

REST with Spring ebook

Introduction

Over the course of this material, we'll explore **how to build a RESTful Service using Spring 3**.

We'll illustrate how to **secure this service** using various mechanisms – starting with the form-based login – standard not necessarily well suited for REST, and moving to more RESTful solutions – Basic and Digest Authentication.

Next, we'll dig into **Discoverability and the HATEOAS Constraint**, and we'll touch on more advanced subjects such as implementing **ETags** functionality in Spring and **error handling** for a REST API.

Finally, we'll discuss best practices for how to **evolve and version a RESTful Service** and we'll end with a high level look on **how to test** the API.

Table of Contents

- [I. Bootstrap the basic Web Project with Spring 3](#)
- [II. Build the REST API with Spring 3 and Java Config](#)
- [III. Security for REST – via Login Form](#)
- [IV. Security for REST – via Basic Authentication](#)
- [V. Security for REST – via Digest Authentication](#)
- [VI. Security for REST – Basic and Digest Authentication together](#)
- [VII. REST API Discoverability and HATEOAS](#)
- [VIII. HATEOAS for a Spring REST Service](#)
- [IX. ETags for REST with Spring](#)
- [X. REST Pagination in Spring](#)
- [XI. Error Handling for REST with Spring 3](#)
- [XII. Versioning a REST API](#)
- [XIII. Testing REST with multiple MIME types](#)

I. Bootstrap a basic Web Project with Spring 3

1. Overview

The section will focus on **bootstrapping the initial web application**, discussing how to make the jump from XML to Java without having to completely migrate the entire XML configuration.

2. The Maven *pom.xml*

```

1 <project xmlns="http://maven.apache.org/POM/4.0.0"
2   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
3   xsi:schemaLocation="
4     http://maven.apache.org/POM/4.0.0
5     http://maven.apache.org/xsd/maven-4.0.0.xsd">
6   <modelVersion>4.0.0</modelVersion>
7   <groupId>org</groupId>
8   <artifactId>rest</artifactId>
9   <version>0.0.1-SNAPSHOT</version>
10  <packaging>war</packaging>
11
12  <dependencies>
13
14    <dependency>
15      <groupId>org.springframework</groupId>
16      <artifactId>spring-webmvc</artifactId>
17      <version>${spring.version}</version>
18      <exclusions>
19        <exclusion>
20          <artifactId>commons-logging</artifactId>
21          <groupId>commons-logging</groupId>
22        </exclusion>
23      </exclusions>
24    </dependency>
25    <dependency>
26      <groupId>cglib</groupId>
27      <artifactId>cglib-nodep</artifactId>
28      <version>${cglib.version}</version>
29      <scope>runtime</scope>
30    </dependency>
31
32  </dependencies>
33
34  <build>
35    <finalName>rest</finalName>
36
37    <plugins>
38      <plugin>
39        <groupId>org.apache.maven.plugins</groupId>
40        <artifactId>maven-compiler-plugin</artifactId>
41        <version>3.1</version>
42        <configuration>
43          <source>1.6</source>
44          <target>1.6</target>
45          <encoding>UTF-8</encoding>
46        </configuration>
47      </plugin>
48    </plugins>
49  </build>
50
51  <properties>
52    <spring.version>3.2.2.RELEASE</spring.version>

```



```

53     <cglib.version>2.2.2</cglib.version>
54   </properties>
55
56 </project>

```

2.1. Justification of the *cglib* dependency

You may wonder why *cglib* is a dependency – it turns out there is a valid reason to include it – the entire configuration cannot function without it. If removed, Spring will throw:

Caused by: java.lang.IllegalStateException: CGLIB is required to process @Configuration classes. Either add CGLIB to the classpath or remove the following @Configuration bean definitions

The reason this happens is explained by the way Spring deals with *@Configuration* classes. These classes are effectively beans, and because of this they need to be aware of the Context, and respect scope and other bean semantics. This is achieved by dynamically creating a *cglib* proxy with this awareness for each *@Configuration* class, hence the *cglib* dependency.

Also, because of this, there are a few restrictions for *Configuration* annotated classes:

- Configuration classes **should not be final**
- They should have a constructor with no arguments

2.2. The *cglib* dependency in Spring 3.2

Starting with Spring 3.2, it is **no longer necessary to add *cglib* as an explicit dependency**. This is because Spring is now inlining *cglib* – which will ensure that all class based proxying functionality will work out of the box with Spring 3.2.

The new *cglib* code is placed under the Spring package: *org.springframework.cglib* (replacing the original *net.sf.cglib*). The reason for the package change is to avoid conflicts with any *cglib* versions already existing on the classpath.

Also, the new *cglib* 3.0 is now used, upgraded from the older 2.2 dependency (see this [JIRA issue](#) for more details).

3. The Java based web configuration

```

1  @Configuration
2  @ImportResource( { "classpath:/rest_config.xml" } )
3  @ComponentScan( basePackages = "org.rest" )
4  @PropertySource( { "classpath:rest.properties", "classpath:web.properties" } )
5  public class AppConfig{
6
7      @Bean
8      public static PropertySourcesPlaceholderConfigurer properties() {
9          return new PropertySourcesPlaceholderConfigurer();
10     }
11 }

```

First, the ***@Configuration*** annotation – this is the main artifact used by the Java based Spring configuration; it is itself meta-annotated with *@Component*, which makes the annotated classes **standard beans** and as such, also candidates for component scanning. The main purpose of *@Configuration* classes is to be sources of bean

definitions for the Spring IoC Container. For a more detailed description, see the [official docs](#).

Then, **@ImportResource** is used to import the existing XML based Spring configuration. This may be configuration which is still being migrated from XML to Java, or simply legacy configuration that you wish to keep. Either way, importing it into the Container is essential for a successful migration, allowing small steps without too much risk. The equivalent XML annotation that is replaced is:

```
<import resource="classpath*:./rest_config.xml" />
```

Moving on to **@ComponentScan** – this configures the component scanning directive, effectively replacing the XML:

```
<context:component-scan base-package="org.rest" />
```

As of Spring 3.1, the **@Configuration** are excluded from classpath scanning by default – see [this JIRA issue](#). Before Spring 3.1 though, these classes should have been excluded explicitly:

```
excludeFilters = { @ComponentScan.Filter( Configuration.class ) }
```

The **@Configuration** classes should not be autodiscovered because they are already specified and used by the Container – allowing them to be rediscovered and introduced into the Spring context will result in the following error:

Caused by: org.springframework.context.annotation.ConflictingBeanDefinitionException: Annotation-specified bean name 'webConfig' for bean class [org.rest.spring.AppConfig] conflicts with existing, non-compatible bean definition of same name and class [org.rest.spring.AppConfig]

And finally, using the **@Bean** annotation to configure the **properties support** – **PropertySourcesPlaceholderConfigurer** is initialized in a **@Bean** annotated method, indicating it will produce a Spring bean managed by the Container. This new configuration has replaced the following XML:

```
1 <context:property-placeholder
2 location="classpath:persistence.properties, classpath:web.properties"
3 ignore-unresolvable="true"/>
```

For a more in depth discussion on why it was necessary to manually register the **PropertySourcesPlaceholderConfigurer** bean, see the [Properties with Spring Tutorial](#).

3.1. The *web.xml*

```
1 <?xml version="1.0" encoding="UTF-8"?>
2 <web-app xmlns="
3     http://java.sun.com/xml/ns/javaee"
4     xmlns:web="http://java.sun.com/xml/ns/javaee/web-app_3_0.xsd"
5     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
6     xsi:schemaLocation="
7     http://java.sun.com/xml/ns/javaee
8     http://java.sun.com/xml/ns/javaee/web-app_3_0.xsd"
9     id="rest" version="3.0">
10
11 <context-param>
12 <param-name>contextClass</param-name>
13 <param-value>
14     org.springframework.web.context.support.AnnotationConfigWebApplicationContext
```

```

15     </param-value>
16 </context-param>
17 <context-param>
18     <param-name>contextConfigLocation</param-name>
19     <param-value>org.rest.spring.root</param-value>
20 </context-param>
21 <listener>
22     <listener-class>org.springframework.web.context.ContextLoaderListener</listener-class>
23 </listener>
24
25 <servlet>
26     <servlet-name>rest</servlet-name>
27     <servlet-class>
28         org.springframework.web.servlet.DispatcherServlet
29     </servlet-class>
30     <init-param>
31         <param-name>contextClass</param-name>
32         <param-value>
33             org.springframework.web.context.support.AnnotationConfigWebApplicationContext
34         </param-value>
35     </init-param>
36     <init-param>
37         <param-name>contextConfigLocation</param-name>
38         <param-value>org.rest.spring.rest</param-value>
39     </init-param>
40     <load-on-startup>1</load-on-startup>
41 </servlet>
42 <servlet-mapping>
43     <servlet-name>rest</servlet-name>
44     <url-pattern>/api/*</url-pattern>
45 </servlet-mapping>
46
47 <welcome-file-list>
48     <welcome-file />
49 </welcome-file-list>
50
51 </web-app>

```

First, the root context is defined and configured to use *AnnotationConfigWebApplicationContext* instead of the default *XmlWebApplicationContext*. The newer *AnnotationConfigWebApplicationContext* accepts *@Configuration* annotated classes as input for the Container configuration and is needed in order to set up the Java based context.

Unlike *XmlWebApplicationContext*, it assumes no default configuration class locations, so the “*contextConfigLocation*” *init-param* for the Servlet must be set. This will point to the java package where the *@Configuration* classes are located; the fully qualified name(s) of the classes are also supported.

Next, the *DispatcherServlet* is configured to use the same kind of context, with the only difference that it’s loading configuration classes out of a different package.

Other than this, the *web.xml* doesn’t really change from a XML to a Java based configuration.

4. Conclusion

The presented approach allows for a smooth **migration of the Spring configuration** from XML to Java, mixing the old and the new. This is important for older projects, which may have a lot of XML based configuration that cannot be migrated all at once. This way, the *web.xml* and bootstrapping of the application is the first step in a migration, after which the remaining XML beans can be ported in small increments.

In the meantime, you can check out the [github project](#).

II. Build the REST API with Spring 3 and Java Config

1. Overview

This section shows how to **set up REST in Spring** – the Controller and HTTP response codes, configuration of payload marshalling and content negotiation.

2. Understanding REST in Spring

The Spring framework supports 2 ways of creating RESTful services:

- using MVC with *ModelAndView*
- using HTTP message converters

The ***ModelAndView*** approach is older and much better documented, but also more verbose and configuration heavy. It tries to shoehorn the REST paradigm into the old model, which is not without problems. The Spring team understood this and provided first-class REST support starting with **Spring 3.0**.

The new approach, based on ***HttpMessageConverter*** and **annotations**, is much more lightweight and easy to implement. Configuration is minimal and it provides sensible defaults for what you would expect from a RESTful service. It is however newer and a bit on the light side concerning documentation; what's more, the reference doesn't go out of its way to make the distinction and the tradeoffs between the two approaches as clear as they should be. Nevertheless, this is the way RESTful services should be build after Spring 3.0.

3. The Java configuration

```

1 | @Configuration
2 | @EnableWebMvc
3 | public class WebConfig{
4 |     //
5 | }
```

The new ***@EnableWebMvc*** annotation does a number of useful things – specifically, in the case of REST, it detect the existence of Jackson and JAXB 2 on the classpath and automatically creates and registers default **JSON and XML converters**. The functionality of the annotation is equivalent to the XML version:

```
<mvc:annotation-driven />
```

This is a shortcut, and though it may be useful in many situations, it's not perfect. When more complex configuration is needed, remove the annotation and extend *WebMvcConfigurationSupport* directly.

4. Testing the Spring context

Starting with **Spring 3.1**, [we get](#) first-class testing support for *@Configuration* classes:

```

1  @RunWith( SpringJUnit4ClassRunner.class )
2  @ContextConfiguration( classes = { ApplicationConfig.class, PersistenceConfig.class },
3    loader = AnnotationConfigContextLoader.class )
4  public class SpringTest{
5
6    @Test
7    public void whenSpringContextIsInstantiated_thenNoExceptions(){
8      // When
9    }
10 }

```

The Java configuration classes are simply specified with the `@ContextConfiguration` annotation and the new `AnnotationConfigContextLoader` loads the bean definitions from the `@Configuration` classes.

Notice that the `WebConfig` configuration class was not included in the test because it needs to run in a Servlet context, which is not provided.

5. The Controller

The `@Controller` is the central artifact in the entire Web Tier of the RESTful API. For the purpose of the following examples, the controller is modeling a simple REST resource – `Foo`:

```

1  @Controller
2  @RequestMapping( value = "foo" )
3  class FooController{
4
5    @Autowired
6    IFooService service;
7
8    @RequestMapping( method = RequestMethod.GET )
9    @ResponseBody
10   public List< Foo > getAll(){
11     return service.getAll();
12   }
13
14   @RequestMapping( value =("/{id})", method = RequestMethod.GET )
15   @ResponseBody
16   public Foo get( @PathVariable( "id" ) Long id ){
17     return RestPreconditions.checkNotNull( service.getById( id ) );
18   }
19
20   @RequestMapping( method = RequestMethod.POST )
21   @ResponseStatus( HttpStatus.CREATED )
22   @ResponseBody
23   public Long create( @RequestBody Foo entity ){
24     RestPreconditions.checkNotNullFromRequest( entity );
25     return service.create( entity );
26   }
27
28   @RequestMapping( method = RequestMethod.PUT )
29   @ResponseStatus( HttpStatus.OK )
30   public void update( @RequestBody Foo entity ){
31     RestPreconditions.checkNotNullFromRequest( entity );
32     RestPreconditions.checkNotNull( service.getById( entity.getId() ) );
33     service.update( entity );
34   }
35
36   @RequestMapping( value =("/{id})", method = RequestMethod.DELETE )
37   @ResponseStatus( HttpStatus.OK )
38   public void delete( @PathVariable( "id" ) Long id ){
39     service.deleteById( id );
40   }
41 }
42 }

```

The Controller implementation is **non-public** – this is because there is no need for it to be. Usually the controller

is the last in the chain of dependencies – it receives HTTP requests from the Spring front controller (the *DispatcherServlet*) and simply delegates them forward to a service layer. If there is no use case where the controller has to be injected or manipulated through a direct reference, then I prefer not to declare it as public.

The **request mappings** are straightforward – as with any controller, the actual *value* of the mapping as well as the HTTP *method* are used to determine the target method for the request. **@RequestBody** will bind the parameters of the method to the body of the HTTP request, whereas **@ResponseBody** does the same for the response and return type. They also ensure that the resource will be marshalled and unmarshalled using the correct HTTP converter. **Content negotiation** will take place to choose which one of the active converters will be used, based mostly on the *Accept* header, although other HTTP headers may be used to determine the representation as well.

6. Mapping the HTTP response codes

The status codes of the HTTP response are one of the most important parts of the REST service, and the subject can quickly become very complex. Getting these right can be what makes or breaks the service.

6.1. Unmapped requests

If Spring MVC receives a request which doesn't have a mapping, it considers the request not to be allowed and returns a **405 METHOD NOT ALLOWED** back to the client. It is also good practice to include the **Allow HTTP header** when returning a *405* to the client, in order to specify which operations **are** allowed. This is the standard behavior of Spring MVC and does not require any additional configuration.

6.2. Valid, mapped requests

For any request that does have a mapping, Spring MVC considers the request valid and responds with **200 OK** if no other status code is specified otherwise. It is because of this that controller declares different **@ResponseStatus** for the *create*, *update* and *delete* actions but not for *get*, which should indeed return the default 200 OK.

6.3. Client error

In case of a **client error**, custom exceptions are defined and mapped to the appropriate error codes. Simply throwing these exceptions from any of the layers of the web tier will ensure Spring maps the corresponding status code on the HTTP response.

```

1 | @ResponseStatus( value = HttpStatus.BAD_REQUEST )
2 | public class BadRequestException extends RuntimeException{
3 |     //
4 | }
5 | @ResponseStatus( value = HttpStatus.NOT_FOUND )
6 | public class ResourceNotFoundException extends RuntimeException{
7 |     //
8 | }
```

These exceptions are part of the REST API and, as such, should only be used in the appropriate layers corresponding to REST; if for instance a DAO/DAL layer exist, it should not use the exceptions directly. Note also that these are not **checked exceptions** but **runtime exceptions** – in line with Spring practices and idioms.

6.4. Using `@ExceptionHandler`

Another option to map custom exceptions on specific status codes is to use the `@ExceptionHandler` annotation in the controller. The problem with that approach is that the annotation only applies to the controller in which it is defined, not to the entire Spring Container, which means that it needs to be declared in each controller individually. This quickly becomes cumbersome, especially in more complex applications which many controllers. There are a few **JIRA issues** opened with Spring at this time to handle this and other related limitations: [SPR-8124](#), [SPR-7278](#), [SPR-8406](#).

7. Additional Maven dependencies

In addition to the `spring-webmvc` dependency [required for the standard web application](#), we'll need to set up content marshalling and unmarshalling for the REST API:

```

1  <dependencies>
2    <dependency>
3      <groupId>com.fasterxml.jackson.core</groupId>
4      <artifactId>jackson-databind</artifactId>
5      <version>${jackson.version}</version>
6    </dependency>
7    <dependency>
8      <groupId>javax.xml.bind</groupId>
9      <artifactId>jaxb-api</artifactId>
10     <version>${jaxb-api.version}</version>
11     <scope>runtime</scope>
12   </dependency>
13 </dependencies>
14
15 <properties>
16   <jackson.version>2.2.2</jackson.version>
17   <jaxb-api.version>2.2.9</jaxb-api.version>
18 </properties>

```

These are the libraries used to convert the representation of the REST resource to either **JSON** or **XML**.

8. Conclusion

This section covered the configuration and implementation of a RESTful service using Spring 3.1 and Java based configuration, discussing HTTP response codes, basic content negotiation and marshaling.

In the meantime, check out the [github project](#) – this is an Eclipse based project, so it should be easy to import and run as it is.

III. Security for REST – via Login Form

1. Overview

This section shows how to **Secure a REST Service using Spring and Spring Security 3.1** with Java based configuration. We'll focus on how to set up the Security Configuration specifically for the REST API using a Login and Cookie approach.

2. Spring Security in the *web.xml*

The architecture of Spring Security is based entirely on Servlet Filters and, as such, comes before Spring MVC in regards to the processing of HTTP requests. Keeping this in mind, to begin with, a **filter** needs to be declared in the *web.xml* of the application:

```

1 <filter>
2   <filter-name>springSecurityFilterChain</filter-name>
3   <filter-class>org.springframework.web.filter.DelegatingFilterProxy</filter-class>
4 </filter>
5 <filter-mapping>
6   <filter-name>springSecurityFilterChain</filter-name>
7   <url-pattern>/*</url-pattern>
8 </filter-mapping>

```

The filter must necessarily be named `'springSecurityFilterChain'` to match the default bean created by Spring Security in the container.

Note that the defined filter is not the actual class implementing the security logic but a *DelegatingFilterProxy* with the purpose of delegating the Filter's methods to an internal bean. This is done so that the target bean can still benefit from the Spring context lifecycle and flexibility.

The URL pattern used to configure the Filter is `/*` even though the entire web service is mapped to `/api/*` so that the security configuration has the option to secure other possible mappings as well, if required.

3. The Security Configuration

```

1 <?xml version="1.0" encoding="UTF-8"?>
2 <beans:beans
3   xmlns="http://www.springframework.org/schema/security"
4   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
5   xmlns:beans="http://www.springframework.org/schema/beans"
6   xmlns:sec="http://www.springframework.org/schema/security"
7   xsi:schemaLocation="
8     http://www.springframework.org/schema/security
9     http://www.springframework.org/schema/security/spring-security-3.1.xsd
10    http://www.springframework.org/schema/beans
11    http://www.springframework.org/schema/beans/spring-beans-3.2.xsd">
12
13   <http entry-point-ref="restAuthenticationEntryPoint">
14     <intercept-url pattern="/api/admin/**" access="ROLE_ADMIN"/>
15
16     <form-login authentication-success-handler-ref="mySuccessHandler" />
17
18     <logout />
19   </http>
20
21   <beans:bean id="mySuccessHandler"
22     class="org.rest.security.MySavedRequestAwareAuthenticationSuccessHandler"/>
23
24   <authentication-manager alias="authenticationManager">
25     <authentication-provider>
26       <user-service>
27         <user name="temporary" password="temporary" authorities="ROLE_ADMIN"/>
28         <user name="user" password="user" authorities="ROLE_USER"/>
29       </user-service>
30     </authentication-provider>
31   </authentication-manager>
32
33 </beans:beans>

```

Most of the configuration is done using the **security namespace** – for this to be enabled, the schema locations must be defined and pointed to the correct 3.1 XSD versions. The namespace is designed so that it expresses the common uses of Spring Security while still providing hooks raw beans to accommodate more advanced scenarios.

3.1. The `<http>` element

The `<http>` element is the main container element for HTTP security configuration. In the current implementation, it only secured a single mapping: `/api/admin/**`. Note that the mapping is **relative to the root context** of the web application, not to the `rest` Servlet; this is because the entire security configuration lives in the root Spring context and not in the child context of the Servlet.

3.2. The Entry Point

In a standard web application, the authentication process may be automatically triggered when the client tries to access a secured resource without being authenticated – this is usually done by redirecting to a login page so that the user can enter credentials. However, for a **REST Web Service** this behavior doesn't make much sense – Authentication should only be done by a request to the correct URI and all other requests should simply fail with a **401 UNAUTHORIZED** status code if the user is not authenticated.

Spring Security handles this automatic triggering of the authentication process with the concept of an **Entry Point** – this is a required part of the configuration, and can be injected via the `entry-point-ref` attribute of the `<http>` element. Keeping in mind that this functionality doesn't make sense in the context of the REST Service, the new custom entry point is defined to simply return 401 whenever it is triggered:

```

1 | @Component( "restAuthenticationEntryPoint" )
2 | public class RestAuthenticationEntryPoint implements AuthenticationEntryPoint{
3 |
4 |     @Override
5 |     public void commence( HttpServletRequest request, HttpServletResponse response,
6 |         AuthenticationException authException ) throws IOException{
7 |         response.sendError( HttpServletResponse.SC_UNAUTHORIZED, "Unauthorized" );
8 |     }
9 | }

```

3.3. The Login Form for REST

There are multiple ways to do Authentication for a REST API – one of the default Spring Security provides is **Form Login** – which uses an authentication processing filter – `org.springframework.security.web.authentication.UsernamePasswordAuthenticationFilter`.

The `<form-login>` element will create this filter and will also allow us to set our custom authentication success handler on it. This can also be done manually by using the `<custom-filter>` element to register a filter at the position `FORM_LOGIN_FILTER` – but the namespace support is flexible enough.

Note that for a standard web application, the **auto-config** attribute of the `<http>` element is shorthand syntax for some useful security configuration. While this may be appropriate for some very simple configurations, it doesn't fit and should not be used for a REST API.

3.4. Authentication should return 200 instead of 301

By default, form login will answer a successful authentication request with a **301 MOVED PERMANENTLY** status code; this makes sense in the context of an actual login form which needs to redirect after login. For a RESTful web service however, the desired response for a successful authentication should be **200 OK**.

This is done by injecting a **custom authentication success handler** in the form login filter, to replace the default one. The new handler implements the exact same login as the default *org.springframework.security.web.authentication.SavedRequestAwareAuthenticationSuccessHandler* with one notable difference – the redirect logic is removed:

```

1  public class MySavedRequestAwareAuthenticationSuccessHandler
2      extends SimpleUrlAuthenticationSuccessHandler {
3
4      private RequestCache requestCache = new HttpSessionRequestCache();
5
6      @Override
7      public void onAuthenticationSuccess(HttpServletRequest request, HttpServletResponse response,
8          Authentication authentication) throws ServletException, IOException {
9          SavedRequest savedRequest = requestCache.getRequest(request, response);
10
11         if (savedRequest == null) {
12             clearAuthenticationAttributes(request);
13             return;
14         }
15         String targetUrlParam = getTargetUrlParameter();
16         if (isAlwaysUseDefaultTargetUrl() ||
17             (targetUrlParam != null &&
18              StringUtils.hasText(request.getParameter(targetUrlParam)))) {
19             requestCache.removeRequest(request, response);
20             clearAuthenticationAttributes(request);
21             return;
22         }
23
24         clearAuthenticationAttributes(request);
25     }
26
27     public void setRequestCache(RequestCache requestCache) {
28         this.requestCache = requestCache;
29     }
30 }

```

3.5. The Authentication Manager and Provider

The authentication process uses an **in-memory provider** to perform authentication – this is meant to simplify the configuration as a production implementation of these artifacts is outside the scope of this post.

3.6 Finally – Authentication against the running REST Service

Now let's see how we can authenticate against the REST API – the URL for login is */j_spring_security_check* – and a simple *curl* command performing login would be:

```

1  curl -i -X POST -d j_username=user -d j_password=userPass
2  http://localhost:8080/spring-security-rest/j_spring_security_check

```

This request will return the Cookie which will then be used by any subsequent request against the REST Service.

We can use *curl* to authentication and **store the cookie it receives in a file**:

```

1 | curl -i -X POST -d j_username=user -d j_password=userPass -c /opt/cookies.txt
2 | http://localhost:8080/spring-security-rest/j_spring_security_check

```

Then **we can use the cookie from the file** to do further authenticated requests:

```

1 | curl -i --header "Accept:application/json" -X GET -b /opt/cookies.txt
2 | http://localhost:8080/spring-security-rest/api/foos

```

This authenticated request will correctly **result in a 200 OK**:

```

1 | HTTP/1.1 200 OK
2 | Server: Apache-Coyote/1.1
3 | Content-Type: application/json;charset=UTF-8
4 | Transfer-Encoding: chunked
5 | Date: Wed, 24 Jul 2013 20:31:13 GMT
6 |
7 | [{"id":0,"name":"JbidXc"}]

```

4. Maven and other trouble

The Spring [core dependencies](#) necessary for a web application and for the REST Service have been discussed in detail. For security, we'll need to add: *spring-security-web* and *spring-security-config* – all of these have also been covered in the [Maven for Spring Security](#) tutorial.

It's worth paying close attention to the way Maven will resolve the older Spring dependencies – the resolution strategy will start [causing problems](#) once the security artifacts are added to the pom. To address this problem, some of the core dependencies will need to be overridden in order to keep them at the right version.

5. Conclusion

This section covered the basic security configuration and implementation for a RESTful Service using **Spring Security 3.1**, discussing the *web.xml*, the security configuration, the HTTP status codes for the authentication process and the Maven resolution of the security artifacts.

The implementation of the code presented in this section can be found in [the github project](#) – this is an Eclipse based project, so it should be easy to import and run as it is.

IV. Security for REST – via Basic Authentication

1. Overview

This section shows how to set up, configure and customize **Basic Authentication with Spring**. We're going to built on top of the simple REST API from the previous sections, and secure the application with the Basic Auth mechanism provided by Spring Security.

2. The Spring Security Configuration

The Configuration for Spring Security is still XML:

```

1 | <?xml version="1.0" encoding="UTF-8" ?>

```

```

2 <beans:beans xmlns="http://www.springframework.org/schema/security"
3   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
4   xmlns:beans="http://www.springframework.org/schema/beans"
5   xsi:schemaLocation="
6     http://www.springframework.org/schema/security
7     http://www.springframework.org/schema/security/spring-security-3.1.xsd
8     http://www.springframework.org/schema/beans
9     http://www.springframework.org/schema/beans/spring-beans-3.2.xsd">
10
11 <http use-expressions="true">
12   <intercept-url pattern="/**" access="isAuthenticated()" />
13
14   <http-basic />
15 </http>
16
17 <authentication-manager>
18   <authentication-provider>
19     <user-service>
20       <user name="user1" password="user1Pass" authorities="ROLE_USER" />
21     </user-service>
22   </authentication-provider>
23 </authentication-manager>
24
25 </beans:beans>

```

This is one of the last pieces of configuration in Spring that still need XML – [Java Configuration for Spring Security](#) is still a work in progress.

What is relevant here is the `<http-basic>` element inside the main `<http>` element of the configuration – this is enough to enable Basic Authentication for the entire application. The Authentication Manager is not the focus of this section, so we are using an in memory manager with the user and password defined in plaintext.

The `web.xml` of the web application enabling Spring Security has already been discussed in the previous sections.

3. Consuming The Secured Application

The `curl` command is our go to tool for consuming the secured application.

First, let's try to request a `Foo` resource without providing any security credentials:

```
1 | curl -i http://localhost:8080/spring-security-rest-basic-auth/api/foos/1
```

We get back the expected `401 Unauthorized` and [the Authentication Challenge](#):

```

1 HTTP/1.1 401 Unauthorized
2 Server: Apache-Coyote/1.1
3 Set-Cookie: JSESSIONID=E5A8D3C16B65A0A007CFAACAE6916B; Path=/spring-security-rest-basic-auth/; HttpOnly
4 WWW-Authenticate: Basic realm="Spring Security Application"
5 Content-Type: text/html;charset=utf-8
6 Content-Length: 1061
7 Date: Wed, 29 May 2013 15:14:08 GMT

```

The browser would interpret this challenge and prompt us for credentials with a simple dialog, but since we're using `curl`, this isn't the case.

Now, let's request the same resource – the homepage – but **provide the credentials** to access it as well:

```
1 | curl -i --user user1:user1Pass http://localhost:8080/spring-security-rest-basic-auth/api/foos/1
```

Now, the response from the server is *200 OK* along with a *Cookie*:

```

1 HTTP/1.1 200 OK
2 Server: Apache-Coyote/1.1
3 Set-Cookie: JSESSIONID=301225C7AE7C74B0892887389996785D; Path=/spring-security-rest-basic-auth/; HttpOnly
4 Content-Type: text/html;charset=ISO-8859-1
5 Content-Language: en-US
6 Content-Length: 90
7 Date: Wed, 29 May 2013 15:19:38 GMT

```

From the browser, the application can be consumed normally – the only difference is that a login page is no longer a hard requirement since all browsers support Basic Authentication and use a dialog to prompt the user for credentials.

4. Further Configuration – The Entry Point

By default, the *BasicAuthenticationEntryPoint* provisioned by Spring Security returns a full html page for a *401 Unauthorized* response back to the client. This html representation of the error renders well in a browser, but it not well suited for other scenarios, such as a REST API where a json representation may be preferred.

The namespace is flexible enough for this new requirement as well – to address this – the entry point can be overridden:

```

1 | <http-basic entry-point-ref="myBasicAuthenticationEntryPoint" />

```

The new entry point is defined as a standard bean:

```

1 | @Component
2 | public class MyBasicAuthenticationEntryPoint extends BasicAuthenticationEntryPoint {
3 |
4 |     @Override
5 |     public void commence
6 |         (HttpServletRequest request, HttpServletResponse response, AuthenticationException authEx)
7 |         throws IOException, ServletException {
8 |         response.setHeader("WWW-Authenticate", "Basic realm=\"" + getRealmName() + "\"");
9 |         response.setStatus(HttpServletResponse.SC_UNAUTHORIZED);
10 |         PrintWriter writer = response.getWriter();
11 |         writer.println("HTTP Status 401 - " + authEx.getMessage());
12 |     }
13 |
14 |     @Override
15 |     public void afterPropertiesSet() throws Exception {
16 |         setRealmName("BaeIdung");
17 |         super.afterPropertiesSet();
18 |     }
19 | }

```

By writing directly to the HTTP Response we now have full control over the format of the response body.

5. The Maven Dependencies

The Maven dependencies for Spring Security have been discussed before in the [Spring Security with Maven article](#) – we will need both *spring-security-web* and *spring-security-config* available at runtime.

6. Conclusion

In this example we secured the existing REST application with Spring Security and Basic Authentication. We

discussed the XML configuration and we consumed the application with simple curl commands. Finally took control of the exact error message format – moving from the standard HTML error page to a custom text or json format.

The implementation of this Spring examples can be found in [the github project](#) – this is an Eclipse based project, so it should be easy to import and run as it is.

V. Security for REST – via Digest Authentication

1. Overview

This section illustrates how to set up, configure and customize Digest Authentication with Spring. Similar to the previous section about Basic Authentication, we're going to built on top of the Spring REST Service from the previous sections, and secure the application with the Digest Auth mechanism provided by Spring Security.

2. The Security XML Configuration

First thing to understand about the configuration is that, while Spring Security does have full out of the box support for the Digest authentication mechanism, this support is **not as well integrated into the namespace** as Basic Authentication was.

In this case, we need to manually **define the raw beans** that are going to make up the security configuration – the *DigestAuthenticationFilter* and the *DigestAuthenticationEntryPoint*:

```

1  <?xml version="1.0" encoding="UTF-8"?>
2  <beans:beans xmlns="http://www.springframework.org/schema/security"
3      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
4      xmlns:beans="http://www.springframework.org/schema/beans"
5      xsi:schemaLocation="
6          http://www.springframework.org/schema/security
7          http://www.springframework.org/schema/security/spring-security-3.1.xsd
8          http://www.springframework.org/schema/beans
9          http://www.springframework.org/schema/beans/spring-beans-3.2.xsd">
10
11     <beans:bean id="digestFilter"
12         class="org.springframework.security.web.authentication.www.DigestAuthenticationFilter">
13         <beans:property name="userService" ref="userService" />
14         <beans:property name="authenticationEntryPoint" ref="digestEntryPoint" />
15     </beans:bean>
16     <beans:bean id="digestEntryPoint"
17         class="org.springframework.security.web.authentication.www.DigestAuthenticationEntryPoint">
18         <beans:property name="realmName" value="Contacts Realm via Digest Authentication" />
19         <beans:property name="key" value="acegi" />
20     </beans:bean>
21
22     <!-- the security namespace configuration -->
23     <http use-expressions="true" entry-point-ref="digestEntryPoint">
24         <intercept-url pattern="/*" access="isAuthenticated()" />
25
26         <custom-filter ref="digestFilter" after="BASIC_AUTH_FILTER" />
27     </http>
28
29     <authentication-manager>
30         <authentication-provider>
31             <user-service id="userService">
32                 <user name="user1" password="user1Pass" authorities="ROLE_USER" />
33             </user-service>
34         </authentication-provider>

```



```

35 |         </authentication-manager>
36 |
37 | </beans:beans>

```

Next, we need to integrate these beans into the overall security configuration – and in this case, the namespace is still flexible enough to allow us to do that.

The first part of this is pointing to the custom entry point bean, via the *entry-point-ref* attribute of the main *<http>* element.

The second part is **adding the newly defined digest filter into the security filter chain**. Since this filter is functionally equivalent to the *BasicAuthenticationFilter*, we are using the same relative position in the chain – this is specified by the *BASIC_AUTH_FILTER* alias in the overall [Spring Security Standard Filters](#).

Finally, notice that the Digest Filter is configured to **point to the user service bean** – and here, the namespace is again very useful as it allows us to specify a bean name for the default user service created by the *<user-service>* element:

```

1 | <user-service id="userService">

```

3. Consuming the Secured Application

We're going to be using **the *curl* command** to consume the secured application and understand how a client can interact with it.

Let's start by requesting a *Foo* Resource – **without providing security credentials** in the request:

```

1 | curl -i http://localhost:8080/spring-security-rest-digest-auth/api/foos/1

```

As expected, we get back a response with a *401 Unauthorized* status code:

```

1 | HTTP/1.1 401 Unauthorized
2 | Server: Apache-Coyote/1.1
3 | Set-Cookie: JSESSIONID=CF0233CEE737576C43B12FBB6C62A67E; Path=/spring-security-rest-digest-auth/; HttpOnly
4 | WWW-Authenticate: Digest realm="Contacts Realm via Digest Authentication", qop="auth",
5 |   nonce="MTM3MzYzODE2NTg3OTc0MmYxN2JkOWYxZTc4MzdmMzBiN2Q0YmY0ZTU0N2RkZg=="
6 | Content-Type: text/html;charset=utf-8
7 | Content-Length: 1061
8 | Date: Fri, 12 Jul 2013 14:04:25 GMT

```

If this request were sent by the browser, the authentication challenge would prompt the user for credentials using a simple user/password dialog.

Let's now **provide the correct credentials** and send the request again:

```

1 | curl -i --digest --user user1:user1Pass http://localhost:8080/spring-security-rest-digest-auth/api/foos/1

```

Notice that we are enabling Digest Authentication for the *curl* command via the *-digest* flag.

The first response from the server will be the same – the *401 Unauthorized* – but the challenge will now be interpreted and acted upon by a second request – which will succeed with a *200 OK*:

```

1 | HTTP/1.1 401 Unauthorized
2 | Server: Apache-Coyote/1.1

```

```

3 Set-Cookie: JSESSIONID=A961E0D09484F58D5885AE6D02C99445; Path=/spring-security-rest-digest-auth/;
4 HttpOnly
5 WWW-Authenticate: Digest realm="Contacts Realm via Digest Authentication", qop="auth",
6   nonce="MTM3MzYzODgyOTczMTo3YjM4OWQzMGU0YTgwZDg0YmYwZjRlZWJjMDQzZWZkOA=="
7 Content-Type: text/html;charset=utf-8
8 Content-Length: 1061
9 Date: Fri, 12 Jul 2013 14:15:29 GMT
10
11 HTTP/1.1 200 OK
12 Server: Apache-Coyote/1.1
13 Set-Cookie: JSESSIONID=55F996B6839E9E06CD546D8F840027C4; Path=/spring-security-rest-digest-auth/;
14 HttpOnly
15 Content-Type: text/html;charset=ISO-8859-1
16 Content-Language: en-US
17 Content-Length: 90
   Date: Fri, 12 Jul 2013 14:15:29 GMT
   {}

```

A final note on this interaction is that a client can **preemptively send the correct *Authorization* header** with the first request, and thus entirely **avoid the server security challenge** and the second request.

4. The Maven Dependencies

The security dependencies are discussed in depth in the [Spring Security Maven tutorial](#). In short, we will need to define *spring-security-web* and *spring-security-config* as dependencies in our *pom.xml*.

5. Conclusion

In this section we introduced security into a simple Spring REST API by leveraging the Digest Authentication support in the framework.

The implementation of these examples can be found in [the github project](#) – this is an Eclipse based project, so it should be easy to import and run as it is.

Finally, in the next section we'll see that there is no reason to choose between Basic and Digest authentication – **both can be set up simultaneously on the same URI structure**. The client can then pick between the two mechanisms when consuming the web application.

VI. Security for REST – Basic and Digest Authentication together

1. Overview

This section discusses how to **set up both Basic and Digest Authentication on the same URI structure of a REST API**. In the two previous sections, we discussed Basic and Digest authentication individually – so now we're going to focus on how to set up both mechanisms on the same URI structure.

2. Configuration of Basic Authentication

The main reason that form based authentication is not ideal for a RESTful Service is that Spring Security will

make use of Sessions – this is of course state on the server, so **the statelessness constraints in REST** is practically ignored.

We'll start by setting up Basic Authentication – first we remove the old custom entry point and filter from the main `<http>` security element:

```

1 | <http create-session="stateless">
2 |   <intercept-url pattern="/api/admin/**" access="ROLE_ADMIN" />
3 |
4 |   <http-basic />
5 | </http>

```

Note how support for basic authentication has been added with a single configuration line – `<http-basic />` – which handles the creation and wiring of both the `BasicAuthenticationFilter` and the `BasicAuthenticationEntryPoint`.

2.1. Satisfying the stateless constraint – getting rid of sessions

One of the main constraints of the RESTful architectural style is that the client-server communication is fully **stateless**, as the [original dissertation](#) reads:

□ 5.1.3 Stateless

*We next add a constraint to the client-server interaction: communication must be stateless in nature, as in the client-stateless-server (CSS) style of Section 3.4.3 (Figure 5-3), such that each request from client to server must contain all of the information necessary to understand the request, and cannot take advantage of any stored context on the server. **Session state is therefore kept entirely on the client.***

The concept of **Session** on the server is one with a long history in Spring Security, and removing it entirely has been difficult until now, especially when configuration was done by using the namespace. However, Spring Security 3.1 [augments](#) the namespace configuration with a **new stateless option** for session creation, which effectively guarantees that no session will be created or used by Spring. What this new option does is completely removes all session related filters from the security filter chain, ensuring that authentication is performed for each request.

3. Configuration of Digest Authentication

Starting with the previous configuration, the filter and entry point necessary to set up digest authentication will be defined as beans. Then, the **digest entry point** will override the one created by `<http-basic>` behind the scenes. Finally, the custom **digest filter** will be introduced in the security filter chain using the *after* semantics of the security namespace to position it directly after the basic authentication filter.

```

1 | <http create-session="stateless" entry-point-ref="digestEntryPoint">
2 |   <intercept-url pattern="/api/admin/**" access="ROLE_ADMIN" />
3 |

```

```

4     <http-basic />
5     <custom-filter ref="digestFilter" after="BASIC_AUTH_FILTER" />
6 </http>
7
8 <beans:bean id="digestFilter" class=
9     "org.springframework.security.web.authentication.www.DigestAuthenticationFilter">
10    <beans:property name="userService" ref="userService" />
11    <beans:property name="authenticationEntryPoint" ref="digestEntryPoint" />
12 </beans:bean>
13
14 <beans:bean id="digestEntryPoint" class=
15     "org.springframework.security.web.authentication.www.DigestAuthenticationEntryPoint">
16    <beans:property name="realmName" value="Contacts Realm via Digest Authentication"/>
17    <beans:property name="key" value="acegi" />
18 </beans:bean>
19
20 <authentication-manager>
21   <authentication-provider>
22     <user-service id="userService">
23       <user name="eparaschiv" password="eparaschiv" authorities="ROLE_ADMIN" />
24       <user name="user" password="user" authorities="ROLE_USER" />
25     </user-service>
26   </authentication-provider>
27 </authentication-manager>

```

Unfortunately there is no [support](#) in the security namespace to automatically configure the digest authentication the way basic authentication can be configured with `<http-basic>`. Because of that, the necessary beans had to be defined and wired manually into the security configuration.

4. Supporting both authentication protocols in the same RESTful service

Basic or Digest authentication alone can be easily implemented in Spring Security 3.x; it is supporting both of them for the same RESTful web service, on the same URI mappings that introduces a new level of complexity into the configuration and testing of the service.

4.1. Anonymous request

With both basic and digest filters in the security chain, the way a **anonymous request** – a request containing no authentication credentials (*Authorization* HTTP header) – is processed by Spring Security is – the two authentication filters will find **no credentials** and will continue execution of the filter chain. Then, seeing how the request wasn't authenticated, an *AccessDeniedException* is thrown and caught in the *ExceptionTranslationFilter*, which commences the digest entry point, prompting the client for credentials.

The responsibilities of both the basic and digest filters are very narrow – they will continue to execute the security filter chain if they are unable to identify the type of authentication credentials in the request. It is because of this that Spring Security can have the flexibility to be configured with support for multiple authentication protocols on the same URI.

When a request is made containing the correct authentication credentials – either basic or digest – that protocol will be rightly used. However, for an anonymous request, the client will get prompted only for digest authentication credentials. This is because the digest entry point is configured as the main and single entry point of the Spring Security chain; as such **digest authentication can be considered the default**.

4.2. Request with authentication credentials

A **request with credentials** for Basic authentication will be identified by the *Authorization* header starting with the prefix “Basic”. When processing such a request, the credentials will be decoded in the basic authentication filter and the request will be authorized. Similarly, a request with credentials for Digest authentication will use the prefix “Digest” for its *Authorization* header.

5. Testing both scenarios

The tests will consume the REST service by creating a new resource after authenticating with either basic or digest:

```

1  @Test
2  public void givenAuthenticatedByBasicAuth_whenAResourceIsCreated_then201IsReceived(){
3      // Given
4      // When
5      Response response = given()
6          .auth().preemptive().basic( ADMIN_USERNAME, ADMIN_PASSWORD )
7          .contentType( HttpConstants.MIME_JSON ).body( new Foo( randomAlphabetic( 6 ) ) )
8          .post( paths.getFooURL() );
9
10     // Then
11     assertThat( response.getStatusCode(), is( 201 ) );
12 }
13 @Test
14 public void givenAuthenticatedByDigestAuth_whenAResourceIsCreated_then201IsReceived(){
15     // Given
16     // When
17     Response response = given()
18         .auth().digest( ADMIN_USERNAME, ADMIN_PASSWORD )
19         .contentType( HttpConstants.MIME_JSON ).body( new Foo( randomAlphabetic( 6 ) ) )
20         .post( paths.getFooURL() );
21
22     // Then
23     assertThat( response.getStatusCode(), is( 201 ) );
24 }

```

Note that the test using basic authentication adds credentials to the request **preemptively**, regardless if the server has challenged for authentication or not. This is to ensure that the server doesn't need to challenge the client for credentials, because if it did, the challenge would be for Digest credentials, since that is the default.

6. Conclusion

This section covered the configuration and implementation of both Basic and Digest authentication for a RESTful service, using mostly Spring Security 3 namespace support as well as some new features added by Spring Security 3.1.

For the full implementation, check out the [github project](#).

VII. REST API Discoverability and HATEOAS

1. Overview

The section will focus on **Discoverability of the REST API, HATEOAS** and practical scenarios driven by tests.

2. Introducing REST Discoverability

Discoverability of an API is a topic that doesn't get enough well deserved attention, and as a consequence very few APIs get it right. It is also something that, if done right, can make the API not only RESTful and usable but also elegant.

To understand **discoverability**, one needs to understand that constraint that is Hypermedia As The Engine Of Application State (HATEOAS); this constraint of a REST API is about full discoverability of actions/transitions on a Resource from Hypermedia (Hypertext really), as the **only driver** of application state. If interaction is to be **driven** by the API through the conversation itself, concretely via Hypertext, then there can be **no documentation**, as that would coerce the client to make assumptions that are in fact outside of the context of the API.

Also, continuing this logical train of thought, the only way an API can indeed be considered RESTful is if it is fully discoverable **from the root** and with **no prior knowledge** – meaning the client should be able to navigate the API by doing a GET on the root. Moving forward, all state changes are driven by the client using the available and discoverable transitions that the REST API provides in representations (hence *Representational State Transfer*).

In conclusion, the server should be descriptive enough to instruct the client how to use the API via Hypertext only, which, in the case of a HTTP conversation, may be the **Link** header.

3. Discoverability Scenarios (Driven by tests)

So what does it mean for a REST service to be **discoverable**? Throughout this section, we will test individual traits of discoverability using Junit, [rest-assured](#) and [Hamcrest](#). Since [the REST Service has been previously secured](#), each test first need to [authenticate](#) before consuming the API.

Some [utilities](#) for parsing the *Link* header of the response are also necessary:

```

1  public static List<String> parseLinkHeader(String linkHeader) {
2      List<String> linkHeaders = Lists.newArrayList();
3      String[] links = linkHeader.split(", ");
4      for (String link : links) {
5          int positionOfSeparator = link.indexOf(';');
6          linkHeaders.add(link.substring(1, positionOfSeparator - 1));
7      }
8      return linkHeaders;
9  }
10 public static String parseSingleLinkHeader(String linkHeader) {
11     int positionOfSeparator = linkHeader.indexOf(';');
12     return linkHeader.substring(1, positionOfSeparator - 1);
13 }

```

3.1. Discover the valid HTTP methods

When a REST Service is consumed with an **invalid HTTP method**, the response should be a **405 METHOD NOT ALLOWED**; in addition, it should also help the client **discover** the valid HTTP methods that are allowed for that particular Resource, using the **Allow** HTTP Header in the response:

```

2 public void
3   whenInvalidPOSTIsSentToValidURIOfResource_thenAllowHeaderListsTheAllowedActions(){
4     // Given
5     String uriOfExistingResource = restTemplate.createResource();
6
7     // When
8     Response res = givenAuthenticated().post( uriOfExistingResource );
9
10    // Then
11    String allowHeader = res.getHeader( HttpHeaders.ALLOW );
12    assertThat( allowHeader, AnyOf.<String> anyOf(
13      containsString("GET"), containsString("PUT"), containsString("DELETE") ) );
14  }

```

3.2. Discover the URI of newly created Resource

The operation of creating a new Resource should always include the URI of the newly created resource in the response, using the **Location** HTTP Header. If the client does a GET on that URI, the resource should be available:

```

1 @Test
2 public void whenResourceIsCreated_thenUriOfTheNewlyCreatedResourceIsDiscoverable(){
3     // When
4     Foo unpersistedResource = new Foo( randomAlphabetic( 6 ) );
5     Response createResponse = givenAuthenticated().contentType( MIME_JSON )
6       .body( unpersistedResource ).post( paths.getFooURL() );
7     String uriOfNewlyCreatedResource = createResp
8       .getHeader( HttpHeaders.LOCATION );
9
10    // Then
11    Response response = givenAuthenticated()
12      .header( HttpHeaders.ACCEPT, MIME_JSON ).get( uriOfNewlyCreatedResource );
13
14    Foo resourceFromServer = response.body().as( Foo.class );
15    assertThat( unpersistedResource, equalTo( resourceFromServer ) );
16  }

```

The test follows a simple scenario: a new *Foo* resource is created and the HTTP response is used to **discover the URI** where the Resource is now accessible. The tests then goes one step further and does a GET on that URI to retrieve the resource and compares it to the original, to make sure that it has been correctly persisted.

3.3. Discover the URI to GET All Resources of that type

When we GET any particular *Foo* resource, we should be able to **discover** what we can do next: we can list all the available *Foo* resources. Thus, the operation of retrieving an resource should always include in its response the URI where to get all the resources of that type, again making use of the **Link** header:

```

1 @Test
2 public void whenResourceIsRetrieved_thenUriToGetAllResourcesIsDiscoverable(){
3     // Given
4     String uriOfExistingResource = restTemplate.createResource();
5
6     // When
7     Response getResponse = givenAuthenticated().get( uriOfExistingResource );
8
9     // Then
10    String uriToAllResources = HTTPLinkHeaderUtils.extractURIByRel
11      ( getResponse.getHeader( "Link" ), "collection" );
12
13    Response getAllResponse = givenAuthenticated().get( uriToAllResources );
14    assertThat( getAllResponse.getStatusCode(), is( 200 ) );
15  }

```

The test tackles a thorny subject of Link Relations in REST: the URI to retrieve all resources uses the

rel="collection" semantics. This type of link relation has not yet been standardized, but is already [in use](#) by several microformats and proposed for standardization. Usage of non-standard link relations opens up the discussion about microformats and richer semantics in RESTful web services.

4. Other potential discoverable URIs and microformats

Other URIs could potentially be discovered via the **Link** header, but there is only so much the existing types of link relations allow without moving to a richer semantic markup such as [defining custom link relations](#), the [Atom Publishing Protocol](#) or [microformats](#).

For example it would be good if the client could discover the URI to create new resources when doing a GET on a specific resource; unfortunately there is no link relation to model **create** semantics. Luckily it is standard practice that the URI for creation is the same as the URI to GET all resources of that type, with the only difference being the POST HTTP method.

5. Conclusion

This section covered the some of the traits of discoverability in the context of a REST web service, discussing HTTP method discovery, the relation between create and get, discovery of the URI to get all resources, etc.

For the full implementation, check out the [github project](#).

VIII. HATEOAS for a Spring REST Service

1. Overview

This section will focus on the **implementation of discoverability in a Spring REST Service** and on satisfying the HATEOAS constraint.

2. Decouple Discoverability through events

Discoverability as a **separate aspect or concern** of the web layer should be decoupled from the controller handling the HTTP request. In order to do so, the Controller will fire off events for all the actions that require additional manipulation of the HTTP response.

First, the events:

```

1 public class SingleResourceRetrieved extends ApplicationEvent {
2     private final HttpServletResponse response;
3     private final HttpServletRequest request;
4
5     public SingleResourceRetrieved(Object source,
6         HttpServletRequest request, HttpServletResponse response) {
7         super(source);
8
9         this.request = request;
10        this.response = response;
11    }
12

```



```

13     public HttpServletResponse getResponse() {
14         return response;
15     }
16     public HttpServletRequest getRequest() {
17         return request;
18     }
19 }
20 public class ResourceCreated extends ApplicationEvent {
21     private final HttpServletResponse response;
22     private final HttpServletRequest request;
23     private final long idOfNewResource;
24
25     public ResourceCreated(Object source,
26         HttpServletRequest request, HttpServletResponse response, long idOfNewResource) {
27         super(source);
28
29         this.request = request;
30         this.response = response;
31         this.idOfNewResource = idOfNewResource;
32     }
33
34     public HttpServletResponse getResponse() {
35         return response;
36     }
37     public HttpServletRequest getRequest() {
38         return request;
39     }
40     public long getIdOfNewResource() {
41         return idOfNewResource;
42     }
43 }

```

The Controller:

```

1  @Autowired
2  private ApplicationEventPublisher eventPublisher;
3
4  @RequestMapping( value = "admin/foo/{id}",method = RequestMethod.GET )
5  @ResponseBody
6  public Foo get( @PathVariable( "id" ) Long id,
7      HttpServletRequest request, HttpServletResponse response ){
8      Foo resourceById = Preconditions.checkNotNull( service.getById( id ) );
9
10     eventPublisher.publishEvent( new SingleResourceRetrieved( this, request, response ) );
11     return resourceById;
12 }
13 @RequestMapping( value = "admin/foo",method = RequestMethod.POST )
14 @ResponseStatus( HttpStatus.CREATED )
15 public void create( @RequestBody Foo resource,
16     HttpServletRequest request, HttpServletResponse response ){
17     Preconditions.checkNotNull( resource );
18     Long idOfCreatedResource = service.create( resource );
19
20     eventPublisher.publishEvent(
21         new ResourceCreated( this, request, response, idOfCreatedResource ) );
22 }

```

These events can then be handled by any number of **decoupled listeners**, each focusing on it's own particular case and each moving towards satisfying the overall HATEOAS constraint.

Also, the listeners should be the last objects in the call stack and no direct access to them is necessary; as such they are not public.

3. Make the URI of a newly created resource discoverable

As discussed in the previous section on HATEOAS, the operation of creating a new resource should return the URI of that resource in the **Location** HTTP header of the response. :

```

1  @Component
2  class ResourceCreatedDiscoverabilityListener
3      implements ApplicationListener< ResourceCreated >{
4
5      @Override
6      public void onApplicationEvent( ResourceCreated resourceCreatedEvent ){
7          Preconditions.checkNotNull( resourceCreatedEvent );
8
9          HttpServletRequest request = resourceCreatedEvent.getRequest();
10         HttpServletResponse response = resourceCreatedEvent.getResponse();
11         long idOfNewResource = resourceCreatedEvent.getIdOfNewResource();
12
13         addLinkHeaderOnResourceCreation( request, response, idOfNewResource );
14     }
15     void addLinkHeaderOnResourceCreation
16     ( HttpServletRequest request, HttpServletResponse response, long idOfNewResource ){
17         String requestUrl = request.getRequestURL().toString();
18         URI uri = new UriTemplate( "{requestUrl}/{idOfNewResource}" ).
19         expand( requestUrl, idOfNewResource );
20         response.setHeader( "Location", uri.toASCIIString() );
21     }
22 }

```

Unfortunately, dealing with the low level request and response objects is inevitable even in Spring 3.1, because first class support for specifying the *Location* is [still in the works](#).

4. Get of single resource

Retrieving a single resource should allow the client to discover the URI to get all resources of that particular type:

```

1  @Component
2  class SingleResourceRetrievedDiscoverabilityListener
3      implements ApplicationListener< SingleResourceRetrieved >{
4
5      @Override
6      public void onApplicationEvent( SingleResourceRetrieved resourceRetrievedEvent ){
7          Preconditions.checkNotNull( resourceRetrievedEvent );
8
9          HttpServletRequest request = resourceRetrievedEvent.getRequest();
10         HttpServletResponse response = resourceRetrievedEvent.getResponse();
11         addLinkHeaderOnSingleResourceRetrieval( request, response );
12     }
13     void addLinkHeaderOnSingleResourceRetrieval
14     ( HttpServletRequest request, HttpServletResponse response ){
15         StringBuffer requestURL = request.getRequestURL();
16         int positionOfLastSlash = requestURL.lastIndexOf( "/" );
17         String uriForResourceCreation = requestURL.substring( 0, positionOfLastSlash );
18
19         String linkHeaderValue = LinkUtil
20             .createLinkHeader( uriForResourceCreation, "collection" );
21         response.addHeader( LINK_HEADER, linkHeaderValue );
22     }
23 }

```

Note that the semantics of the link relation make use of the “*collection*” relation type, specified and used in [several microformats](#), but not yet standardized.

The *Link* header is one of the most used HTTP header for the purposes of discoverability. The utility to create this header is simple enough:

```

1  public final class LinkUtil {
2      public static String createLinkHeader( final String uri, final String rel ) {
3          return "<" + uri + ">; rel=\"" + rel + "\"";
4      }
5  }

```

5. Discoverability at the root

The root is the entry point in the RESTful web service – it is what the client comes into contact with when consuming the API for the first time. If the HATEOAS constraint is to be considered and implemented throughout, then this is the place to start. The fact that most of the main URIs of the system have to be discoverable from the root shouldn't come as much of a surprise by this point.

This is a sample controller method to provide discoverability at the root:

```

1  @RequestMapping( value = "admin",method = RequestMethod.GET )
2  @ResponseStatus( value = HttpStatus.NO_CONTENT )
3  public void adminRoot( HttpServletRequest request, HttpServletResponse response ){
4      String rootUri = request.getRequestURL().toString();
5
6      URI fooUri = new UriTemplate( "{rootUri}/{resource}" ).expand( rootUri, "foo" );
7      String linkToFoo = LinkUtil.createLinkHeader
8          ( fooUri.toASCIIString(), "collection" );
9      response.addHeader( "Link", linkToFoo );
10 }

```

This is of course an illustration of the concept, to be read in the context of the proof of concept RESTful service of the series. In a more complex system there would be many more links, each with its own semantics defined by the type of [link relation](#).

5.1. Discoverability is not about changing URIs

One of the more common pitfalls related to discoverability is the misunderstanding that, since the URIs are now discoverable, then they can be **subject to change**. This is however simply not the case, and for good reason: first, this is not how the web works – clients will bookmark the URIs and will expect them to work in the future. Second, the client shouldn't have to navigate through the API to get to a certain state that could have been reached directly.

Instead, all URIs of the RESTful web service should be considered **cool URIs**, and cool URIs **don't change**. Instead, **versioning** of the API can be used to solve the problem of a URI reorganization.

6. Caveats of Discoverability

The first goal of discoverability is to make minimal or no use of **documentation** and have the client learn and understand how to use the API via the responses it gets. In fact, this shouldn't be regarded as such a far fetched ideal – it is how we consume every new web page – **without any documentation**. So, if the concept is more problematic in the context of REST, then it must be a matter of technical implementation, not of a question of whether or not it's possible.

That being said, technically, we are still far from the a fully working solution – the specification and framework support are still evolving, and because of that, some compromises may have to be made; these are nevertheless compromises and should be regarded as such.

7. Conclusion

In this section we covered the implementation of some of the traits of discoverability in the context of a RESTful Service with Spring MVC and touched on the concept of discoverability at the root.

For the full implementation, check out the [github project](#).

IX. ETags for REST with Spring

1. Overview

This section will focus on **working with ETags in Spring**, integration testing of the REST API and consumption scenarios with *curl*.

2. REST and ETags

From the official Spring documentation on ETag support:

*An **ETag** (entity tag) is an HTTP response header returned by an HTTP/1.1 compliant web server used to determine change in content at a given URL.*

ETags are used for two things – caching and conditional requests. The **ETag value can be thought as a hash** computed out of the bytes of the Response body. Because a cryptographic hash function is likely used, even the smallest modification of the body will drastically change the output and thus the value of the ETag. This is only true for strong ETags – the protocol does provide a **weak Etag** as well.

Using an **If-* header** turns a standard GET request into a **conditional GET**. The two *If-** headers that are using with ETags are “**If-None-Match**” and “**If-Match**” – each with it's own semantics as discussed later in this section.

3. Client-Server communication with *curl*

A simple Client-Server communication involving ETags can be broken down into the steps:

- **first**, the Client makes a REST API call – the Response includes the ETag header to be stored for further use:

```
curl -H "Accept: application/json" -i http://localhost:8080/rest-sec/api/resources/1
```

```

1 | HTTP/1.1 200 OK
2 | ETag: "f88dd058fe004909615a64f01be66a7"
3 | Content-Type: application/json;charset=UTF-8
4 | Content-Length: 52
```

- **next** request the Client makes to the RESTful API includes the *If-None-Match* request header with the ETag value from the previous step; if the Resource has not changed on the Server, the Response will contain **no body** and a status code of *304 – Not Modified*:

```
curl -H "Accept: application/json" -H 'If-None-Match: "f88dd058fe004909615a64f01be66a7"' -i http://localhost:8080/rest-sec/api/resources/1
```

```

1 | HTTP/1.1 304 Not Modified
2 | ETag: "f88dd058fe004909615a64f01be66a7"
```

- **now**, before retrieving the Resource again, we will change it by performing an update:

```
curl --user admin@fake.com:adminpass -H "Content-Type: application/json" -i
-X PUT --data '{ "id":1, "name":"newRoleName2", "description":"theNewDescription" }'
```

```
http://localhost:8080/rest-sec/api/resources/1
```

```
1 | HTTP/1.1 200 OK
2 | ETag: "d41d8cd98f00b204e9800998ecf8427e"
3 | Content-Length: 0
```

- **finally**, we send out the the last request to retrieve the Privilege again; keep in mind that it has been updated since the last time it was retrieved, so the previous ETag value should no longer work – the response will contain the new data and a new ETag which, again, can be stored for further use:

```
curl -H "Accept: application/json" -H 'If-None-Match: "f88dd058fe004909615a64f01be66a7"' -i
```

```
http://localhost:8080/rest-sec/api/resources/1
```

```
1 | HTTP/1.1 200 OK
2 | ETag: "03cb37ca667706c68c0aad4cb04c3a211"
3 | Content-Type: application/json;charset=UTF-8
4 | Content-Length: 56
```

And there you have it – ETags in the wild and saving bandwidth.

4. ETag support in Spring

On to the Spring support – to use ETag in Spring is extremely easy to set up and completely **transparent** for the application. The support is enabled by adding a simple *Filter* in the *web.xml*:

```
1 | <filter>
2 |   <filter-name>etagFilter</filter-name>
3 |   <filter-class>org.springframework.web.filter.ShallowEtagHeaderFilter</filter-class>
4 | </filter>
5 | <filter-mapping>
6 |   <filter-name>etagFilter</filter-name>
7 |   <url-pattern>/api/*</url-pattern>
8 | </filter-mapping>
```

The filter is mapped on the same URI pattern as the RESTful API itself. The filter itself is the standard implementation of ETag functionality since Spring 3.0.

The implementation is a **shallow** one – the ETag is calculated based on the response, which will **save bandwidth** but **not server performance**. So, a request that will benefit from the ETag support will still be processed as a standard request, consume any resource that it would normally consume (database connections, etc) and only before having it's response returned back to the client will the ETag support kick in.

At that point the ETag will be calculated out of the Response body and set on the Resource itself; also, if the *If-None-Match* header was set on the Request, it will be handled as well.

A **deeper implementation** of the ETag mechanism could potentially provide much greater benefits – such as serving some requests from the cache and not having to perform the computation at all – but the implementation would most definitely not be as simple, nor as pluggable as the shallow approach described here.

5. Testing ETags

Let's start simple – we need to verify that the response of a simple request retrieving a single Resource will actually return the “ETag” header:

```

1  @Test
2  public void givenResourceExists_whenRetrievingResource_thenEtagIsAlsoReturned() {
3      // Given
4      Resource existingResource = getApi().create(new Resource());
5      String uriOfResource = baseUrl + "/" + existingResource.getId();
6
7      // When
8      Response findOneResponse = RestAssured.given()
9          .header("Accept", "application/json").get(uriOfResource);
10
11     // Then
12     assertNotNull(findOneResponse.getHeader(HttpHeaders.ETAG));
13 }

```

Next, we verify the happy path of the ETag behaviour – if the Request to retrieve the Resource from the server uses the correct ETag value, then the Resource is no longer returned.

```

1  @Test
2  public void givenResourceWasRetrieved_whenRetrievingAgainWithEtag_thenNotModifiedReturned() {
3      // Given
4      T existingResource = getApi().create(createNewEntity());
5      String uriOfResource = baseUrl + "/" + existingResource.getId();
6      Response findOneResponse = RestAssured.given()
7          .header("Accept", "application/json").get(uriOfResource);
8      String etagValue = findOneResponse.getHeader(HttpHeaders.ETAG);
9
10     // When
11     Response secondFindOneResponse= RestAssured.given()
12         .header("Accept", "application/json").headers("If-None-Match", etagValue)
13         .get(uriOfResource);
14
15     // Then
16     assertTrue(secondFindOneResponse.getStatusCode() == 304);
17 }

```

Step by step:

- a Resource is first created and then retrieved – the ETag value is stored for further use
- a new retrieve request is sent, this time with the “If-None-Match” header specifying the ETag value previously stored
- on this second request, the server simply returns a **304 Not Modified**, since the Resource itself has indeed not been modified between the two retrieval operations

Finally, we verify the case where the Resource is **changed** between the first and the second retrieval requests:

```

1  @Test
2  public void givenResourceWasRetrieved_whenRetrievingAgainWithEtag_thenNotModifiedReturned() {
3      // Given
4      T existingResource = getApi().create(createNewEntity());
5      String uriOfResource = baseUrl + "/" + existingResource.getId();
6      Response findOneResponse = RestAssured.given()
7          .header("Accept", "application/json").get(uriOfResource);
8      String etagValue = findOneResponse.getHeader(HttpHeaders.ETAG);
9
10     existingResource.setName(randomAlphabetic(6))
11     getApi().update(existingResource.setName(randomString));
12
13     // When
14     Response secondFindOneResponse= RestAssured.given()
15         .header("Accept", "application/json").headers("If-None-Match", etagValue)
16         .get(uriOfResource);
17 }

```

```

18 | // Then
19 | assertTrue(secondFindOneResponse.getStatusCode() == 200);
20 | }

```

Step by step:

- a *Resource* is first created and then retrieved – the *ETag* value is stored for further use
- the same *Resource* is then updated
- a new retrieve request is sent, this time with the “*If-None-Match*” header specifying the *ETag* value previously stored
- on this second request, the server will return a **200 OK** along with the full *Resource*, since the *ETag* value is no longer correct, as the *Resource* has been updated in the meantime

Next, we test the behavior for “*If-Match*” – the *ShallowEtagHeaderFilter* does not have out of the box support for the *If-Match* HTTP header (being tracked on [this JIRA issue](#)), so the following test should fail:

```

1 | @Test
2 | public void givenResourceExists_whenRetrievedWithIfMatchIncorrectEtag_then412IsReceived() {
3 |     // Given
4 |     T existingResource = getApi().create(createNewEntity());
5 |
6 |     // When
7 |     String uriOfResource = baseUrl + "/" + existingResource.getId();
8 |     Response findOneResponse = RestAssured.given().header("Accept", "application/json").
9 |         headers("If-Match", randomAlphabetic(8)).get(uriOfResource);
10 |
11 |     // Then
12 |     assertTrue(findOneResponse.getStatusCode() == 412);
13 | }

```

Step by step:

- a *Resource* is first created
- the *Resource* is then retrieved with the “*If-Match*” header specifying an incorrect *ETag* value – this is a **conditional GET request**
- the server should return a **412 Precondition Failed**

6. ETags are BIG

We have only used ETags for **read operations** – a [RFC exists](#) trying to clarify how implementations should deal with ETags on **write operations** – this is not standard, but is an interesting read.

There are of course other possible uses of the *ETag* mechanism, such as for an [Optimistic Locking Mechanism using Spring 3.1](#) as well as dealing with the [related “Lost Update Problem”](#).

There are also several known [potential pitfalls and caveats](#) to be aware of when using ETags.

7. Conclusion

In this section, we only scratched the surface with what’s possible with Spring and ETags.

For a full implementation of an *ETag* enabled RESTful service, along with integration tests verifying the *ETag* behavior, check out the [github project](#).

X. REST Pagination in Spring

1. Overview

This section will focus on the **implementation of pagination** in a REST API, using Spring MVC and Spring Data.

2. Page as resource vs Page as representation

The first question when designing pagination in the context of a RESTful architecture is whether to consider the **page an actual resource or just a representation of resources**. Treating the page itself as a resource introduces a host of problems such as no longer being able to uniquely identify resources between calls. This, coupled with the fact that, in the persistence layer, the page is not proper entity but a holder that is constructed when necessary, makes the choice straightforward: **the page is part of the representation**.

The next question in the pagination design in the context of REST is where to include the paging information:

- in the **URI path**: `/foo/page/1`
- the **URI query**: `/foo?page=1`

Keeping in mind that **a page is not a resource**, encoding the page information in the URI is no longer an option.

3. Page information in the URI query

Encoding paging information in the **URI query** is the standard way to solve this issue in a REST Service. This approach does however have one **downside** – it cuts into the query space for actual queries:

```
/foo?page=1&size=10
```

4. The Controller

Now, for the implementation – the Spring **MVC Controller for pagination** is straightforward:

```

1  @RequestMapping( value = "admin/foo",params = { "page", "size" },method = GET )
2  @ResponseBody
3  public List< Foo > findPaginated(
4  @RequestParam( "page" ) int page, @RequestParam( "size" ) int size,
5  UriComponentsBuilder uriBuilder, HttpServletResponse response ){
6
7      Page< Foo > resultPage = service.findPaginated( page, size );
8      if( page > resultPage.getTotalPages() ){
9          throw new ResourceNotFoundException();
10     }
11     eventPublisher.publishEvent( new PaginatedResultsRetrievedEvent< Foo >
12         ( Foo.class, uriBuilder, response, page, resultPage.getTotalPages(), size ) );
13
14     return resultPage.getContent();
15 }
```

The two query parameters are injected into the Controller method via `@RequestParam`. The HTTP response and the `UriComponentsBuilder` are injected so that they can be included in the event – both are needed to implement

discoverability.

5. Discoverability for REST pagination

Withing the scope of **pagination**, satisfying the **HATEOAS constraint of REST** means enabling the client of the API to discover the *next* and *previous* pages based on the current page in the navigation. For this purpose, the **Link HTTP header** will be used, coupled with the **official** “*next*”, “*prev*”, “*first*” and “*last*” link relation types.

In REST, **Discoverability is a cross cutting concern**, applicable not only to specific operations but to types of operations. For example, each time a Resource is created, the URI of that resource should be discoverable by the client. Since this requirement is relevant for the creation of ANY Resource, it should be dealt with separately and decoupled from the main Controller flow.

With Spring, this **decoupling is achieved with events**, as was thoroughly discussed in the previous section focused on Discoverability of the REST Service. In the case of pagination, the event – *PaginatedResultsRetrievedEvent* – was fired in the Controller, and discoverability is achieved in a listener for this event:

```

1  void addLinkHeaderOnPagedResourceRetrieval(
2  UriComponentsBuilder uriBuilder, HttpServletResponse response,
3  Class clazz, int page, int totalPages, int size ){
4
5      String resourceName = clazz.getSimpleName().toString().toLowerCase();
6      uriBuilder.path( "/admin/" + resourceName );
7
8      StringBuilder linkHeader = new StringBuilder();
9      if( hasNextPage( page, totalPages ) ){
10         String uriNextPage = constructNextPageUri( uriBuilder, page, size );
11         linkHeader.append( createLinkHeader( uriForNextPage, REL_NEXT ) );
12     }
13     if( hasPreviousPage( page ) ){
14         String uriPrevPage = constructPrevPageUri( uriBuilder, page, size );
15         appendCommaIfNecessary( linkHeader );
16         linkHeader.append( createLinkHeader( uriForPrevPage, REL_PREV ) );
17     }
18     if( hasFirstPage( page ) ){
19         String uriFirstPage = constructFirstPageUri( uriBuilder, size );
20         appendCommaIfNecessary( linkHeader );
21         linkHeader.append( createLinkHeader( uriForFirstPage, REL_FIRST ) );
22     }
23     if( hasLastPage( page, totalPages ) ){
24         String uriLastPage = constructLastPageUri( uriBuilder, totalPages, size );
25         appendCommaIfNecessary( linkHeader );
26         linkHeader.append( createLinkHeader( uriForLastPage, REL_LAST ) );
27     }
28     response.addHeader( HttpConstants.LINK_HEADER, linkHeader.toString() );
29 }

```

In short, the listener logic checks if the navigation allows for a next, previous, first and last pages and, if it does, adds the relevant URIs to the Link HTTP Header. It also makes sure that the link relation type is the correct one – “next”, “prev”, “first” and “last”. This is the single responsibility of the listener ([the full code here](#)).

6. Test Driving Pagination

Both the main logic of pagination and discoverability are covered by small, focused integration tests; the **rest-assured library** is used to consume the REST service and to verify the results.

These are a few example of pagination integration tests; for a full test suite, check out the github project (link at the end of the section):

```

1  @Test
2  public void whenResourcesAreRetrievedPaged_then200IsReceived(){
3      Response response = givenAuth().get( paths.getFooURL() + "?page=1&size=10" );
4
5      assertThat( response.getStatusCode(), is( 200 ) );
6  }
7  @Test
8  public void
9  whenPageOfResourcesAreRetrievedOutOfBounds_then404IsReceived(){
10     Response response = givenAuth().get(
11         paths.getFooURL() + "?page=" + randomNumeric( 5 ) + "&size=10" );
12
13     assertThat( response.getStatusCode(), is( 404 ) );
14 }
15 @Test
16 public void
17 givenResourcesExist_whenFirstPageIsRetrieved_thenPageContainsResources(){
18     restTemplate.createResource();
19
20     Response response = givenAuth().get( paths.getFooURL() + "?page=1&size=10" );
21
22     assertFalse( response.body().as( List.class ).isEmpty() );
23 }

```

7. Test Driving Pagination Discoverability

Testing that **pagination is discoverable** by a client is relatively straightforward, although there is a lot of ground to cover. The tests are focused on the **position of the current page in navigation** and the different URIs that should be discoverable from each position:

```

1  @Test
2  public void whenFirstPageOfResourcesAreRetrieved_thenSecondPageIsNext(){
3      Response response = givenAuth().get( paths.getFooURL()+"?page=0&size=10" );
4
5      String uriToNextPage = extractURIByRel( response.getHeader( LINK ), REL_NEXT );
6      assertEquals( paths.getFooURL()+"?page=1&size=10", uriToNextPage );
7  }
8  @Test
9  public void whenFirstPageOfResourcesAreRetrieved_thenNoPreviousPage(){
10     Response response = givenAuth().get( paths.getFooURL()+"?page=0&size=10" );
11
12     String uriToPrevPage = extractURIByRel( response.getHeader( LINK ), REL_PREV );
13     assertNull( uriToPrevPage );
14 }
15 @Test
16 public void whenSecondPageOfResourcesAreRetrieved_thenFirstPageIsPrevious(){
17     Response response = givenAuth().get( paths.getFooURL()+"?page=1&size=10" );
18
19     String uriToPrevPage = extractURIByRel( response.getHeader( LINK ), REL_PREV );
20     assertEquals( paths.getFooURL()+"?page=0&size=10", uriToPrevPage );
21 }
22 @Test
23 public void whenLastPageOfResourcesIsRetrieved_thenNoNextPageIsDiscoverable(){
24     Response first = givenAuth().get( paths.getFooURL()+"?page=0&size=10" );
25     String uriToLastPage = extractURIByRel( first.getHeader( LINK ), REL_LAST );
26
27     Response response = givenAuth().get( uriToLastPage );
28
29     String uriToNextPage = extractURIByRel( response.getHeader( LINK ), REL_NEXT );
30     assertNull( uriToNextPage );
31 }

```

These are just a few examples of integration tests consuming the RESTful Service.

8. Getting All Resources

On the same topic of pagination and discoverability, the choice must be made if a client is allowed to **retrieve all the Resources in the system** at once, or if the client **MUST** ask for them paginated.

If the choice is made that the client cannot retrieve all Resources with a single request, and pagination is not optional but required, then several options are available for the **response to a get all request**.

One option is to return a **404 (Not Found)** and use the **Link header** to make the first page discoverable:

`Link=<http://localhost:8080/rest/api/admin/foo?page=0&size=10>; rel="first",
<http://localhost:8080/rest/api/admin/foo?page=103&size=10>; rel="last"`

Another option is to return redirect – **303 (See Other)** – to the first page of the pagination.

A third option is to return a **405 (Method Not Allowed)** for the GET request.

9. REST Paging with *Range* HTTP headers

A relatively different way of doing pagination is to work with the **HTTP Range headers** – *Range*, *Content-Range*, *If-Range*, *Accept-Ranges* – and **HTTP status codes** – 206 (*Partial Content*), 413 (*Request Entity Too Large*), 416 (*Requested Range Not Satisfiable*). One view on this approach is that the HTTP Range extensions were not intended for pagination, and that they should be managed by the Server, not by the Application.

Implementing pagination based on the HTTP Range header extensions is nevertheless technically possible, although not nearly as common as the implementation discussed in this section.

10. Conclusion

We illustrated and focused on how to implement Pagination in a REST API using Spring 3, and we discussed how to set up and test Discoverability. For a full implementation of pagination, check out the [github project](#).

XI. Error Handling for REST with Spring 3

1. Overview

This section will focus on **how to implement Exception Handling with Spring for a REST API**. We'll look at the older solutions available before Spring 3.2 and then at the new Spring 3.2 support.

The main goal of this section is to show how to best map Exceptions in the application to HTTP Status Codes. Which status codes are suitable for which scenarios is not in scope, neither is the syntax of REST Error Representation.

Before Spring 3.2, the two main approaches to handling exceptions in a Spring MVC application were: *HandlerExceptionResolver* and the *@ExceptionHandler* annotation. Spring 3.2 introduced the new *@ControllerAdvice* annotation to address the limitations of the previous two solutions.

All of these do have one thing in common – they deal with the **separation of concerns** very well: the standard application code can throw exception normally to indicate a failure of some kind – exceptions which will then be handled via any of the following.

2. Via Controller level *@ExceptionHandler*

Defining a Controller level method annotated with *@ExceptionHandler* is very easy:

```

1 public class FooController{
2     ...
3     @ExceptionHandler({ CustomException1.class, CustomException2.class })
4     public void handleException() {
5         //
6     }
7 }

```

That's all well and good, but this approach does have one major drawback – the *@ExceptionHandler* annotated method is **only active for that particular Controller**, not globally for the entire application. Of course, this makes it not well suited for a generic exception handling mechanism.

A common solution for this is to have all Controllers in the application **extending a Base Controller class** – however, this can be a problem for applications where, for whatever reasons, the Controllers cannot be made to extend from such a class. For example, the Controllers may already extend from another base class which may be in another jar or not directly modifiable, or may themselves not be directly modifiable.

Next, we'll look at another way to solve the exception handling problem – one that is global and does not include any changes to existing artifacts such as Controllers.

3. Via *HandlerExceptionResolver*

In order to implement a **uniform exception handling mechanism** in our REST API, we'll need to work with an *HandlerExceptionResolver* – this will resolve any exceptions thrown at runtime by the application. Before going for a custom resolver, let's go over the existing implementations.

3.1. *ExceptionHandlerExceptionResolver*

This resolver was introduced in Spring 3.1 and is enabled by default in the *DispatcherServlet*. This is actually the core component of how the *@ExceptionHandler* mechanism presented earlier works.

3.2. *DefaultHandlerExceptionResolver*

This resolver was introduced in Spring 3.0 and is enabled by default in the *DispatcherServlet*. It is used to resolve standard Spring exceptions to their corresponding HTTP Status Codes, namely Client error – 4xx and Server error – 5xx status codes. [Here is the full list](#) of the Spring Exceptions it handles, and how these are mapped to status codes.

While it does set the Status Code of the Response properly, one **limitation** of this resolver is that it doesn't set anything to the body of the Response. However, in the context of a REST API, the Status Code is really **not enough information** to present to the Client – the response has to have a body as well, to allow the application to give additional information about the cause of the failure.

This can be solved by configuring View resolution and rendering error content through *ModelAndView*, but the solution is clearly not optimal – which is why a better option has been made available with Spring 3.2 – we'll talk about that in the latter part of this section.

3.3. *ResponseStatusExceptionHandler*

This resolver was also introduced in Spring 3.0 and is enabled by default in the *DispatcherServlet*. It's main responsibility is to use the **@ResponseStatus** annotation available on custom exceptions and to map these exceptions to HTTP status codes.

Such a custom exception may look like:

```

1  @ResponseStatus(value = HttpStatus.NOT_FOUND)
2  public final class ResourceNotFoundException extends RuntimeException {
3      public ResourceNotFoundException() {
4          super();
5      }
6      public ResourceNotFoundException(String message, Throwable cause) {
7          super(message, cause);
8      }
9      public ResourceNotFoundException(String message) {
10         super(message);
11     }
12     public ResourceNotFoundException(Throwable cause) {
13         super(cause);
14     }
15 }

```

Same as the *DefaultHandlerExceptionHandler*, this resolver is **limited** in the way it deals with the body of the response – it does map the Status Code on the response, but the body is still null.

3.4. *SimpleMappingExceptionHandler* and *AnnotationMethodHandlerExceptionHandler*

The *SimpleMappingExceptionHandler* has been around for quite some time – it comes out of the older Spring MVC model and is **not very relevant for a REST Service**. It is used to map exception class names to **view** names.

The *AnnotationMethodHandlerExceptionHandler* was introduced in Spring 3.0 to handle exceptions through the **@ExceptionHandler** annotation, but has been **deprecated** by *ExceptionHandlerResolver* as of Spring 3.2.

3.5. Custom *HandlerExceptionHandler*

The combination of *DefaultHandlerExceptionHandler* and *ResponseStatusExceptionHandler* goes a long way towards providing a good error handling mechanism for a Spring RESTful Service – but the major limitation – no control over the body of the response – justifies creating a **new exception resolver**.

So, one goal for the new resolver is to enable setting a more informative response body – one that would also

conform to the type of representation requested by the client (as specified by the *Accept* header):

```

1  @Component
2  public class RestResponseStatusExceptionHandler extends AbstractHandlerExceptionResolver {
3
4      @Override
5      protected ModelAndView doResolveException
6      (HttpServletRequest request, HttpServletResponse response, Object handler, Exception ex) {
7          try {
8              if (ex instanceof IllegalArgumentException) {
9                  return handleIllegalArgument((IllegalArgumentException) ex, response, handler);
10             }
11             ...
12         } catch (ExceptionHandlerException) {
13             logger.warn("Handling of [" + ex.getClass().getName() + "]
14                 resulted in Exception", handlerException);
15         }
16         return null;
17     }
18
19     private ModelAndView handleIllegalArgument
20     (IllegalArgumentException ex, HttpServletResponse response) throws IOException {
21         response.sendError(HttpServletResponse.SC_CONFLICT);
22         String accept = request.getHeader(HttpHeaders.ACCEPT);
23         ...
24         return new ModelAndView();
25     }
26 }

```

One detail to notice here is the Request itself is available, so the application can consider the value of the *Accept* header sent by the client. For example, if the client asks for *application/json* then, in case of an error condition, the application should still return a response body encoded with *application/json*.

The other important implementation detail is that a *ModelAndView* is returned – this is the **body of the response** and it will allow the application to set whatever is necessary on it.

This approach is a consistent and easily configurable mechanism for the error handling of a Spring REST Service. It does however have **limitations**: it's interacting with the low level *HttpServletResponse* and it fits into the old MVC model which uses *ModelAndView* – so there's still room for improvement.

4. Via new **@ControllerAdvice** (Spring 3.2 Only)

Spring 3.2 brings support for a global *@ExceptionHandler* with the new *@ControllerAdvice* annotation. This enables a mechanism that breaks away from the older MVC model and makes use of *ResponseEntity* along with the type safety and flexibility of *@ExceptionHandler*:

```

1  @ControllerAdvice
2  public class RestResponseEntityExceptionHandler extends ResponseEntityExceptionHandler {
3
4      @ExceptionHandler(value = { IllegalArgumentException.class, IllegalStateException.class })
5      protected ResponseEntity<Object> handleConflict(RuntimeException ex, WebRequest request) {
6          String bodyOfResponse = "This should be application specific";
7          return handleExceptionInternal(ex, bodyOfResponse,
8              new HttpHeaders(), HttpStatus.CONFLICT, request);
9      }
10 }

```

The new annotation allows the multiple scattered *@ExceptionHandler* from before to be consolidated into a **single, global error handling component**.

The actual mechanism is extremely simple but also very flexible:

- it allows full control over the body of the response as well as the status code
- it allows mapping of several exceptions to the same method, to be handled together
- it makes good use of the newer RESTful *ResponseEntity* response

5. Conclusion

We discussed here several ways to implement an exception handling mechanism for a REST API in Spring, starting with the older mechanism and continuing with the new Spring 3.2 support.

For a full implementation of these exception handling mechanisms working in a real-world REST Service, check out the [github project](#).

XII. Versioning a REST API

1. The Problem

Evolving a REST API is a difficult problem – one for which many options are available. This section discusses through some of these options.

2. What is in the Contract?

Before anything else, we need to answer one simple question: ***What is the Contract between the API and the Client?***

2.1. URIs part of the Contract?

Let's first consider **the URI structure of the REST API** – is that part of the contract? Should clients bookmark, hardcode and generally rely on URIs of the API?

If they would, then the interaction of the Client with the REST Service would no longer be driven by the Service itself, but by what [Roy Fielding calls](#) ***out-of-band information***:

- *A REST API should be entered with no prior knowledge beyond the initial URI (bookmark) and set of standardized media types that are appropriate for the intended audience...Failure here implies that out-of-band information is driving interaction instead of hypertext.*

So clearly **URIs are not part of the contract!** The client should only know a single URI – the entry point to the API – all other URIs should be discovered while consuming the API.

2.2. Media Types part of the Contract?

What about **the Media Type information** used for the representations of Resources – are these part of the

contract between the Client and the Service?

In order to successfully consume the API, the Client **must have prior knowledge of these Media Types** – in fact, the definition of these media types represents the entire contract, and is where the REST Service should focus the most:

- *A REST API should spend almost all of its descriptive effort in defining the media type(s) used for representing resources and driving application state, or in defining extended relation names and/or hypertext-enabled mark-up for existing standard media types.*

So the **Media Type definitions are part of the contract** and should be prior knowledge for the client that consumes the API – this is where standardization comes in.

We now have a good idea of what the contract is, let's move on to how to actually tackle the versioning problem.

3. High Level Options

Let's now discuss the high level approaches to versioning the REST API:

- **URI Versioning** – version the URI space using version indicators
- **Media Type Versioning** – version the Representation of the Resource

When we introduce **the version in the URI space**, the Representations of Resources are considered immutable, so when changes need to be introduced in the API, a new URI space needs to be created.

For example, say an API publishes the following resources – users and privileges:

```
1 http://host/v1/users
2 http://host/v1/privileges
```

Now, let's consider that a breaking change in the *users* API requires **a second version to be introduced**:

```
1 http://host/v2/users
2 http://host/v2/privileges
```

When we **version the Media Type and extend the language**, we go through Content Negotiation based on this header. The REST API would make use of custom **vendor MIME media types** instead of generic media types such as *application/json*. It is these media types that are going to be versioned instead of the URIs.

For example:

```
1 ==>
2 GET /users/3 HTTP/1.1
3 Accept: application/vnd.myname.v1+json
4 <==
5 HTTP/1.1 200 OK
6 Content-Type: application/vnd.myname.v1+json
7 {
8     "user": {
9         "name": "John Smith"
```



```
10 | }  
11 | }
```

What is important to understand here is that **the client makes no assumptions about the structure of the response** beyond what is defined in the media type. This is why generic media types are not ideal – these **do not provide enough semantic information** and force the client to use require additional hints to process the actual representation of the resource.

An exception to this is using some other way of uniquely identifying the semantics of the content – such as an XML schema.

4. Advantages and Disadvantages

Now that we have a clear concept of what is part of the Contract between the Client and the Service, as well as a high level overview of the options to version the API, let's discuss the advantages and disadvantages of each approach.

First, introducing version identifiers in the URI leads to a **very large URI footprint**. This is due to the fact that any breaking change in any of the published APIs will introduce a whole new tree of representations for the entire API. Over time, this becomes a burden to maintain as well as a problem for the client – which now has more options to choose from.

Version identifiers in the URI is also **severely inflexible** – there is no way to simply evolve the API of a single Resource, or a small subset of the overall API. As we mentioned before, this is an all or nothing approach – if part of the API moves to the new version, then the entire API has to move along with it. This also makes upgrading clients from v1 to v2 a major undertaking – which leads to slower upgrades and much longer sunset periods for the old versions.

HTTP Caching is also a major concern when it comes to versioning.

From the **perspective of proxy caches in the middle**, each approach has advantages and disadvantages: if the URI is versioned, then the cache will need to keep multiple copies of each Resource – one for every version of the API. This puts load on the cache and decreases the cache hit rate, since different clients will use different versions. Also, some cache invalidation mechanisms will no longer work. If the media type is the one that is versioned, then both the Client and the Service need to support the [Vary HTTP header](#) to indicate that there are multiple versions being cached.

From the **perspective of client caching** however, the solution that versions the media type involves slightly more work than the one where URIs contain the version identifier. This is because it's simply easier to cache something when its key is an URL than a media type.

Let's end this section with defining some goals (straight out of [API Evolution](#)):

- keep compatible changes out of names
- avoid new major versions
- makes changes backwards-compatible
- think about forwards-compatibility

5. Possible Changes to the API

Next, let's consider the types of changes to the REST API – these are introduced here:

- representation format changes
- resource changes

5.1. Adding to the Representation of a Resource

The format documentation of the media type should be designed with forward compatibility in mind; specifically – a client should ignore information that it doesn't understand (which JSON does better than XML).

Now, adding information in the Representation of a resource **will not break existing clients** if these are correctly implemented.

To continue our earlier **example**, adding the *amount* in the representation of the *user* will not be a breaking change:

```

1 | {
2 |   "user": {
3 |     "name": "John Smith",
4 |     "amount": "300"
5 |   }
6 | }
```

5.2. Removing or changing an existing Representation

Removing, renaming or generally restructuring information in the design of existing representations is a breaking change for clients, because they already understand and rely on the old format.

This is where Content Negotiation comes in – for such changes, **a new vendor MIME media type** needs to be introduced.

Let's continue with the previous **example** – say we want to break the *name* of the *user* into *firstname* and *lastname*:

```

1 | ==>
2 | GET /users/3 HTTP/1.1
3 | Accept: application/vnd.myname.v2+json
4 | <==
5 | HTTP/1.1 200 OK
6 | Content-Type: application/vnd.myname.v2+json
7 | {
8 |   "user": {
9 |     "firstname": "John",
10 |    "lastname": "Smith",
11 |    "amount": "300"
12 |   }
13 | }
```

As such, this does represent an incompatible change for the Client – which will have to request the new Representation and understand the new semantics, but the URI space will remain stable and will not be affected.

5.3. Major Semantic Changes

These are changes in the meaning of the Resources, the relations between them or what they map to in the backend. This kind of changes may require a new media type, or they may require publishing a new, sibling Resource next to the old one and making use of linking to point to it.

While this sounds like using version identifiers in the URI all over again, the important distinction is that the new Resource **is published independently of any other Resources in the API** and will not fork the entire API at the root.

The REST API should adhere to **the HATEOAS constraint** – most of the URIs should be DISCOVERED by Clients, not hardcoded. Changing such an URI should not be considered an incompatible change – the new URI can replace the old one and Clients will be able to re-discover the URI and still function.

It is worth noting however that, while using version identifiers in the URI is problematic for all of these reasons, **it is not un-RESTful** in any way.

6. Conclusion

This section tried to provide an overview of the very diverse and difficult problem of **evolving a REST Service**. We discussed the two common solutions, advantages and disadvantages of each one, and ways to reason about these approaches in the context of REST. The material concludes by making the case for the second solution – **versioning the media types**, while examining the possible changes to a RESTful API.

7. Further Reading

Usually these reading resources are linked throughout the content of the section, but in these cases, there are simply too many good ones:

- [REST APIs must be hypertext-driven](#)
- [API Evolution](#)
- [Linking for a HTTP API](#)
- [Compatibility Strategies](#)

XIII. Testing REST with multiple MIME types

1. Overview

This final section will focus on **testing a REST Service with multiple Media Types/representations**. We will illustrate how to write integration tests capable of switching between the multiple types of representations that the REST API supports. The goal is to be able to run the exact same test consuming the exact same URIs of the service, just asking for a different Media Type.

2. Goals

Any REST API needs to expose its Resources as representations using some sort of Media Type, and in many cases more than a single one. **The client will set the *Accept* header** to choose the type of representation it asks for from the service.

Since the Resource can have multiple representations, the server will have to implement a mechanism responsible with choosing the right representation – also known as **Content Negotiation**. Thus, if the client asks for *application/xml*, then it should get an XML representation of the Resource, and if it asks for *application/json*, then it should get JSON.

3. Testing Infrastructure

We'll begin by defining a simple interface for a **marshaller** – this will be the main abstraction that will allow the test to switch between different Media Types:

```
1 public interface IMarshaller {
2     ...
3     String getMime();
4 }
```

Then we need a way to initialize the right marshaller based on some form of external configuration. For this mechanism, we will use a Spring *FactoryBean* to initialize the marshaller and a simple **property** to determine which marshaller to use:

```
1 @Component
2 @Profile("test")
3 public class TestMarshallerFactory implements FactoryBean<IMarshaller> {
4
5     @Autowired
6     private Environment env;
7
8     public IMarshaller getObject() {
9         String testMime = env.getProperty("test.mime");
10        if (testMime != null) {
11            switch (testMime) {
12                case "json":
13                    return new JacksonMarshaller();
14                case "xml":
15                    return new XStreamMarshaller();
16                default:
17                    throw new IllegalStateException();
18            }
19        }
20
21        return new JacksonMarshaller();
22    }
23
24    public Class<IMarshaller> getObjectType() {
25        return IMarshaller.class;
26    }
27
28    public boolean isSingleton() {
29        return true;
30    }
31 }
```

Let's look over this:

- first, the new *Environment* abstraction introduced in Spring 3.1 is used here – for more on this check out the [detailed article on using Properties with Spring](#)
- the ***test.mime* property** is retrieved from the environment and used to determine which marshaller to create

- some Java 7 *switch on String* syntax at work here
- next, the **default marshaller**, in case the property is not defined at all, is going to be the Jackson marshaller for JSON support
- finally – this *BeanFactory* is only active in a test scenario, as the new *@Profile* support, also introduced in Spring 3.1 is used

That's it – the mechanism is able to switch between marshallers based on whatever the value of the *test.mime* property is.

4. The JSON and XML Marshallers

Moving on, we'll need the actual marshaller implementation – one for each supported Media Type.

For JSON we'll use *Jackson* as the underlying library:

```

1 | public class JacksonMarshaller implements IMarshaller {
2 |     private ObjectMapper objectMapper;
3 |
4 |     public JacksonMarshaller() {
5 |         super();
6 |         objectMapper = new ObjectMapper();
7 |     }
8 |
9 |     ...
10 |
11 | @Override
12 | public String getMimeType() {
13 |     return MediaType.APPLICATION_JSON.toString();
14 | }
15 | }

```

For the XML support, the marshaller uses *XStream*:

```

1 | public class XStreamMarshaller implements IMarshaller {
2 |     private XStream xstream;
3 |
4 |     public XStreamMarshaller() {
5 |         super();
6 |         xstream = new XStream();
7 |     }
8 |
9 |     ...
10 |
11 | public String getMimeType() {
12 |     return MediaType.APPLICATION_XML.toString();
13 | }
14 | }

```

Note that these marshallers are not define as Spring components themselves. The reason for that is they will be bootstrapped into the Spring context by the *TestMarshallerFactory*, so there is no need to make them components directly.

5. Consuming the Service with both JSON and XML

At this point we should be able to run a full integration test against the deployed service. Using the marshaller is straightforward – an *IMarshaller* is simply injected directly into the test:

```

1 | @ActiveProfiles({ "test" })

```

```

2 | public abstract class SomeRestLiveTest {
3 |
4 |     @Autowired
5 |     private IMarshaller marshaller;
6 |
7 |     // tests
8 |     ...
9 |
10| }

```

The exact marshaller that will be injected by Spring will of course be decided by the value of *test.mime* property; this could be picked up from a properties file or simply set on the test environment manually. If however a value is **not provided** for this property, the *TestMarshallerFactory* will simply fall back on the default marshaller – the JSON marshaller.

6. Maven and Jenkins

If Maven is set up to run integration tests against an already deployed REST Service, then it can be run like this:

```
1 | mvn test -Dtest.mime=xml
```

Or, if this the build uses the *integration-test* phase of the Maven lifecycle:

```
1 | mvn integration-test -Dtest.mime=xml
```

For more details about how to use these phases and how to set up the a Maven build so that it will bind the deployment of the application to the *pre-integration-test* goal, run the integration tests in the *integration-test* goal and then shut down the deployed service in on *post-integration-test*, see the [Integration Testing with Maven](#) article.

With **Jenkins**, the job must be configured with:

```
1 | This build is parametrized
```

And the *String parameter*: **test.mime=xml** added.

A common Jenkins configuration would be having to jobs running the same set of integration tests against the deployed service – one with XML and the other with JSON representations.

6. Conclusion

This section showed how to properly test a REST API. Most APIs do publish their resources under multiple representations, so testing all of these representations is vital, and using the exact same tests for that is just cool.

For a full implementation of this mechanism in actual integration tests verifying both the XML and JSON representations of all Resources, check out the [github project](#).

Take Aways

At the end of this practical guide, you should be able to **implement a fully mature REST API** using Spring and Spring Security.