getAttribute("sub") + " is the subject");
}
```

Looking Up Attributes Via SpEL

Of course, this also means that attributes can be accessed via SpEL.

For example, if using `@EnableReactiveMethodSecurity` so that you can use `@PreAuthorize` annotations, you can do:

```
@PreAuthorize("principal?.attributes['sub'] == 'foo'")
public Mono<String> forFoosEyesOnly() {
    return Mono.just("foo");
}
```

Overriding or Replacing Boot Auto Configuration

There are two `@Beans` that Spring Boot generates on Resource Server's behalf.

The first is a `SecurityWebFilterChain` that configures the app as a resource server. When use Opaque Token, this `SecurityWebFilterChain` looks like:

```
@Bean
SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
    http
        .authorizeExchange()
            .anyExchange().authenticated()
            .and()
            .oauth2ResourceServer(ServerHttpSecurity.OAuth2ResourceServerSpec::opaqueToken)
        .return http.build();
}
```

If the application doesn't expose a `SecurityWebFilterChain` bean, then Spring Boot will expose the above default one.

Replacing this is as simple as exposing the bean within the application:

```
@EnableWebFluxSecurity
public class MyCustomSecurityConfiguration {
    @Bean
    SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
        http
            .authorizeExchange()
                .pathMatchers("/messages/**").hasAuthority("SCOPE_message:read")
                .anyExchange().authenticated()
                .and()
            .oauth2ResourceServer()
                .opaqueToken()
                    .introspector(myIntrospector());
        return http.build();
    }
}
```

The above requires the scope of `message:read` for any URL that starts with `/messages/`.

Methods on the `oauth2ResourceServer` DSL will also override or replace auto configuration.

For example, the second `@Bean` Spring Boot creates is a `ReactiveOpaqueTokenIntrospector`, which decodes `String` tokens into validated instances of `OAuth2AuthenticatedPrincipal`:

```
@Bean
public ReactiveOpaqueTokenIntrospector introspector() {
    return new NimbusReactiveOpaqueTokenIntrospector(introspectionUri, clientId, clientSecret);
}
```

If the application doesn't expose a `ReactiveOpaqueTokenIntrospector` bean, then Spring Boot will expose the above default one.

And its configuration can be overridden using `introspectionUri()` and `introspectionClientCredentials()` or replaced using `introspector()`.

Using `introspectionUri()`

An authorization server's Introspection Uri can be configured [as a configuration property](#) or it can be supplied in the DSL:

```
@EnableWebFluxSecurity
public class DirectlyConfiguredIntrospectionUri {
    @Bean
    SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
        http
            .authorizeExchange()
                .anyExchange().authenticated()
                .and()
            .oauth2ResourceServer()
                .opaqueToken()
                    .introspectionUri("https://idp.example.com/introspect")
                    .introspectionClientCredentials("client", "secret");
        return http.build();
    }
}
```

Using `introspectionUri()` takes precedence over any configuration property.

Using `introspector()`

More powerful than `introspectionUri()` is `introspector()`, which will completely replace any Boot auto configuration of `ReactiveOpaqueTokenIntrospector`:

```

@EnableWebFluxSecurity
public class DirectlyConfiguredIntrospector {
    @Bean
    SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
        http
            .authorizeExchange()
                .anyExchange().authenticated()
                .and()
            .oauth2ResourceServer()
                .opaqueToken()
                    .introspector(myCustomIntrospector());
        return http.build();
    }
}

```

This is handy when deeper configuration, like [authority mapping](#) or [JWT revocation](#) is necessary.

Exposing a ReactiveOpaqueTokenIntrospector @Bean

Or, exposing a `ReactiveOpaqueTokenIntrospector @Bean` has the same effect as `introspector()`:

```

@Bean
public ReactiveOpaqueTokenIntrospector introspector() {
    return new NimbusOpaqueTokenIntrospector(introspectionUri, clientId, clientSecret);
}

```

Configuring Authorization

An OAuth 2.0 Introspection endpoint will typically return a `scope` attribute, indicating the scopes (or authorities) it's been granted, for example:

```
{ ..., "scope" : "messages contacts" }
```

When this is the case, Resource Server will attempt to coerce these scopes into a list of granted authorities, prefixing each scope with the string "SCOPE_".

This means that to protect an endpoint or method with a scope derived from an Opaque Token, the corresponding expressions should include this prefix:

```

@EnableWebFluxSecurity
public class MappedAuthorities {
    @Bean
    SecurityWebFilterChain springSecurityFilterChain(ServerHttpSecurity http) {
        http
            .authorizeExchange(exchange -> exchange
                .pathMatchers("/contacts/**").hasAuthority("SCOPE_contacts")
                .pathMatchers("/messages/**").hasAuthority("SCOPE_messages")
                .anyExchange().authenticated()
            )
            .oauth2ResourceServer(ServerHttpSecurity.OAuth2ResourceServerSpec::opaqueToken);
        return http.build();
    }
}

```

Or similarly with method security:

```

@PreAuthorize("hasAuthority('SCOPE_messages')")
public Flux<Message> getMessages(...) {}

```

Extracting Authorities Manually

By default, Opaque Token support will extract the scope claim from an introspection response and parse it into individual `GrantedAuthority` instances.

For example, if the introspection response were:

```
{
  "active" : true,
  "scope" : "message:read message:write"
}
```

Then Resource Server would generate an Authentication with two authorities, one for `message:read` and the other for `message:write`.

This can, of course, be customized using a custom `ReactiveOpaqueTokenIntrospector` that takes a look at the attribute set and converts in its own way:

```
public class CustomAuthoritiesOpaqueTokenIntrospector implements ReactiveOpaqueTokenIntrospector {
    private ReactiveOpaqueTokenIntrospector delegate =
        new NimbusReactiveOpaqueTokenIntrospector("https://idp.example.org/
introspect", "client", "secret");

    public Mono<OAuth2AuthenticatedPrincipal> introspect(String token) {
        return this.delegate.introspect(token)
            .map(principal -> new DefaultOAuth2AuthenticatedPrincipal(
                principal.getName(), principal.getAttributes(), extractAuthorities(principal)));
    }

    private Collection<GrantedAuthority> extractAuthorities(OAuth2AuthenticatedPrincipal principal) {
        List<String> scopes = principal.getAttribute(OAuth2IntrospectionClaimNames.SCOPE);
        return scopes.stream()
            .map(SimpleGrantedAuthority::new)
            .collect(Collectors.toList());
    }
}
```

Thereafter, this custom introspector can be configured simply by exposing it as a `@Bean`:

```
@Bean
public ReactiveOpaqueTokenIntrospector introspector() {
    return new CustomAuthoritiesOpaqueTokenIntrospector();
}
```

Using Introspection with JWTs

A common question is whether or not introspection is compatible with JWTs. Spring Security's Opaque Token support has been designed to not care about the format of the token — it will gladly pass any token to the introspection endpoint provided.

So, let's say that you've got a requirement that requires you to check with the authorization server on each request, in case the JWT has been revoked.

Even though you are using the JWT format for the token, your validation method is introspection, meaning you'd want to do:

```
spring:
  security:
    oauth2:
      resourceserver:
        opaque-token:
          introspection-uri: https://idp.example.org/introspection
          client-id: client
          client-secret: secret
```

In this case, the resulting Authentication would be `BearerTokenAuthentication`. Any attributes in the corresponding `OAuth2AuthenticatedPrincipal` would be whatever was returned by the introspection endpoint.

But, let's say that, oddly enough, the introspection endpoint only returns whether or not the token is active. Now what?

In this case, you can create a custom `ReactiveOpaqueTokenIntrospector` that still hits the endpoint, but then updates the returned principal to have the JWT's claims as the attributes:

```
public class JwtOpaqueTokenIntrospector implements ReactiveOpaqueTokenIntrospector {
    private ReactiveOpaqueTokenIntrospector delegate =
        new NimbusReactiveOpaqueTokenIntrospector("https://idp.example.org/
introspect", "client", "secret");
    private ReactiveJwtDecoder jwtDecoder = new NimbusReactiveJwtDecoder(new ParseOnlyJWTProcessor());

    public Mono<OAuth2AuthenticatedPrincipal> introspect(String token) {
        return this.delegate.introspect(token)
            .flatMap(principal -> this.jwtDecoder.decode(token))
            .map(jwt -> new DefaultOAuth2AuthenticatedPrincipal(jwt.getClaims(), NO_AUTHORITIES));
    }

    private static class ParseOnlyJWTProcessor implements Converter<JWT, Mono<JWTClaimsSet>> {
        public Mono<JWTClaimsSet> convert(JWT jwt) {
            try {
                return Mono.just(jwt.getJWTClaimsSet());
            } catch (Exception e) {
                return Mono.error(e);
            }
        }
    }
}
```

Thereafter, this custom introspector can be configured simply by exposing it as a `@Bean`:

```
@Bean
public ReactiveOpaqueTokenIntrospector introspector() {
    return new JwtOpaqueTokenIntrospector();
}
```

Calling a `/userinfo` Endpoint

Generally speaking, a Resource Server doesn't care about the underlying user, but instead about the authorities that have been granted.

That said, at times it can be valuable to tie the authorization statement back to a user.

If an application is also using `spring-security-oauth2-client`, having set up the appropriate `ClientRegistrationRepository`, then this is quite simple with a custom `OpaqueTokenIntrospector`. This implementation below does three things:

- Delegates to the introspection endpoint, to affirm the token's validity
- Looks up the appropriate client registration associated with the `/userinfo` endpoint
- Invokes and returns the response from the `/userinfo` endpoint

```

public class UserInfoOpaqueTokenIntrospector implements ReactiveOpaqueTokenIntrospector {
    private final ReactiveOpaqueTokenIntrospector delegate =
        new NimbusReactiveOpaqueTokenIntrospector("https://idp.example.org/
introspect", "client", "secret");
    private final ReactiveOAuth2UserService<OAuth2UserRequest, OAuth2User> oauth2UserService =
        new DefaultReactiveOAuth2UserService();

    private final ReactiveClientRegistrationRepository repository;

    // ... constructor

    @Override
    public Mono<OAuth2AuthenticatedPrincipal> introspect(String token) {
        return
Mono.zip(this.delegate.introspect(token), this.repository.findByRegistrationId("registration-id"))
        .map(t -> {
            OAuth2AuthenticatedPrincipal authorized = t.getT1();
            ClientRegistration clientRegistration = t.getT2();
            Instant issuedAt = authorized.getAttribute(ISSUED_AT);
            Instant expiresAt =
authorized.getAttribute(OAuth2IntrospectionClaimNames.EXPIRES_AT);
            OAuth2AccessToken accessToken = new OAuth2AccessToken(BEARER, token, issuedAt,
expiresAt);
            return new OAuth2UserRequest(clientRegistration, accessToken);
        })
        .flatMap(this.oauth2UserService::loadUser);
    }
}

```

If you aren't using `spring-security-oauth2-client`, it's still quite simple. You will simply need to invoke the `/userinfo` with your own instance of `WebClient`:

```

public class UserInfoOpaqueTokenIntrospector implements ReactiveOpaqueTokenIntrospector {
    private final ReactiveOpaqueTokenIntrospector delegate =
        new NimbusReactiveOpaqueTokenIntrospector("https://idp.example.org/
introspect", "client", "secret");
    private final WebClient rest = WebClient.create();

    @Override
    public Mono<OAuth2AuthenticatedPrincipal> introspect(String token) {
        return this.delegate.introspect(token)
            .map(this::makeUserInfoRequest);
    }
}

```

Either way, having created your `ReactiveOpaqueTokenIntrospector`, you should publish it as a `@Bean` to override the defaults:

```

@Bean
ReactiveOpaqueTokenIntrospector introspector() {
    return new UserInfoOpaqueTokenIntrospector(...);
}

```

Bearer Token Propagation

Now that you're in possession of a bearer token, it might be handy to pass that to downstream services. This is quite simple with [ServerBearerExchangeFilterFunction](#), which you can see in the following example:

```

@Bean
public WebClient rest() {
    return WebClient.builder()
        .filter(new ServerBearerExchangeFilterFunction())
        .build();
}

```

When the above `WebClient` is used to perform requests, Spring Security will look up the current Authentication and extract any [AbstractOAuth2Token](#) credential. Then, it will propagate that token in the Authorization header.

For example:

```
this.rest.get()
    .uri("https://other-service.example.com/endpoint")
    .retrieve()
    .bodyToMono(String.class)
```

Will invoke the <https://other-service.example.com/endpoint>, adding the bearer token Authorization header for you.

In places where you need to override this behavior, it's a simple matter of supplying the header yourself, like so:

```
this.rest.get()
    .uri("https://other-service.example.com/endpoint")
    .headers(headers -> headers.setBearerAuth(overridingToken))
    .retrieve()
    .bodyToMono(String.class)
```

In this case, the filter will fall back and simply forward the request onto the rest of the web filter chain.

Note

Unlike the [OAuth 2.0 Client filter function](#), this filter function makes no attempt to renew the token, should it be expired. To obtain this level of support, please use the OAuth 2.0 Client filter.

24. @RegisteredOAuth2AuthorizedClient

Spring Security allows resolving an access token using `@RegisteredOAuth2AuthorizedClient`.

Note

A working example can be found in [OAuth 2.0 WebClient WebFlux sample](#).

After configuring Spring Security for [OAuth2 Login](#) or as an [OAuth2 Client](#), an `OAuth2AuthorizedClient` can be resolved using the following:

```
@GetMapping("/explicit")
Mono<String> explicit(@RegisteredOAuth2AuthorizedClient("client-id") OAuth2AuthorizedClient
    authorizedClient) {
    // ...
}
```

This integrates into Spring Security to provide the following features:

- Spring Security will automatically refresh expired tokens (if a refresh token is present)
- If an access token is requested and not present, Spring Security will automatically request the access token.
 - For `authorization_code` this involves performing the redirect and then replaying the original request
 - For `client_credentials` the token is simply requested and saved

If the user authenticated using `oauth2Login()`, then the `client-id` is optional. For example, the following would work:

```
@GetMapping("/implicit")
Mono<String> implicit(@RegisteredOAuth2AuthorizedClient OAuth2AuthorizedClient authorizedClient) {
    // ...
}
```

This is convenient if the user always authenticates with OAuth2 Login and an access token from the same authorization server is needed.

25. Reactive X.509 Authentication

Similar to [Servlet X.509 authentication](#), reactive x509 authentication filter allows extracting an authentication token from a certificate provided by a client.

Below is an example of a reactive x509 security configuration:

```
@Bean
public SecurityWebFilterChain securityWebFilterChain(ServerHttpSecurity http) {
    http
        .x509(withDefaults())
        .authorizeExchange(exchanges ->
            exchanges
                .anyExchange().permitAll()
        );
    return http.build();
}
```

In the configuration above, when neither `principalExtractor` nor `authenticationManager` is provided defaults will be used. The default principal extractor is `SubjectDnX509PrincipalExtractor` which extracts the CN (common name) field from a certificate provided by a client. The default authentication manager is `ReactivePreAuthenticatedAuthenticationManager` which performs user account validation, checking that user account with a name extracted by `principalExtractor` exists and it is not locked, disabled, or expired.

The next example demonstrates how these defaults can be overridden.

```
@Bean
public SecurityWebFilterChain securityWebFilterChain(ServerHttpSecurity http) {
    SubjectDnX509PrincipalExtractor principalExtractor =
        new SubjectDnX509PrincipalExtractor();

    principalExtractor.setSubjectDnRegex("OU=(.*?)(?:,|/|$)");

    ReactiveAuthenticationManager authenticationManager = authentication -> {
        authentication.setAuthenticated("Trusted Org Unit".equals(authentication.getName()));
        return Mono.just(authentication);
    };

    http
        .x509(x509 ->
            x509
                .principalExtractor(principalExtractor)
                .authenticationManager(authenticationManager)
            )
        .authorizeExchange(exchanges ->
            exchanges
                .anyExchange().authenticated()
        );
    return http.build();
}
```

In this example, a username is extracted from the OU field of a client certificate instead of CN, and account lookup using `ReactiveUserDetailsService` is not performed at all. Instead, if the provided certificate issued to an OU named "Trusted Org Unit", a request will be authenticated.

For an example of configuring Netty and WebClient or curl command-line tool to use mutual TLS and enable X.509 authentication, please refer to <https://github.com/spring-projects/spring-security/tree/master/samples/boot/webflux-x509>.

26. WebClient

Note

The following documentation is for use within Reactive environments. For Servlet environments, refer to [WebClient for Servlet](#) environments.

Spring Framework has built in support for setting a Bearer token.

```
webClient.get()
    .headers(h -> h.setBearerAuth(token))
    ...
```

Spring Security builds on this support to provide additional benefits:

- Spring Security will automatically refresh expired tokens (if a refresh token is present)
- If an access token is requested and not present, Spring Security will automatically request the access token.
 - For `authorization_code` this involves performing the redirect and then replaying the original request
 - For `client_credentials` the token is simply requested and saved
- Support for the ability to transparently include the current OAuth token or explicitly select which token should be used.

26.1 WebClient OAuth2 Setup

The first step is ensuring to setup the `WebClient` correctly. An example of setting up `WebClient` in a fully reactive environment can be found below:

```
@Bean
WebClient webClient(ReactiveClientRegistrationRepository clientRegistrations,
    ServerOAuth2AuthorizedClientRepository authorizedClients) {
    ServerOAuth2AuthorizedClientExchangeFilterFunction oauth =
        new ServerOAuth2AuthorizedClientExchangeFilterFunction(clientRegistrations,
            authorizedClients);
    // (optional) explicitly opt into using the oauth2Login to provide an access token implicitly
    // oauth.setDefaultOAuth2AuthorizedClient(true);
    // (optional) set a default ClientRegistration.registrationId
    // oauth.setDefaultClientRegistrationId("client-registration-id");
    return WebClient.builder()
        .filter(oauth)
        .build();
}
```

26.2 Implicit OAuth2AuthorizedClient

If we set `defaultOAuth2AuthorizedClient` to `true` in our setup and the user authenticated with `oauth2Login` (i.e. OIDC), then the current authentication is used to automatically provide the access token. Alternatively, if we set `defaultClientRegistrationId` to a valid `ClientRegistration` id, that registration is used to provide the access token. This is convenient, but in environments where not all endpoints should get the access token, it is dangerous (you might provide the wrong access token to an endpoint).

```

Mono<String> body = this.webClient
    .get()
    .uri(this.uri)
    .retrieve()
    .bodyToMono(String.class);

```

26.3 Explicit OAuth2AuthorizedClient

The `OAuth2AuthorizedClient` can be explicitly provided by setting it on the requests attributes. In the example below we resolve the `OAuth2AuthorizedClient` using Spring WebFlux or Spring MVC argument resolver support. However, it does not matter how the `OAuth2AuthorizedClient` is resolved.

```

@GetMapping("/explicit")
Mono<String> explicit(@RegisteredOAuth2AuthorizedClient("client-id") OAuth2AuthorizedClient
authorizedClient) {
    return this.webClient
        .get()
        .uri(this.uri)
        .attributes(oauth2AuthorizedClient(authorizedClient))
        .retrieve()
        .bodyToMono(String.class);
}

```

26.4 clientRegistrationId

Alternatively, it is possible to specify the `clientRegistrationId` on the request attributes and the `WebClient` will attempt to lookup the `OAuth2AuthorizedClient`. If it is not found, one will automatically be acquired.

```

Mono<String> body = this.webClient
    .get()
    .uri(this.uri)
    .attributes(clientRegistrationId("client-id"))
    .retrieve()
    .bodyToMono(String.class);

```


27. EnableReactiveMethodSecurity

Spring Security supports method security using [Reactor's Context](#) which is setup using `ReactiveSecurityContextHolder`. For example, this demonstrates how to retrieve the currently logged in user's message.

Note

For this to work the return type of the method must be a `org.reactivestreams.Publisher` (i.e. `Mono/Flux`). This is necessary to integrate with Reactor's Context.

```
Authentication authentication = new TestingAuthenticationToken("user", "password", "ROLE_USER");

Mono<String> messageByUsername = ReactiveSecurityContextHolder.getContext()
    .map(SecurityContext::getAuthentication)
    .map(Authentication::getName)
    .flatMap(this::findMessageByUsername)
    // In a WebFlux application the `subscriberContext` is automatically setup using
    `ReactorContextWebFilter`
    .subscriberContext(ReactiveSecurityContextHolder.withAuthentication(authentication));

StepVerifier.create(messageByUsername)
    .expectNext("Hi user")
    .verifyComplete();
```

with `this::findMessageByUsername` defined as:

```
Mono<String> findMessageByUsername(String username) {
    return Mono.just("Hi " + username);
}
```

Below is a minimal method security configuration when using method security in reactive applications.

```
@EnableReactiveMethodSecurity
public class SecurityConfig {
    @Bean
    public MapReactiveUserDetailsService userDetailsService() {
        User.UserBuilder userBuilder = User.withDefaultPasswordEncoder();
        UserDetails rob = userBuilder.username("rob")
            .password("rob")
            .roles("USER")
            .build();
        UserDetails admin = userBuilder.username("admin")
            .password("admin")
            .roles("USER", "ADMIN")
            .build();
        return new MapReactiveUserDetailsService(rob, admin);
    }
}
```

Consider the following class:

```
@Component
public class HelloWorldMessageService {
    @PreAuthorize("hasRole('ADMIN')")
    public Mono<String> findMessage() {
        return Mono.just("Hello World!");
    }
}
```

Combined with our configuration above, `@PreAuthorize("hasRole('ADMIN')")` will ensure that `findByMessage` is only invoked by a user with the role `ADMIN`. It is important to note that any of the

expressions in standard method security work for `@EnableReactiveMethodSecurity`. However, at this time we only support return type of `Boolean` or `boolean` of the expression. This means that the expression must not block.

When integrating with Chapter 21, *WebFlux Security*, the Reactor Context is automatically established by Spring Security according to the authenticated user.

```
@EnableWebFluxSecurity
@EnableReactiveMethodSecurity
public class SecurityConfig {

    @Bean
    SecurityWebFilterChain springWebFilterChain(ServerHttpSecurity http) throws Exception {
        return http
            // Demonstrate that method security works
            // Best practice to use both for defense in depth
            .authorizeExchange(exchanges ->
                exchanges
                    .anyExchange().permitAll()
            )
            .httpBasic(withDefaults())
            .build();
    }

    @Bean
    MapReactiveUserDetailsService userDetailsService() {
        User.UserBuilder userBuilder = User.withDefaultPasswordEncoder();
        UserDetails rob = userBuilder.username("rob")
            .password("rob")
            .roles("USER")
            .build();
        UserDetails admin = userBuilder.username("admin")
            .password("admin")
            .roles("USER", "ADMIN")
            .build();
        return new MapReactiveUserDetailsService(rob, admin);
    }
}
```

You can find a complete sample in [hellowebflux-method](#)

28. Reactive Test Support

28.1 Testing Reactive Method Security

For example, we can test our example from Chapter 27, *EnableReactiveMethodSecurity* using the same setup and annotations we did in Section 18.1, “Testing Method Security”. Here is a minimal sample of what we can do:

```
@RunWith(SpringRunner.class)
@ContextConfiguration(classes = HelloWebfluxMethodApplication.class)
public class HelloWorldMessageServiceTests {
    @Autowired
    HelloWorldMessageService messages;

    @Test
    public void messagesWhenNotAuthenticatedThenDenied() {
        StepVerifier.create(this.messages.findMessage())
            .expectError(AccessDeniedException.class)
            .verify();
    }

    @Test
    @WithMockUser
    public void messagesWhenUserThenDenied() {
        StepVerifier.create(this.messages.findMessage())
            .expectError(AccessDeniedException.class)
            .verify();
    }

    @Test
    @WithMockUser(roles = "ADMIN")
    public void messagesWhenAdminThenOk() {
        StepVerifier.create(this.messages.findMessage())
            .expectNext("Hello World!")
            .verifyComplete();
    }
}
```

28.2 WebTestClientSupport

Spring Security provides integration with `WebTestClient`. The basic setup looks like this:

```
@RunWith(SpringRunner.class)
@ContextConfiguration(classes = HelloWebfluxMethodApplication.class)
public class HelloWebfluxMethodApplicationTests {
    @Autowired
    ApplicationContext context;

    WebTestClient rest;

    @Before
    public void setup() {
        this.rest = WebTestClient
            .bindToApplicationContext(this.context)
            // add Spring Security test Support
            .apply(springSecurity())
            .configureClient()
            .filter(basicAuthentication())
            .build();
    }
    // ...
}
```

Authentication

After applying the Spring Security support to `WebTestClient` we can use either annotations or `mutateWith` support. For example:

```
@Test
public void messageWhenNotAuthenticated() throws Exception {
    this.rest
        .get()
        .uri("/message")
        .exchange()
        .expectStatus().isUnauthorized();
}

// --- WithMockUser ---

@Test
@WithMockUser
public void messageWhenWithMockUserThenForbidden() throws Exception {
    this.rest
        .get()
        .uri("/message")
        .exchange()
        .expectStatus().isEqualTo(HttpStatus.FORBIDDEN);
}

@Test
@WithMockUser(roles = "ADMIN")
public void messageWhenWithMockAdminThenOk() throws Exception {
    this.rest
        .get()
        .uri("/message")
        .exchange()
        .expectStatus().isOk()
        .expectBody(String.class).isEqualTo("Hello World!");
}

// --- mutateWith mockUser ---

@Test
public void messageWhenMutateWithMockUserThenForbidden() throws Exception {
    this.rest
        .mutateWith(mockUser())
        .get()
        .uri("/message")
        .exchange()
        .expectStatus().isEqualTo(HttpStatus.FORBIDDEN);
}

@Test
public void messageWhenMutateWithMockAdminThenOk() throws Exception {
    this.rest
        .mutateWith(mockUser().roles("ADMIN"))
        .get()
        .uri("/message")
        .exchange()
        .expectStatus().isOk()
        .expectBody(String.class).isEqualTo("Hello World!");
}
```

CSRF Support

Spring Security also provides support for CSRF testing with `WebTestClient`. For example:

```

this.rest
  // provide a valid CSRF token
  .mutateWith(csrf())
  .post()
  .uri("/login")
  ...

```

Testing Bearer Authentication

In order to make an authorized request on a resource server, you need a bearer token. If your resource server is configured for JWTs, then this would mean that the bearer token needs to be signed and then encoded according to the JWT specification. All of this can be quite daunting, especially when this isn't the focus of your test.

Fortunately, there are a number of simple ways that you can overcome this difficulty and allow your tests to focus on authorization and not on representing bearer tokens. We'll look at two of them now:

`mockJwt()` `WebTestClientConfigurer`

The first way is via a `WebTestClientConfigurer`. The simplest of these would look something like this:

```

client
  .mutateWith(mockJwt()).get().uri("/endpoint").exchange();

```

What this will do is create a mock `JWT`, passing it correctly through any authentication APIs so that it's available for your authorization mechanisms to verify.

By default, the `JWT` that it creates has the following characteristics:

```

{
  "headers" : { "alg" : "none" },
  "claims" : {
    "sub" : "user",
    "scope" : "read"
  }
}

```

And the resulting `JWT`, were it tested, would pass in the following way:

```

assertThat(jwt.getTokenValue()).isEqualTo("token");
assertThat(jwt.getHeaders().get("alg")).isEqualTo("none");
assertThat(jwt.getSubject()).isEqualTo("sub");
GrantedAuthority authority = jwt.getAuthorities().iterator().next();
assertThat(authority.getAuthority()).isEqualTo("read");

```

These values can, of course be configured.

Any headers or claims can be configured with their corresponding methods:

```

client
  .mutateWith(jwt(jwt -> jwt.header("kid", "one")
    .claim("iss", "https://idp.example.org")))
  .get().uri("/endpoint").exchange();

```

```

client
  .mutateWith(jwt(jwt -> jwt.claims(claims -> claims.remove("scope"))))
  .get().uri("/endpoint").exchange();

```

The `scope` and `scp` claims are processed the same way here as they are in a normal bearer token request. However, this can be overridden simply by providing the list of `GrantedAuthority` instances that you need for your test:

```
client
    .mutateWith(jwt().authorities(new SimpleGrantedAuthority("SCOPE_messages")))
    .get().uri("/endpoint").exchange();
```

Or, if you have a custom `Jwt to Collection<GrantedAuthority>` converter, you can also use that to derive the authorities:

```
client
    .mutateWith(jwt().authorities(new MyConverter()))
    .get().uri("/endpoint").exchange();
```

You can also specify a complete `Jwt`, for which [Jwt.Builder](#) comes quite handy:

```
Jwt jwt = Jwt.withTokenValue("token")
    .header("alg", "none")
    .claim("sub", "user")
    .claim("scope", "read");

client
    .mutateWith(jwt(jwt))
    .get().uri("/endpoint").exchange();
```

`authentication()` `WebTestClientConfigurer`

The second way is by using the `authentication()` Mutator. Essentially, you can instantiate your own `JwtAuthenticationToken` and provide it in your test, like so:

```
Jwt jwt = Jwt.withTokenValue("token")
    .header("alg", "none")
    .claim("sub", "user")
    .build();
Collection<GrantedAuthority> authorities = AuthorityUtils.createAuthorityList("SCOPE_read");
JwtAuthenticationToken token = new JwtAuthenticationToken(jwt, authorities);

client
    .mutateWith(authentication(token))
    .get().uri("/endpoint").exchange();
```

Note that as an alternative to these, you can also mock the `ReactiveJwtDecoder` bean itself with a `@MockBean` annotation.

29. RSocket Security

Spring Security's RSocket support relies on a `SocketAcceptorInterceptor`. The main entry point into security is found in the `PayloadSocketAcceptorInterceptor` which adapts the RSocket APIs to allow intercepting a `PayloadExchange` with `PayloadInterceptor` implementations.

You can find a few sample applications that demonstrate the code below:

- Hello RSocket [hellorsocket](#)
- [Spring Flights](#)

29.1 Minimal RSocket Security Configuration

You can find a minimal RSocket Security configuration below:

```
@Configuration
@EnableRSocketSecurity
public class HelloRSocketSecurityConfig {

    @Bean
    public MapReactiveUserDetailsService userDetailsService() {
        UserDetails user = User.withDefaultPasswordEncoder()
            .username("user")
            .password("user")
            .roles("USER")
            .build();
        return new MapReactiveUserDetailsService(user);
    }
}
```

This configuration enables [basic authentication](#) and sets up [rsocket-authorization](#) to require an authenticated user for any request.

29.2 Adding SecuritySocketAcceptorInterceptor

For Spring Security to work we need to apply `SecuritySocketAcceptorInterceptor` to the `ServerRSocketFactory`. This is what connects our `PayloadSocketAcceptorInterceptor` we created with the RSocket infrastructure. In a Spring Boot application this can be done using the following code.

```
@Bean
ServerRSocketFactoryCustomizer springSecurityRSocketSecurity(
    SecuritySocketAcceptorInterceptor interceptor) {
    return builder -> builder.addSocketAcceptorPlugin(interceptor);
}
```

29.3 RSocket Authentication

RSocket authentication is performed with `AuthenticationPayloadInterceptor` which acts as a controller to invoke a `ReactiveAuthenticationManager` instance.

Authentication at Setup vs Request Time

Generally, authentication can occur at setup time and/or request time.

Authentication at setup time makes sense in a few scenarios. A common scenarios is when a single user (i.e. mobile connection) is leveraging an RSocket connection. In this case only a single user is leveraging the connection, so authentication can be done once at connection time.

In a scenario where the RSocket connection is shared it makes sense to send credentials on each request. For example, a web application that connects to an RSocket server as a downstream service would make a single connection that all users leverage. In this case, if the RSocket server needs to perform authorization based on the web application's users credentials per request makes sense.

In some scenarios authentication at setup and per request makes sense. Consider a web application as described previously. If we need to restrict the connection to the web application itself, we can provide a credential with a `SETUP` authority at connection time. Then each user would have different authorities but not the `SETUP` authority. This means that individual users can make requests but not make additional connections.

Basic Authentication

Spring Security has early support for [RSocket's Basic Authentication Metadata Extension](#).

The RSocket receiver can decode the credentials using `BasicAuthenticationPayloadExchangeConverter` which is automatically setup using the `basicAuthentication` portion of the DSL. An explicit configuration can be found below.

```
@Bean
PayloadSocketAcceptorInterceptor rsocketInterceptor(RSocketSecurity rsocket) {
    rsocket
        .authorizePayload(authorize ->
            authorize
                .anyRequest().authenticated()
                .anyExchange().permitAll()
            )
        .basicAuthentication(Customizer.withDefaults());
    return rsocket.build();
}
```

The RSocket sender can send credentials using `BasicAuthenticationEncoder` which can be added to Spring's `RSocketStrategies`.

```
RSocketStrategies.Builder strategies = ...;
strategies.encoder(new BasicAuthenticationEncoder());
```

It can then be used to send a username and password to the receiver in the setup:

```
UsernamePasswordMetadata credentials = new UsernamePasswordMetadata("user", "password");
Mono<RSocketRequester> requester = RSocketRequester.builder()
    .setupMetadata(credentials, UsernamePasswordMetadata.BASIC_AUTHENTICATION_MIME_TYPE)
    .rsocketStrategies(strategies.build())
    .connectTcp(host, port);
```

Alternatively or additionally, a username and password can be sent in a request.

```
Mono<RSocketRequester> requester;
UsernamePasswordMetadata credentials = new UsernamePasswordMetadata("user", "password");

public Mono<AirportLocation> findRadar(String code) {
    return this.requester.flatMap(req ->
        req.route("find.radar.{code}", code)
            .metadata(credentials, UsernamePasswordMetadata.BASIC_AUTHENTICATION_MIME_TYPE)
            .retrieveMono(AirportLocation.class)
    );
}
```


JWT

Spring Security has early support for [RSocket's Bearer Token Authentication Metadata Extension](#). The support comes in the form of authenticating a JWT (determining the JWT is valid) and then using the JWT to make authorization decisions.

The RSocket receiver can decode the credentials using `BearerPayloadExchangeConverter` which is automatically setup using the `jwt` portion of the DSL. An example configuration can be found below:

```
@Bean
PayloadSocketAcceptorInterceptor rsocketInterceptor(RSocketSecurity rsocket) {
    rsocket
        .authorizePayload(authorize ->
            authorize
                .anyRequest().authenticated()
                .anyExchange().permitAll()
            )
        .jwt(Customizer.withDefaults());
    return rsocket.build();
}
```

The configuration above relies on the existence of a `ReactiveJwtDecoder` `@Bean` being present. An example of creating one from the issuer can be found below:

```
@Bean
ReactiveJwtDecoder jwtDecoder() {
    return ReactiveJwtDecoders
        .fromIssuerLocation("https://example.com/auth/realms/demo");
}
```

The RSocket sender does not need to do anything special to send the token because the value is just a simple `String`. For example, the token can be sent at setup time:

```
String token = ...;
Mono<RSocketRequester> requester = RSocketRequester.builder()
    .setupMetadata(token, BearerTokenMetadata.BEARER_AUTHENTICATION_MIME_TYPE)
    .connectTcp(host, port);
```

Alternatively or additionally, the token can be sent in a request.

```
Mono<RSocketRequester> requester;
String token = ...;

public Mono<AirportLocation> findRadar(String code) {
    return this.requester.flatMap(req ->
        req.route("find.radar.{code}", code)
            .metadata(token, BearerTokenMetadata.BEARER_AUTHENTICATION_MIME_TYPE)
            .retrieveMono(AirportLocation.class)
    );
}
```

29.4 RSocket Authorization

RSocket authorization is performed with `AuthorizationPayloadInterceptor` which acts as a controller to invoke a `ReactiveAuthorizationManager` instance. The DSL can be used to setup authorization rules based upon the `PayloadExchange`. An example configuration can be found below:

```
rsocket
  .authorizePayload(authorize ->
    authz
      .setup().hasRole("SETUP") ❶
      .route("fetch.profile.me").authenticated() ❷
      .matcher(payloadExchange -> isMatch(payloadExchange)) ❸
        .hasRole("CUSTOM")
      .route("fetch.profile.{username}") ❹
        .access((authentication, context) -> checkFriends(authentication, context))
      .anyRequest().authenticated() ❺
      .anyExchange().permitAll() ❻
    )
  )
```

- ❶ Setting up a connection requires the authority `ROLE_SETUP`
- ❷ If the route is `fetch.profile.me` authorization only requires the user be authenticated
- ❸ In this rule we setup a custom matcher where authorization requires the user to have the authority `ROLE_CUSTOM`
- ❹ This rule leverages custom authorization. The matcher expresses a variable with the name `username` that is made available in the `context`. A custom authorization rule is exposed in the `checkFriends` method.
- ❺ This rule ensures that request that does not already have a rule will require the user to be authenticated. A request is where the metadata is included. It would not include additional payloads.
- ❻ This rule ensures that any exchange that does not already have a rule is allowed for anyone. In this example, it means that payloads that have no metadata have no authorization rules.

It is important to understand that authorization rules are performed in order. Only the first authorization rule that matches will be invoked.