

Spring Data MongoDB - Reference Documentation

1.2.0.RELEASE

MarkPollack, ThomasRisberg, OliverGierke, CostinLeau, JonBrisbin

Copyright ©

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

Table of Contents

| | |
|---|----|
| Preface | v |
| I. Introduction | 1 |
| 1. Why Spring Data - Document? | 2 |
| 2. Requirements | 3 |
| 3. Additional Help Resources | 4 |
| 3.1. Support | 4 |
| Community Forum | 4 |
| Professional Support | 4 |
| 3.2. Following Development | 4 |
| 4. Repositories | 5 |
| 4.1. Introduction | 5 |
| 4.2. Core concepts | 5 |
| 4.3. Query methods | 6 |
| Defining repository interfaces | 7 |
| Fine tuning repository definition | 7 |
| Defining query methods | 8 |
| Query lookup strategies | 8 |
| Query creation | 9 |
| Special parameter handling | 10 |
| Creating repository instances | 10 |
| XML Configuration | 10 |
| JavaConfig | 11 |
| Standalone usage | 12 |
| 4.4. Custom implementations | 12 |
| Adding behaviour to single repositories | 12 |
| Adding custom behaviour to all repositories | 13 |
| 4.5. Extensions | 15 |
| Domain class web binding for Spring MVC | 15 |
| Web pagination | 18 |
| Repository populators | 19 |
| II. Reference Documentation | 21 |
| 5. MongoDB support | 22 |
| 5.1. Getting Started | 22 |
| 5.2. Examples Repository | 25 |
| 5.3. Connecting to MongoDB with Spring | 25 |
| Registering a Mongo instance using Java based metadata | 25 |
| Registering a Mongo instance using XML based metadata | 26 |
| The MongoClientFactory interface | 27 |
| Registering a MongoClientFactory instance using Java based metadata | 28 |
| Registering a MongoClientFactory instance using XML based metadata | 29 |
| 5.4. Introduction to MongoTemplate | 30 |
| Instantiating MongoTemplate | 31 |
| WriteResultChecking Policy | 32 |
| WriteConcern | 32 |
| WriteConcernResolver | 32 |
| 5.5. Saving, Updating, and Removing Documents | 33 |
| How the '_id' field is handled in the mapping layer | 35 |

| | |
|--|----|
| Type mapping | 36 |
| Methods for saving and inserting documents | 36 |
| Which collection will my documents be saved into? | 37 |
| Inserting or saving individual objects | 37 |
| Inserting several objects in a batch | 38 |
| Updating documents in a collection | 38 |
| Methods for executing updates for documents | 38 |
| Methods for the Update class | 38 |
| Upserting documents in a collection | 39 |
| Finding and Upserting documents in a collection | 39 |
| Methods for removing documents | 40 |
| 5.6. Querying Documents | 40 |
| Querying documents in a collection | 41 |
| Methods for the Criteria class | 41 |
| Methods for the Query class | 42 |
| Methods for querying for documents | 43 |
| GeoSpatial Queries | 43 |
| Geo near queries | 45 |
| 5.7. Map-Reduce Operations | 45 |
| Example Usage | 46 |
| 5.8. Group Operations | 47 |
| Example Usage | 48 |
| 5.9. Overriding default mapping with custom converters | 49 |
| Saving using a registered Spring Converter | 50 |
| Reading using a Spring Converter | 50 |
| Registering Spring Converters with the MongoConverter | 50 |
| Converter disambiguation | 51 |
| 5.10. Index and Collection management | 51 |
| Methods for creating an Index | 52 |
| Accessing index information | 52 |
| Methods for working with a Collection | 52 |
| 5.11. Executing Commands | 53 |
| Methods for executing commands | 53 |
| 5.12. Lifecycle Events | 53 |
| 5.13. Exception Translation | 54 |
| 5.14. Execution callbacks | 55 |
| 5.15. GridFS support | 55 |
| 6. MongoDB repositories | 58 |
| 6.1. Introduction | 58 |
| 6.2. Usage | 58 |
| 6.3. Query methods | 60 |
| Geo-spatial repository queries | 61 |
| MongoDB JSON based query methods and field restriction | 62 |
| Type-safe Query methods | 63 |
| 7. Mapping | 65 |
| 7.1. Convention based Mapping | 65 |
| How the '_id' field is handled in the mapping layer | 65 |
| 7.2. Mapping Configuration | 66 |
| 7.3. Metadata based Mapping | 68 |
| Mapping annotation overview | 69 |

| | |
|--|----|
| Compound Indexes | 71 |
| Using DBRefs | 71 |
| Mapping Framework Events | 72 |
| Overriding Mapping with explicit Converters | 72 |
| 8. Cross Store support | 74 |
| 8.1. Cross Store Configuration | 74 |
| 8.2. Writing the Cross Store Application | 76 |
| 9. Logging support | 79 |
| 9.1. MongoDB Log4j Configuration | 79 |
| 10. JMX support | 80 |
| 10.1. MongoDB JMX Configuration | 80 |
| III. Appendix | 82 |
| A. Namespace reference | 83 |
| A.1. The <code><repositories /></code> element | 83 |
| B. Repository query keywords | 84 |
| B.1. Supported query keywords | 84 |

Preface

The Spring Data MongoDB project applies core Spring concepts to the development of solutions using the MongoDB document style data store. We provide a "template" as a high-level abstraction for storing and querying documents. You will notice similarities to the JDBC support in the Spring Framework.

Part I. Introduction

This document is the reference guide for Spring Data - Document Support. It explains Document module concepts and semantics and the syntax for various stores namespaces.

This section provides some basic introduction to Spring and Document database. The rest of the document refers only to Spring Data Document features and assumes the user is familiar with document databases such as MongoDB and CouchDB as well as Spring concepts.

1 Knowing Spring

Spring Data uses Spring framework's [core](#) functionality, such as the [IoC](#) container, [type conversion system](#), [expression language](#), [JMX integration](#), and portable [DAO exception hierarchy](#). While it is not important to know the Spring APIs, understanding the concepts behind them is. At a minimum, the idea behind IoC should be familiar for whatever IoC container you choose to use.

The core functionality of the MongoDB and CouchDB support can be used directly, with no need to invoke the IoC services of the Spring Container. This is much like `JdbcTemplate` which can be used 'standalone' without any other services of the Spring container. To leverage all the features of Spring Data document, such as the repository support, you will need to configure some parts of the library using Spring.

To learn more about Spring, you can refer to the comprehensive (and sometimes disarming) documentation that explains in detail the Spring Framework. There are a lot of articles, blog entries and books on the matter - take a look at the Spring framework [home page](#) for more information.

2 Knowing NoSQL and Document databases

NoSQL stores have taken the storage world by storm. It is a vast domain with a plethora of solutions, terms and patterns (to make things worth even the term itself has multiple [meanings](#)). While some of the principles are common, it is crucial that the user is familiar to some degree with the stores supported by DATADOC. The best way to get acquainted to this solutions is to read their documentation and follow their examples - it usually doesn't take more then 5-10 minutes to go through them and if you are coming from an RDBMS-only background many times these exercises can be an eye opener.

The jumping off ground for learning about MongoDB is www.mongodb.org. Here is a list of other useful resources.

- The [online shell](#) provides a convenient way to interact with a MongoDB instance in combination with the online [tutorial](#).
 - MongoDB [Java Language Center](#)
 - Several [books](#) available for purchase
 - Karl Seguin's online book: "[The Little MongoDB Book](#)"
-

1. Why Spring Data - Document?

The Spring Framework is the leading full-stack Java/JEE application framework. It provides a lightweight container and a non-invasive programming model enabled by the use of dependency injection, AOP, and portable service abstractions.

[NoSQL](#) storages provide an alternative to classical RDBMS for horizontal scalability and speed. In terms of implementation, Document stores represent one of the most popular types of stores in the NoSQL space. The document database supported by Spring Data are MongoDB and CouchDB, though just MongoDB integration has been released to date.

The goal of the Spring Data Document (or DATADOC) framework is to provide an extension to the Spring programming model that supports writing applications that use Document databases. The Spring framework has always promoted a POJO programming model with a strong emphasis on portability and productivity. These values are carried over into Spring Data Document.

Notable features that are used in Spring Data Document from the Spring framework are the Features that particular, features from the Spring framework that are used are the Conversion Service, JMX Exporters, portable Data Access Exception hierarchy, Spring Expression Language, and Java based IoC container configuration. The programming model follows the familiar Spring 'template' style, so if you are familiar with Spring template classes such as JdbcTemplate, JmsTemplate, RestTemplate, you will feel right at home. For example, MongoTemplate removes much of the boilerplate code you would have to write when using the MongoDB driver to save POJOs as well as a rich java based query interface to retrieve POJOs. The programming model also offers a new Repository approach in which the Spring container will provide an implementation of a Repository based solely off an interface definition which can also include custom finder methods.

2. Requirements

Spring Data Document 1.x binaries requires JDK level 6.0 and above, and [Spring Framework](#) 3.0.x and above.

In terms of document stores, [MongoDB](#) preferably version 1.6.5 or later or [CouchDB](#) 1.0.1 or later are required.

3. Additional Help Resources

Learning a new framework is not always straight forward. In this section, we try to provide what we think is an easy to follow guide for starting with Spring Data Document module. However, if you encounter issues or you are just looking for an advice, feel free to use one of the links below:

3.1 Support

There are a few support options available:

Community Forum

The Spring Data [forum](#) is a message board for all Spring Data (not just Document) users to share information and help each other. Note that registration is needed *only* for posting.

Professional Support

Professional, from-the-source support, with guaranteed response time, is available from [SpringSource](#), the company behind Spring Data and Spring.

3.2 Following Development

For information on the Spring Data Mongo source code repository, nightly builds and snapshot artifacts please see the [Spring Data Mongo homepage](#).

You can help make Spring Data best serve the needs of the Spring community by interacting with developers through the Spring Community [forums](#). To follow developer activity look for the mailing list information on the Spring Data Mongo homepage.

If you encounter a bug or want to suggest an improvement, please create a ticket on the Spring Data issue [tracker](#).

To stay up to date with the latest news and announcements in the Spring eco system, subscribe to the Spring Community [Portal](#).

Lastly, you can follow the SpringSource Data [blog](#) or the project team on Twitter ([SpringData](#))

4. Repositories

4.1 Introduction

Implementing a data access layer of an application has been cumbersome for quite a while. Too much boilerplate code had to be written. Domain classes were anemic and not designed in a real object oriented or domain driven manner.

Using both of these technologies makes developers life a lot easier regarding rich domain model's persistence. Nevertheless the amount of boilerplate code to implement repositories especially is still quite high. So the goal of the repository abstraction of Spring Data is to reduce the effort to implement data access layers for various persistence stores significantly.

The following chapters will introduce the core concepts and interfaces of Spring Data repositories in general for detailed information on the specific features of a particular store consult the later chapters of this document.

Note

As this part of the documentation is pulled in from Spring Data Commons we have to decide for a particular module to be used as example. The configuration and code samples in this chapter are using the JPA module. Make sure you adapt e.g. the XML namespace declaration, types to be extended to the equivalents of the module you're actually using.

4.2 Core concepts

The central interface in Spring Data repository abstraction is `Repository` (probably not that much of a surprise). It is typeable to the domain class to manage as well as the id type of the domain class. This interface mainly acts as marker interface to capture the types to deal with and help us when discovering interfaces that extend this one. Beyond that there's `CrudRepository` which provides some sophisticated functionality around CRUD for the entity being managed.

```

public interface CrudRepository<T, ID extends Serializable>
    extends Repository<T, ID> {

    <S extends T> S save(S entity);

    T findOne(ID primaryKey);

    Iterable<T> findAll();

    Long count();

    void delete(T entity);

    boolean exists(ID primaryKey);

    // ... more functionality omitted.
}

```

- ❶ Saves the given entity.
- ❷ Returns the entity identified by the given id.
- ❸ Returns all entities.
- ❹ Returns the number of entities.
- ❺ Deletes the given entity.
- ❻ Returns whether an entity with the given id exists.

Example 4.1 CrudRepository interface

Usually we will have persistence technology specific sub-interfaces to include additional technology specific methods. We will now ship implementations for a variety of Spring Data modules that implement this interface.

On top of the `CrudRepository` there is a `PagingAndSortingRepository` abstraction that adds additional methods to ease paginated access to entities:

```

public interface PagingAndSortingRepository<T, ID extends Serializable> extends
    CrudRepository<T, ID> {

    Iterable<T> findAll(Sort sort);

    Page<T> findAll(Pageable pageable);

}

```

Example 4.2 PagingAndSortingRepository

Accessing the second page of `User` by a page size of 20 you could simply do something like this:

```

PagingAndSortingRepository<User, Long> repository = // ... get access to a bean
Page<User> users = repository.findAll(new PageRequest(1, 20));

```

4.3 Query methods

Next to standard CRUD functionality repositories are usually queries on the underlying datastore. With Spring Data declaring those queries becomes a four-step process:

1. Declare an interface extending `Repository` or one of its sub-interfaces and type it to the domain class it shall handle.

```
public interface PersonRepository extends Repository<User, Long> { ... }
```

2. Declare query methods on the interface.

```
List<Person> findByLastname(String lastname);
```

3. Setup Spring to create proxy instances for those interfaces.

```
<?xml version="1.0" encoding="UTF-8"?>
<beans:beans xmlns:beans="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="http://www.springframework.org/schema/data/jpa"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://www.springframework.org/schema/data/jpa
    http://www.springframework.org/schema/data/jpa/spring-jpa.xsd">

  <repositories base-package="com.acme.repositories" />

</beans>
```

Note

Note that we use the JPA namespace here just by example. If you're using the repository abstraction for any other store you need to change this to the appropriate namespace declaration of your store module which should be exchanging `jpa` in favor of e.g. `mongodb`.

4. Get the repository instance injected and use it.

```
public class SomeClient {

  @Autowired
  private PersonRepository repository;

  public void doSomething() {
    List<Person> persons = repository.findByLastname("Matthews");
  }
}
```

At this stage we barely scratched the surface of what's possible with the repositories but the general approach should be clear. Let's go through each of these steps and figure out details and various options that you have at each stage.

Defining repository interfaces

As a very first step you define a domain class specific repository interface. It's got to extend `Repository` and be typed to the domain class and an ID type. If you want to expose CRUD methods for that domain type, extend `CrudRepository` instead of `Repository`.

Fine tuning repository definition

Usually you will have your repository interface extend `Repository`, `CrudRepository` or `PagingAndSortingRepository`. If you don't like extending Spring Data interfaces at all you can also annotate your repository interface with `@RepositoryDefinition`. Extending `CrudRepository` will expose a complete set of methods to manipulate your entities. If you would rather be selective about

the methods being exposed, simply copy the ones you want to expose from `CrudRepository` into your domain repository.

```
interface MyBaseRepository<T, ID extends Serializable> extends Repository<T, ID> {
    T findOne(ID id);
    T save(T entity);
}

interface UserRepository extends MyBaseRepository<User, Long> {
    User findByEmailAddress(EmailAddress emailAddress);
}
```

Example 4.3 Selectively exposing CRUD methods

In the first step we define a common base interface for all our domain repositories and expose `findOne(...)` as well as `save(...)`. These methods will be routed into the base repository implementation of the store of your choice because they are matching the method signatures in `CrudRepository`. So our `UserRepository` will now be able to save users, find single ones by id as well as triggering a query to find `Users` by their email address.

Defining query methods

Query lookup strategies

The next thing we have to discuss is the definition of query methods. There are two main ways that the repository proxy is able to come up with the store specific query from the method name. The first option is to derive the query from the method name directly, the second is using some kind of additionally created query. What detailed options are available pretty much depends on the actual store, however, there's got to be some algorithm that decides what actual query is created.

There are three strategies available for the repository infrastructure to resolve the query. The strategy to be used can be configured at the namespace through the `query-lookup-strategy` attribute. However, it might be the case that some of the strategies are not supported for specific datastores. Here are your options:

CREATE

This strategy will try to construct a store specific query from the query method's name. The general approach is to remove a given set of well-known prefixes from the method name and parse the rest of the method. Read more about query construction in the section called "Query creation".

USE_DECLARED_QUERY

This strategy tries to find a declared query which will be used for execution first. The query could be defined by an annotation somewhere or declared by other means. Please consult the documentation of the specific store to find out what options are available for that store. If the repository infrastructure does not find a declared query for the method at bootstrap time it will fail.

CREATE_IF_NOT_FOUND (default)

This strategy is actually a combination of `CREATE` and `USE_DECLARED_QUERY`. It will try to lookup a declared query first but create a custom method name based query if no declared query was found. This is the default lookup strategy and thus will be used if you don't configure anything explicitly. It allows quick query definition by method names but also custom tuning of these queries by introducing declared queries as needed.

Query creation

The query builder mechanism built into Spring Data repository infrastructure is useful to build constraining queries over entities of the repository. We will strip the prefixes `find...By`, `read...By`, as well as `get...By` from the method and start parsing the rest of it. The introducing clause can contain further expressions such as a `Distinct` to set a distinct flag on the query to be created. However, the first `By` acts as delimiter to indicate the start of the actual criterias. At a very basic level you can define conditions on entity properties and concatenate them with `AND` and `OR`.

```
public interface PersonRepository extends Repository<User, Long> {

    List<Person> findByEmailAddressAndLastname(EmailAddress emailAddress, String lastname);

    // Enables the distinct flag for the query
    List<Person> findDistinctPeopleByLastnameOrFirstname(String lastname, String firstname);
    List<Person> findPeopleDistinctByLastnameOrFirstname(String lastname, String firstname);

    // Enabling ignoring case for an individual property
    List<Person> findByLastnameIgnoreCase(String lastname);
    // Enabling ignoring case for all suitable properties
    List<Person> findByLastnameAndFirstnameAllIgnoreCase(String lastname, String firstname);

    // Enabling static ORDER BY for a query
    List<Person> findByLastnameOrderByFirstnameAsc(String lastname);
    List<Person> findByLastnameOrderByFirstnameDesc(String lastname);
}
```

Example 4.4 Query creation from method names

The actual result of parsing that method will of course depend on the persistence store we create the query for, however, there are some general things to notice. The expressions are usually property traversals combined with operators that can be concatenated. As you can see in the example you can combine property expressions with `And` and `Or`. Beyond that you also get support for various operators like `Between`, `LessThan`, `GreaterThan`, `Like` for the property expressions. As the operators supported can vary from datastore to datastore please consult the according part of the reference documentation.

As you can see the method parser also supports setting an ignore case flag for individual properties (e.g. `findByLastnameIgnoreCase(...)`) or for all properties of a type that support ignoring case (i.e. usually `Strings`, e.g. `findByLastnameAndFirstnameAllIgnoreCase(...)`). Whether ignoring cases is supported may differ from store to store, so consult the relevant sections of the store specific query method reference docs.

Static ordering can be applied by appending an `OrderBy` clause to the query method referencing a property and providing a sorting direction (`Asc` or `Desc`). To create a query method that supports dynamic sorting have a look at the section called “Special parameter handling”.

Property expressions

Property expressions can just refer to a direct property of the managed entity (as you just saw in the example above). On query creation time we already make sure that the parsed property is at a property of the managed domain class. However, you can also define constraints by traversing nested properties. Assume `Persons` have `Addresses` with `ZipCodes`. In that case a method name of

```
List<Person> findByAddressZipCode(ZipCode zipCode);
```

will create the property traversal `x.address.zipCode`. The resolution algorithm starts with interpreting the entire part (`AddressZipCode`) as property and checks the domain class for a property with that name (uncapitalized). If it succeeds it just uses that. If not it starts splitting up the source at the camel case parts from the right side into a head and a tail and tries to find the according property, e.g. `AddressZip` and `Code`. If we find a property with that head we take the tail and continue building the tree down from there. As in our case the first split does not match we move the split point to the left (`Address, ZipCode`).

Although this should work for most cases, there might be cases where the algorithm could select the wrong property. Suppose our `Person` class has an `addressZip` property as well. Then our algorithm would match in the first split round already and essentially choose the wrong property and finally fail (as the type of `addressZip` probably has no code property). To resolve this ambiguity you can use `_` inside your method name to manually define traversal points. So our method name would end up like so:

```
List<Person> findByAddress_ZipCode(ZipCode zipCode);
```

Special parameter handling

To hand parameters to your query you simply define method parameters as already seen in the examples above. Besides that we will recognize certain specific types to apply pagination and sorting to your queries dynamically.

```
Page<User> findByLastname(String lastname, Pageable pageable);

List<User> findByLastname(String lastname, Sort sort);

List<User> findByLastname(String lastname, Pageable pageable);
```

Example 4.5 Using Pageable and Sort in query methods

The first method allows you to pass a `org.springframework.data.domain.Pageable` instance to the query method to dynamically add paging to your statically defined query. Sorting options are handed via the `Pageable` instance too. If you only need sorting, simply add an `org.springframework.data.domain.Sort` parameter to your method. As you also can see, simply returning a `List` is possible as well. We will then not retrieve the additional metadata required to build the actual `Page` instance but rather simply restrict the query to lookup only the given range of entities.

Note

To find out how many pages you get for a query entirely we have to trigger an additional count query. This will be derived from the query you actually trigger by default.

Creating repository instances

So now the question is how to create instances and bean definitions for the repository interfaces defined.

XML Configuration

The easiest way to do so is by using the Spring namespace that is shipped with each Spring Data module that supports the repository mechanism. Each of those includes a `repositories` element that allows you to simply define a base package that Spring will scan for you.

```
<?xml version="1.0" encoding="UTF-8"?>
<beans:beans xmlns:beans="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="http://www.springframework.org/schema/data/jpa"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://www.springframework.org/schema/data/jpa
    http://www.springframework.org/schema/data/jpa/spring-jpa.xsd">

  <repositories base-package="com.acme.repositories" />

</beans:beans>
```

In this case we instruct Spring to scan `com.acme.repositories` and all its sub packages for interfaces extending `Repository` or one of its sub-interfaces. For each interface found it will register the persistence technology specific `FactoryBean` to create the according proxies that handle invocations of the query methods. Each of these beans will be registered under a bean name that is derived from the interface name, so an interface of `UserRepository` would be registered under `userRepository`. The `base-package` attribute allows the use of wildcards, so that you can have a pattern of scanned packages.

Using filters

By default we will pick up every interface extending the persistence technology specific `Repository` sub-interface located underneath the configured base package and create a bean instance for it. However, you might want finer grained control over which interfaces bean instances get created for. To do this we support the use of `<include-filter />` and `<exclude-filter />` elements inside `<repositories />`. The semantics are exactly equivalent to the elements in Spring's context namespace. For details see [Spring reference documentation](#) on these elements.

E.g. to exclude certain interfaces from instantiation as repository, you could use the following configuration:

```
<repositories base-package="com.acme.repositories">
  <context:exclude-filter type="regex" expression=".*SomeRepository" />
</repositories>
```

This would exclude all interfaces ending in `SomeRepository` from being instantiated.

Example 4.6 Using exclude-filter element

JavaConfig

The repository infrastructure can also be triggered using a store-specific `@Enable ${store}Repositories` annotation on a JavaConfig class. For an introduction into Java based configuration of the Spring container please have a look at the reference documentation.²

A sample configuration to enable Spring Data repositories would look something like this.

²JavaConfig in the Spring reference documentation - <http://static.springsource.org/spring/docs/3.1.x/spring-framework-reference/html/beans.html#beans-java>


```
@Configuration
@EnableJpaRepositories("com.acme.repositories")
class ApplicationConfiguration {

    @Bean
    public EntityManagerFactory entityManagerFactory() {
        // ...
    }
}
```

Example 4.7 Sample annotation based repository configuration

Note that the sample uses the JPA specific annotation which would have to be exchanged depending on which store module you actually use. The same applies to the definition of the `EntityManagerFactory` bean. Please consult the sections covering the store-specific configuration.

Standalone usage

You can also use the repository infrastructure outside of a Spring container usage. You will still need to have some of the Spring libraries on your classpath but you can generally setup repositories programmatically as well. The Spring Data modules providing repository support ship a persistence technology specific `RepositoryFactory` that can be used as follows:

```
RepositoryFactorySupport factory = ... // Instantiate factory here
UserRepository repository = factory.getRepository(UserRepository.class);
```

Example 4.8 Standalone usage of repository factory

4.4 Custom implementations

Adding behaviour to single repositories

Often it is necessary to provide a custom implementation for a few repository methods. Spring Data repositories easily allow you to provide custom repository code and integrate it with generic CRUD abstraction and query method functionality. To enrich a repository with custom functionality you have to define an interface and an implementation for that functionality first and let the repository interface you provided so far extend that custom interface.

```
interface UserRepositoryCustom {

    public void someCustomMethod(User user);
}
```

Example 4.9 Interface for custom repository functionality

```
class UserRepositoryImpl implements UserRepositoryCustom {

    public void someCustomMethod(User user) {
        // Your custom implementation
    }
}
```

Note that the implementation itself does not depend on Spring Data and can be a regular Spring bean. So you can use standard dependency injection behaviour to inject references to other beans, take part in aspects and so on.

Example 4.10 Implementation of custom repository functionality

```
public interface UserRepository extends CrudRepository<User, Long>, UserRepositoryCustom {

    // Declare query methods here

}
```

Let your standard repository interface extend the custom one. This makes CRUD and custom functionality available to clients.

Example 4.11 Changes to the your basic repository interface

Configuration

If you use namespace configuration the repository infrastructure tries to autodetect custom implementations by looking up classes in the package we found a repository using the naming conventions appending the namespace element's attribute `repository-impl-postfix` to the classname. This suffix defaults to `Impl`.

```
<repositories base-package="com.acme.repository" />

<repositories base-package="com.acme.repository" repository-impl-postfix="FooBar" />
```

Example 4.12 Configuration example

The first configuration example will try to lookup a class `com.acme.repository.UserRepositoryImpl` to act as custom repository implementation, where the second example will try to lookup `com.acme.repository.UserRepositoryFooBar`.

Manual wiring

The approach above works perfectly well if your custom implementation uses annotation based configuration and autowiring entirely as it will be treated as any other Spring bean. If your custom implementation bean needs some special wiring you simply declare the bean and name it after the conventions just described. We will then pick up the custom bean by name rather than creating an instance.

```
<repositories base-package="com.acme.repository" />

<beans:bean id="userRepositoryImpl" class="...">
    <!-- further configuration -->
</beans:bean>
```

Example 4.13 Manual wiring of custom implementations (I)

Adding custom behaviour to all repositories

In other cases you might want to add a single method to all of your repository interfaces. So the approach just shown is not feasible. The first step to achieve this is adding an intermediate interface to declare the shared behaviour

```
public interface MyRepository<T, ID extends Serializable>
    extends JpaRepository<T, ID> {

    void sharedCustomMethod(ID id);

}
```

Example 4.14 An interface declaring custom shared behaviour

Now your individual repository interfaces will extend this intermediate interface instead of the `Repository` interface to include the functionality declared. The second step is to create an

implementation of this interface that extends the persistence technology specific repository base class which will then act as a custom base class for the repository proxies.

Note

The default behaviour of the Spring `<repositories />` namespace is to provide an implementation for all interfaces that fall under the `base-package`. This means that if left in it's current state, an implementation instance of `MyRepository` will be created by Spring. This is of course not desired as it is just supposed to act as an intermediary between `Repository` and the actual repository interfaces you want to define for each entity. To exclude an interface extending `Repository` from being instantiated as a repository instance it can either be annotated it with `@NoRepositoryBean` or moved out side of the configured `base-package`.

```
public class MyRepositoryImpl<T, ID extends Serializable>
    extends SimpleJpaRepository<T, ID> implements MyRepository<T, ID> {

    private EntityManager entityManager;

    // There are two constructors to choose from, either can be used.
    public MyRepositoryImpl(Class<T> domainClass, EntityManager entityManager) {
        super(domainClass, entityManager);

        // This is the recommended method for accessing inherited class dependencies.
        this.entityManager = entityManager;
    }

    public void sharedCustomMethod(ID id) {
        // implementation goes here
    }
}
```

Example 4.15 Custom repository base class

The last step is to create a custom repository factory to replace the default `RepositoryFactoryBean` that will in turn produce a custom `RepositoryFactory`. The new repository factory will then provide your `MyRepositoryImpl` as the implementation of any interfaces that extend the `Repository` interface, replacing the `SimpleJpaRepository` implementation you just extended.

```
public class MyRepositoryFactoryBean<R extends JpaRepository<T, I>, T, I extends
Serializable>
    extends JpaRepositoryFactoryBean<R, T, I> {

    protected RepositoryFactorySupport createRepositoryFactory(EntityManager entityManager)
    {

        return new MyRepositoryFactory(entityManager);
    }

    private static class MyRepositoryFactory<T, I extends Serializable> extends
JpaRepositoryFactory {

        private EntityManager entityManager;

        public MyRepositoryFactory(EntityManager entityManager) {
            super(entityManager);

            this.entityManager = entityManager;
        }

        protected Object getTargetRepository(RepositoryMetadata metadata) {

            return new MyRepositoryImpl<T, I>((Class<T>) metadata.getDomainClass(),
entityManager);
        }

        protected Class<?> getRepositoryBaseClass(RepositoryMetadata metadata) {

            // The RepositoryMetadata can be safely ignored, it is used by the
JpaRepositoryFactory
            //to check for QueryDslJpaRepository's which is out of scope.
            return MyRepository.class;
        }
    }
}
```

Example 4.16 Custom repository factory bean

Finally you can either declare beans of the custom factory directly or use the `factory-class` attribute of the Spring namespace to tell the repository infrastructure to use your custom factory implementation.

```
<repositories base-package="com.acme.repository"
    factory-class="com.acme.MyRepositoryFactoryBean" />
```

Example 4.17 Using the custom factory with the namespace

4.5 Extensions

This chapter documents a set of Spring Data extensions that enable Spring Data usage in a variety of contexts. Currently most of the integration is targeted towards Spring MVC.

Domain class web binding for Spring MVC

Given you are developing a Spring MVC web applications you typically have to resolve domain class ids from URLs. By default it's your task to transform that request parameter or URL part into the domain class to hand it layers below then or execute business logic on the entities directly. This should look something like this:

```
@Controller
@RequestMapping("/users")
public class UserController {

    private final UserRepository userRepository;

    @Autowired
    public UserController(UserRepository userRepository) {
        Assert.notNull(repository, "Repository must not be null!");
        userRepository = userRepository;
    }

    @RequestMapping("/{id}")
    public String showUserForm(@PathVariable("id") Long id, Model model) {

        // Do null check for id
        User user = userRepository.findOne(id);
        // Do null check for user

        model.addAttribute("user", user);
        return "user";
    }
}
```

First you pretty much have to declare a repository dependency for each controller to lookup the entity managed by the controller or repository respectively. Beyond that looking up the entity is boilerplate as well as it's always a `findOne(...)` call. Fortunately Spring provides means to register custom converting components that allow conversion between a `String` value to an arbitrary type.

PropertyEditors

For versions up to Spring 3.0 simple Java `PropertyEditors` had to be used. Thus, we offer a `DomainClassPropertyEditorRegistrar`, that will look up all Spring Data repositories registered in the `ApplicationContext` and register a custom `PropertyEditor` for the managed domain class

```
<bean class="...web.servlet.mvc.annotation.AnnotationMethodHandlerAdapter">
    <property name="webBindingInitializer">
        <bean class="...web.bind.support.ConfigurableWebBindingInitializer">
            <property name="propertyEditorRegistrars">

                <bean class="org.springframework.data.repository.support.DomainClassPropertyEditorRegistrar"
                />
            </property>
        </bean>
    </property>
</bean>
```

If you have configured Spring MVC like this you can turn your controller into the following that reduces a lot of the clutter and boilerplate.

```
@Controller
@RequestMapping("/users")
public class UserController {

    @RequestMapping("/{id}")
    public String showUserForm(@PathVariable("id") User user, Model model) {

        model.addAttribute("user", user);
        return "userForm";
    }
}
```

ConversionService

As of Spring 3.0 the `PropertyEditor` support is superseded by a new conversion infrastructure that leaves all the drawbacks of `PropertyEditors` behind and uses a stateless X to Y conversion approach. We now ship with a `DomainClassConverter` that pretty much mimics the behaviour of `DomainClassPropertyEditorRegistrar`. To configure, simply declare a bean instance and pipe the `ConversionService` being used into it's constructor:

```
<mvc:annotation-driven conversion-service="conversionService" />

<bean class="org.springframework.data.repository.support.DomainClassConverter">
    <constructor-arg ref="conversionService" />
</bean>
```

If you're using `JavaConfig` you can simply extend `WebMvcConfigurationSupport` and hand the `FormattingConversionService` the configuration superclass provides into the `DomainClassConverter` instance you create.

```
class WebConfiguration extends WebMvcConfigurationSupport {

    // Other configuration omitted

    @Bean
    public DomainClassConverter<?> domainClassConverter() {
        return new DomainClassConverter<FormattingConversionService>(mvcConversionService());
    }
}
```

Web pagination

```
@Controller
@RequestMapping("/users")
public class UserController {

    // DI code omitted

    @RequestMapping
    public String showUsers(Model model, HttpServletRequest request) {

        int page = Integer.parseInt(request.getParameter("page"));
        int pageSize = Integer.parseInt(request.getParameter("pageSize"));

        Pageable pageable = new PageRequest(page, pageSize);

        model.addAttribute("users", userService.getUsers(pageable));
        return "users";
    }
}
```

As you can see the naive approach requires the method to contain an `HttpServletRequest` parameter that has to be parsed manually. We even omitted an appropriate failure handling which would make the code even more verbose. The bottom line is that the controller actually shouldn't have to handle the functionality of extracting pagination information from the request. So we include a `PageableArgumentResolver` that will do the work for you.

```
<bean class="...web.servlet.mvc.annotation.AnnotationMethodHandlerAdapter">
  <property name="customArgumentResolvers">
    <list>
      <bean class="org.springframework.data.web.PageableArgumentResolver" />
    </list>
  </property>
</bean>
```

This configuration allows you to simplify controllers down to something like this:

```
@Controller
@RequestMapping("/users")
public class UserController {

    @RequestMapping
    public String showUsers(Model model, Pageable pageable) {

        model.addAttribute("users", userRepository.findAll(pageable));
        return "users";
    }
}
```

The `PageableArgumentResolver` will automatically resolve request parameters to build a `PageRequest` instance. By default it will expect the following structure for the request parameters:

Table 4.1. Request parameters evaluated by `PageableArgumentResolver`

| | |
|-----------|---|
| page | The page you want to retrieve |
| page.size | The size of the page you want to retrieve |

| | |
|----------------------------|---|
| <code>page.sort</code> | The property that should be sorted by |
| <code>page.sort.dir</code> | The direction that should be used for sorting |

In case you need multiple `Pageables` to be resolved from the request (for multiple tables e.g.) you can use Spring's `@Qualifier` annotation to distinguish one from another. The request parameters then have to be prefixed with `${qualifier}_`. So a method signature like this:

```
public String showUsers(Model model,
    @Qualifier("foo") Pageable first,
    @Qualifier("bar") Pageable second) { ... }
```

you'd have to populate `foo_page` and `bar_page` and the according subproperties.

Defaulting

The `PageableArgumentResolver` will use a `PageRequest` with the first page and a page size of 10 by default and will use that in case it can't resolve a `PageRequest` from the request (because of missing parameters e.g.). You can configure a global default on the bean declaration directly. In case you might need controller method specific defaults for the `Pageable` simply annotate the method parameter with `@PageableDefaults` and specify page (through `pageNumber`), page size (through `value`) as well as sort (the list of properties to sort by) as well as `sortDir` (the direction to sort by) as annotation attributes:

```
public String showUsers(Model model,
    @PageableDefaults(pageNumber = 0, value = 30) Pageable pageable) { ... }
```

Repository populators

If you have been working with the JDBC module of Spring you're probably familiar with the support to populate a `DataSource` using SQL scripts. A similar abstraction is available on the repositories level although we don't use SQL as data definition language as we need to be store independent of course. Thus the populators support XML (through Spring's OXM abstraction) and JSON (through Jackson) to define data for the repositories to be populated with.

Assume you have a file `data.json` with the following content:

```
[ { "_class" : "com.acme.Person",
  "firstname" : "Dave",
  "lastname" : "Matthews" },
  { "_class" : "com.acme.Person",
  "firstname" : "Carter",
  "lastname" : "Beauford" } ]
```

Example 4.18 Data defined in JSON

You can easily populate you repositories by using the populator elements of the repository namespace provided in Spring Data Commons. To get the just shown data be populated to your `PersonRepository` all you need to do is the following:


```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:repository="http://www.springframework.org/schema/data/repository"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://www.springframework.org/schema/data/repository
    http://www.springframework.org/schema/data/repository/spring-repository.xsd">

  <repository:jackson-populator location="classpath:data.json" />

</beans>
```

Example 4.19 Declaring a Jackson repository populator

This declaration causes the data.json file being read, deserialized by a Jackson ObjectMapper. The type the JSON object will be unmarshalled to will be determined by inspecting the `_class` attribute of the JSON document. We will eventually select the appropriate repository being able to handle the object just deserialized.

To rather use XML to define the repositories shall be populated with you can use the `unmarshaller-populator` you hand one of the marshaller options Spring OXM provides you with.

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:repository="http://www.springframework.org/schema/data/repository"
  xmlns:oxm="http://www.springframework.org/schema/oxm"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://www.springframework.org/schema/data/repository
    http://www.springframework.org/schema/data/repository/spring-repository.xsd
    http://www.springframework.org/schema/oxm
    http://www.springframework.org/schema/oxm/spring-oxm.xsd">

  <repository:unmarshaller-populator location="classpath:data.json" unmarshaller-
    ref="unmarshaller" />

  <oxm:jaxb2-marshaller contextPath="com.acme" />

</beans>
```

Example 4.20 Declaring an unmarshalling repository populator (using JAXB)

Part II. Reference Documentation

Document Structure

This part of the reference documentation explains the core functionality offered by Spring Data Document.

Chapter 5, *MongoDB support* introduces the MongoDB module feature set.

Chapter 6, *MongoDB repositories* introduces the repository support for MongoDB.

5. MongoDB support

The MongoDB support contains a wide range of features which are summarized below.

- Spring configuration support using Java based `@Configuration` classes or an XML namespace for a Mongo driver instance and replica sets
- `MongoTemplate` helper class that increases productivity performing common Mongo operations. Includes integrated object mapping between documents and POJOs.
- Exception translation into Spring's portable Data Access Exception hierarchy
- Feature Rich Object Mapping integrated with Spring's Conversion Service
- Annotation based mapping metadata but extensible to support other metadata formats
- Persistence and mapping lifecycle events
- Java based Query, Criteria, and Update DSLs
- Automatic implementation of Repository interfaces including support for custom finder methods.
- QueryDSL integration to support type-safe queries.
- Cross-store persistence - support for JPA Entities with fields transparently persisted/retrieved using MongoDB
- Log4j log appender
- GeoSpatial integration

For most tasks you will find yourself using `MongoTemplate` or the Repository support that both leverage the rich mapping functionality. `MongoTemplate` is the place to look for accessing functionality such as incrementing counters or ad-hoc CRUD operations. `MongoTemplate` also provides callback methods so that it is easy for you to get a hold of the low level API artifacts such as `org.mongodb.DB` to communicate directly with MongoDB. The goal with naming conventions on various API artifacts is to copy those in the base MongoDB Java driver so you can easily map your existing knowledge onto the Spring APIs.

5.1 Getting Started

Spring MongoDB support requires MongoDB 1.4 or higher and Java SE 5 or higher. The latest production release (2.0.x as of this writing) is recommended. An easy way to bootstrap setting up a working environment is to create a Spring based project in [STS](#).

First you need to set up a running MongoDB server. Refer to the [MongoDB Quick Start guide](#) for an explanation on how to startup a MongoDB instance. Once installed starting MongoDB is typically a matter of executing the following command: `MONGO_HOME/bin/mongod`

To create a Spring project in STS go to File -> New -> Spring Template Project -> Simple Spring Utility Project --> press Yes when prompted. Then enter a project and a package name such as `org.springframework.mongodb.example`.

Then add the following to pom.xml dependencies section.

```
<dependencies>

  <!-- other dependency elements omitted -->

  <dependency>
    <groupId>org.springframework.data</groupId>
    <artifactId>spring-data-mongodb</artifactId>
    <version>1.1.0.RELEASE</version>
  </dependency>

</dependencies>
```

Also change the version of Spring in the pom.xml to be

```
<spring.framework.version>3.1.2.RELEASE</spring.framework.version>
```

You will also need to add the location of the Spring Milestone repository for maven to your pom.xml which is at the same level of your <dependencies/> element

```
<repositories>
  <repository>
    <id>spring-milestone</id>
    <name>Spring Maven MILESTONE Repository</name>
    <url>http://repo.springsource.org/libs-milestone</url>
  </repository>
</repositories>
```

The repository is also [browseable here](#).

You may also want to set the logging level to `DEBUG` to see some additional information, edit the `log4j.properties` file to have

```
log4j.category.org.springframework.data.document.mongodb=DEBUG
log4j.appender.stdout.layout.ConversionPattern=%d{ABSOLUTE} %5p %40.40c:%4L - %m%n
```

Create a simple `Person` class to persist

```
package org.springframework.mongodb.example;

public class Person {

    private String id;
    private String name;
    private int age;

    public Person(String name, int age) {
        this.name = name;
        this.age = age;
    }

    public String getId() {
        return id;
    }
    public String getName() {
        return name;
    }
    public int getAge() {
        return age;
    }

    @Override
    public String toString() {
        return "Person [id=" + id + ", name=" + name + ", age=" + age + "]";
    }
}
```

And a main application to run

```
package org.springframework.mongodb.example;

import static org.springframework.data.mongodb.core.query.Criteria.where;

import org.apache.commons.logging.Log;
import org.apache.commons.logging.LogFactory;
import org.springframework.data.mongodb.core.MongoOperations;
import org.springframework.data.mongodb.core.MongoTemplate;
import org.springframework.data.mongodb.core.query.Query;

import com.mongodb.Mongo;

public class MongoApp {

    private static final Log log = LogFactory.getLog(MongoApp.class);

    public static void main(String[] args) throws Exception {

        MongoOperations mongoOps = new MongoTemplate(new Mongo(), "database");

        mongoOps.insert(new Person("Joe", 34));

        log.info(mongoOps.findOne(new Query(where("name").is("Joe")), Person.class));

        mongoOps.dropCollection("person");
    }
}
```

This will produce the following output

```
10:01:32,062 DEBUG apping.MongoPersistentEntityIndexCreator: 80 - Analyzing class class
org.springframework.example.Person for index information.
10:01:32,265 DEBUG framework.data.mongodb.core.MongoTemplate: 631 - insertDBObject
containing fields: [_class, age, name] in collection: Person
10:01:32,765 DEBUG framework.data.mongodb.core.MongoTemplate:1243 - findOne using query:
{ "name" : "Joe"} in db.collection: database.Person
10:01:32,953 INFO org.springframework.mongodb.example.MongoApp: 25 - Person
[id=4ddbba3c0be56b7e1b210166, name=Joe, age=34]
10:01:32,984 DEBUG framework.data.mongodb.core.MongoTemplate: 375 - Dropped collection
[database.person]
```

Even in this simple example, there are few things to take notice of

- You can instantiate the central helper class of Spring Mongo, [MongoTemplate](#), using the standard `com.mongodb.Mongo` object and the name of the database to use.
- The mapper works against standard POJO objects without the need for any additional metadata (though you can optionally provide that information. See [here](#)).
- Conventions are used for handling the id field, converting it to be a `ObjectId` when stored in the database.
- Mapping conventions can use field access. Notice the `Person` class has only getters.
- If the constructor argument names match the field names of the stored document, they will be used to instantiate the object

5.2 Examples Repository

There is an [github repository with several examples](#) that you can download and play around with to get a feel for how the library works.

5.3 Connecting to MongoDB with Spring

One of the first tasks when using MongoDB and Spring is to create a `com.mongodb.Mongo` object using the IoC container. There are two main ways to do this, either using Java based bean metadata or XML based bean metadata. These are discussed in the following sections.

Note

For those not familiar with how to configure the Spring container using Java based bean metadata instead of XML based metadata see the high level introduction in the reference docs [here](#) as well as the detailed documentation [here](#).

Registering a Mongo instance using Java based metadata

An example of using Java based bean metadata to register an instance of a `com.mongodb.Mongo` is shown below

```
@Configuration
public class AppConfig {

    /**
     * Use the standard Mongo driver API to create a com.mongodb.Mongo instance.
     */
    public @Bean Mongo mongo() throws UnknownHostException {
        return new Mongo("localhost");
    }
}
```

Example 5.1 Registering a com.mongodb.Mongo object using Java based bean metadata

This approach allows you to use the standard `com.mongodb.Mongo` API that you may already be used to using but also pollutes the code with the `UnknownHostException` checked exception. The use of the checked exception is not desirable as Java based bean metadata uses methods as a means to set object dependencies, making the calling code cluttered.

An alternative is to register an instance of `com.mongodb.Mongo` instance with the container using Spring's `MongoFactoryBean`. As compared to instantiating a `com.mongodb.Mongo` instance directly, the `FactoryBean` approach does not throw a checked exception and has the added advantage of also providing the container with an `ExceptionHandler` implementation that translates MongoDB exceptions to exceptions in Spring's portable `DataAccessException` hierarchy for data access classes annotated with the `@Repository` annotation. This hierarchy and use of `@Repository` is described in [Spring's DAO support features](#).

An example of a Java based bean metadata that supports exception translation on `@Repository` annotated classes is shown below:

```
@Configuration
public class AppConfig {

    /**
     * Factory bean that creates the com.mongodb.Mongo instance
     */
    public @Bean MongoFactoryBean mongo() {
        MongoFactoryBean mongo = new MongoFactoryBean();
        mongo.setHost("localhost");
        return mongo;
    }
}
```

To access the `com.mongodb.Mongo` object created by the `MongoFactoryBean` in other `@Configuration` or your own classes, use a `"private @Autowired Mongo mongo;"` field.

Example 5.2 Registering a com.mongodb.Mongo object using Spring's MongoFactoryBean and enabling Spring's exception translation support

Registering a Mongo instance using XML based metadata

While you can use Spring's traditional `<beans/>` XML namespace to register an instance of `com.mongodb.Mongo` with the container, the XML can be quite verbose as it is general purpose. XML namespaces are a better alternative to configuring commonly used objects such as the Mongo instance. The mongo namespace allows you to create a Mongo instance server location, replica-sets, and options.

To use the Mongo namespace elements you will need to reference the Mongo schema:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:context="http://www.springframework.org/schema/context"
       xmlns:mongo="http://www.springframework.org/schema/data/mongo"
       xsi:schemaLocation=
         "http://www.springframework.org/schema/context
         http://www.springframework.org/schema/context/spring-context-3.0.xsd
         http://www.springframework.org/schema/data/mongo
         http://www.springframework.org/schema/data/mongo/spring-mongo-1.0.xsd
         http://www.springframework.org/schema/beans
         http://www.springframework.org/schema/beans/spring-beans-3.0.xsd">

    <!-- Default bean name is 'mongo' -->
    <mongo:mongo host="localhost" port="27017"/>

</beans>
```

Example 5.3 XML schema to configure MongoDB

A more advanced configuration with `MongoOptions` is shown below (note these are not recommended values)

```
<beans>

    <mongo:mongo host="localhost" port="27017">
        <mongo:options connections-per-host="8"
            threads-allowed-to-block-for-connection-multiplier="4"
            connect-timeout="1000"
            max-wait-time="1500}"
            auto-connect-retry="true"
            socket-keep-alive="true"
            socket-timeout="1500"
            slave-ok="true"
            write-number="1"
            write-timeout="0"
            write-fsync="true"/>
    </mongo:mongo/>

</beans>
```

Example 5.4 XML schema to configure a `com.mongodb.Mongo` object with `MongoOptions`

A configuration using replica sets is shown below.

```
<mongo:mongo id="replicaSetMongo" replica-set="127.0.0.1:27017,localhost:27018"/>
```

Example 5.5 XML schema to configure `com.mongodb.Mongo` object with Replica Sets

The MongoClientFactory interface

While `com.mongodb.Mongo` is the entry point to the MongoDB driver API, connecting to a specific MongoDB database instance requires additional information such as the database name and an optional username and password. With that information you can obtain a `com.mongodb.DB` object and access all the functionality of a specific MongoDB database instance. Spring provides the `org.springframework.data.mongodb.core.MongoDbFactory` interface shown below to bootstrap connectivity to the database.


```
public interface MongoDBFactory {

    DB getDb() throws DataAccessException;

    DB getDb(String dbName) throws DataAccessException;

}
```

The following sections show how you can use the container with either Java or the XML based metadata to configure an instance of the `MongoDbFactory` interface. In turn, you can use the `MongoDbFactory` instance to configure `MongoTemplate`.

The class `org.springframework.data.mongodb.core.SimpleMongoDbFactory` provides implements the `MongoDbFactory` interface and is created with a standard `com.mongodb.Mongo` instance, the database name and an optional `org.springframework.data.authentication.UserCredentials` constructor argument.

Instead of using the IoC container to create an instance of `MongoTemplate`, you can just use them in standard Java code as shown below.

```
public class MongoApp {

    private static final Log log = LoggerFactory.getLog(MongoApp.class);

    public static void main(String[] args) throws Exception {

        MongoOperations mongoOps = new MongoTemplate(new SimpleMongoDbFactory(new Mongo(),
"database"));

        mongoOps.insert(new Person("Joe", 34));

        log.info(mongoOps.findOne(new Query(where("name").is("Joe")), Person.class));

        mongoOps.dropCollection("person");
    }

}
```

The code in bold highlights the use of `SimpleMongoDbFactory` and is the only difference between the listing shown in the [getting started section](#).

Registering a MongoDBFactory instance using Java based metadata

To register a `MongoDbFactory` instance with the container, you write code much like what was highlighted in the previous code listing. A simple example is shown below

```
@Configuration
public class MongoConfiguration {

    public @Bean MongoDBFactory mongoDbFactory() throws Exception {
        return new SimpleMongoDbFactory(new Mongo(), "database");
    }

}
```

To define the username and password create an instance of `org.springframework.data.authentication.UserCredentials` and pass it into the constructor as shown below. This listing also shows using `MongoDbFactory` register an instance of `MongoTemplate` with the container.

```
@Configuration
public class MongoConfiguration {

    public @Bean MongoClient mongoDbFactory() throws Exception {
        UserCredentials userCredentials = new UserCredentials("joe", "secret");
        return new SimpleMongoDbFactory(new MongoClient(), "database", userCredentials);
    }

    public @Bean MongoClient mongoTemplate() throws Exception {
        return new MongoClient(mongoDbFactory());
    }
}
```

Registering a MongoClientFactory instance using XML based metadata

The mongo namespace provides a convenient way to create a SimpleMongoDbFactory as compared to using the <beans/> namespace. Simple usage is shown below

```
<mongo:db-factory dbname="database">
```

In the above example a com.mongodb.Mongo instance is created using the default host and port number. The SimpleMongoDbFactory registered with the container is identified by the id 'mongoDbFactory' unless a value for the id attribute is specified.

You can also provide the host and port for the underlying com.mongodb.Mongo instance as shown below, in addition to username and password for the database.

```
<mongo:db-factory id="anotherMongoDbFactory"
    host="localhost"
    port="27017"
    dbname="database"
    username="joe"
    password="secret"/>
```

If you need to configure additional options on the com.mongodb.Mongo instance that is used to create a SimpleMongoDbFactory you can refer to an existing bean using the mongo-ref attribute as shown below. To show another common usage pattern, this listing shows the use of a property placeholder to parameterise the configuration and creating MongoClient.

```
<context:property-placeholder location="classpath:/com/myapp/mongodb/config/
mongo.properties"/>

<mongo:mongo host="${mongo.host}" port="${mongo.port}">
  <mongo:options
    connections-per-host="${mongo.connectionsPerHost}"
    threads-allowed-to-block-for-connection-
multiplier="${mongo.threadsAllowedToBlockForConnectionMultiplier}"
    connect-timeout="${mongo.connectTimeout}"
    max-wait-time="${mongo.maxWaitTime}"
    auto-connect-retry="${mongo.autoConnectRetry}"
    socket-keep-alive="${mongo.socketKeepAlive}"
    socket-timeout="${mongo.socketTimeout}"
    slave-ok="${mongo.slaveOk}"
    write-number="1"
    write-timeout="0"
    write-fsync="true"/>
</mongo:mongo>

<mongo:db-factory dbname="database" mongo-ref="mongo"/>

<bean id="anotherMongoTemplate" class="org.springframework.data.mongodb.core.MongoTemplate">
  <constructor-arg name="mongoDbFactory" ref="mongoDbFactory"/>
</bean>
```

5.4 Introduction to MongoTemplate

The class `MongoTemplate`, located in the package `org.springframework.data.document.mongodb`, is the central class of the Spring's MongoDB support providing a rich feature set to interact with the database. The template offers convenience operations to create, update, delete and query for MongoDB documents and provides a mapping between your domain objects and MongoDB documents.

Note

Once configured, `MongoTemplate` is thread-safe and can be reused across multiple instances.

The mapping between MongoDB documents and domain classes is done by delegating to an implementation of the interface `MongoConverter`. Spring provides two implementations, `SimpleMappingConverter` and `MongoMappingConverter`, but you can also write your own converter. Please refer to the section on `MongoConverters` for more detailed information.

The `MongoTemplate` class implements the interface `MongoOperations`. In as much as possible, the methods on `MongoOperations` are named after methods available on the MongoDB driver `Collection` object as to make the API familiar to existing MongoDB developers who are used to the driver API. For example, you will find methods such as "find", "findAndModify", "findOne", "insert", "remove", "save", "update" and "updateMulti". The design goal was to make it as easy as possible to transition between the use of the base MongoDB driver and `MongoOperations`. A major difference in between the two APIs is that `MongoOperations` can be passed domain objects instead of `DBObject` and there are fluent APIs for `Query`, `Criteria`, and `Update` operations instead of populating a `DBObject` to specify the parameters for those operations.

Note

The preferred way to reference the operations on `MongoTemplate` instance is via its interface `MongoOperations`.

The default converter implementation used by `MongoTemplate` is `MongoMappingConverter`. While the `MongoMappingConverter` can make use of additional metadata to specify the mapping of objects to documents it is also capable of converting objects that contain no additional metadata by using some conventions for the mapping of IDs and collection names. These conventions as well as the use of mapping annotations is explained in the [Mapping chapter](#).

Note

In the M2 release `SimpleMappingConverter`, was the default and this class is now deprecated as its functionality has been subsumed by the `MongoMappingConverter`.

Another central feature of `MongoTemplate` is exception translation of exceptions thrown in the MongoDB Java driver into Spring's portable Data Access Exception hierarchy. Refer to the section on [exception translation](#) for more information.

While there are many convenience methods on `MongoTemplate` to help you easily perform common tasks if you should need to access the MongoDB driver API directly to access functionality not explicitly exposed by the `MongoTemplate` you can use one of several Execute callback methods to access underlying driver APIs. The execute callbacks will give you a reference to either a `com.mongodb.Collection` or a `com.mongodb.DB` object. Please see the section [Execution Callbacks](#) for more information.

Now let's look at a examples of how to work with the `MongoTemplate` in the context of the Spring container.

Instantiating MongoTemplate

You can use Java to create and register an instance of `MongoTemplate` as shown below.

```
@Configuration
public class AppConfig {

    public @Bean Mongo mongo() throws Exception {
        return new Mongo("localhost");
    }

    public @Bean MongoTemplate mongoTemplate() throws Exception {
        return new MongoTemplate(mongo(), "mydatabase");
    }
}
```

Example 5.6 Registering a `com.mongodb.Mongo` object and enabling Spring's exception translation support

There are several overloaded constructors of `MongoTemplate`. These are

- **MongoTemplate** (`Mongo mongo`, `String dbName`) - takes the `com.mongodb.Mongo` object and the default database name to operate against.

- **MongoTemplate** (Mongo mongo, String databaseName, UserCredentials userCredentials) - adds the username and password for authenticating with the database.
- **MongoTemplate** (MongoDbFactory mongoDbFactory) - takes a MongoDbFactory object that encapsulated the com.mongodb.Mongo object, database name, and username and password.
- **MongoTemplate** (MongoDbFactory mongoDbFactory, MongoConverter mongoConverter) - adds a MongoConverter to use for mapping.

You can also configure a MongoTemplate using Spring's XML <beans/> schema.

```
<mongo:mongo host="localhost" port="27017"/>

<bean id="mongoTemplate" class="org.springframework.data.mongodb.core.MongoTemplate">
  <constructor-arg ref="mongo"/>
  <constructor-arg name="databaseName" value="geospatial"/>
</bean>
```

Other optional properties that you might like to set when creating a MongoTemplate are the default WriteResultCheckingPolicy, WriteConcern, and ReadPreference.

Note

The preferred way to reference the operations on MongoTemplate instance is via its interface MongoOperations.

WriteResultChecking Policy

When in development it is very handy to either log or throw an exception if the com.mongodb.WriteResult returned from any MongoDB operation contains an error. It is quite common to forget to do this during development and then end up with an application that looks like it runs successfully but in fact the database was not modified according to your expectations. Set MongoTemplate's WriteResultChecking property to an enum with the following values, LOG, EXCEPTION, or NONE to either log the error, throw and exception or do nothing. The default is to use a WriteResultChecking value of NONE.

WriteConcern

You can set the com.mongodb.WriteConcern property that the MongoTemplate will use for write operations if it has not yet been specified via the driver at a higher level such as com.mongodb.Mongo. If MongoTemplate's WriteConcern property is not set it will default to the one set in the MongoDB driver's DB or Collection setting.

WriteConcernResolver

For more advanced cases where you want to set different WriteConcern values on a per-operation basis (for remove, update, insert and save operations), a strategy interface called WriteConcernResolver can be configured on MongoTemplate. Since MongoTemplate is used to persist POJOs, the WriteConcernResolver lets you create a policy that can map a specific POJO class to a WriteConcern value. The WriteConcernResolver interface is shown below.

```
public interface WriteConcernResolver {
    WriteConcern resolve(MongoAction action);
}
```

The passed in argument, `MongoAction`, is what you use to determine the `WriteConcern` value to be used or to use the value of the Template itself as a default. `MongoAction` contains the collection name being written to, the `java.lang.Class` of the POJO, the converted `DBObject`, as well as the operation as an enumeration (`MongoActionOperation`: `REMOVE`, `UPDATE`, `INSERT`, `INSERT_LIST`, `SAVE`) and a few other pieces of contextual information. For example,

```
private class MyAppWriteConcernResolver implements WriteConcernResolver {

    public WriteConcern resolve(MongoAction action) {
        if (action.getEntityClass().getSimpleName().contains("Audit")) {
            return WriteConcern.NONE;
        } else if (action.getEntityClass().getSimpleName().contains("Metadata")) {
            return WriteConcern.JOURNAL_SAFE;
        }
        return action.getDefaultWriteConcern();
    }
}
```

5.5 Saving, Updating, and Removing Documents

`MongoTemplate` provides a simple way for you to save, update, and delete your domain objects and map those objects to documents stored in MongoDB.

Given a simple class such as `Person`

```
public class Person {

    private String id;
    private String name;
    private int age;

    public Person(String name, int age) {
        this.name = name;
        this.age = age;
    }

    public String getId() {
        return id;
    }
    public String getName() {
        return name;
    }
    public int getAge() {
        return age;
    }

    @Override
    public String toString() {
        return "Person [id=" + id + ", name=" + name + ", age=" + age + "]";
    }
}
```

You can save, update and delete the object as shown below.

Note

`MongoOperations` is the interface that `MongoTemplate` implements.

```
package org.springframework.example;

import static org.springframework.data.mongodb.core.query.Criteria.where;
import static org.springframework.data.mongodb.core.query.Update.update;
import static org.springframework.data.mongodb.core.query.Query.query;

import java.util.List;

import org.apache.commons.logging.Log;
import org.apache.commons.logging.LogFactory;
import org.springframework.data.mongodb.core.MongoOperations;
import org.springframework.data.mongodb.core.MongoTemplate;
import org.springframework.data.mongodb.core.SimpleMongoDbFactory;

import com.mongodb.Mongo;

public class MongoApp {

    private static final Log log = LogFactory.getLog(MongoApp.class);

    public static void main(String[] args) throws Exception {

        MongoOperations mongoOps = new MongoTemplate(new SimpleMongoDbFactory(new
Mongo(), "database"));

        Person p = new Person("Joe", 34);

        // Insert is used to initially store the object into the database.
        mongoOps.insert(p);
        log.info("Insert: " + p);

        // Find
        p = mongoOps.findById(p.getId(), Person.class);
        log.info("Found: " + p);

        // Update
        mongoOps.updateFirst(query(where("name").is("Joe")), update("age", 35), Person.class);

        p = mongoOps.findOne(query(where("name").is("Joe")), Person.class);
        log.info("Updated: " + p);

        // Delete
        mongoOps.remove(p);

        // Check that deletion worked
        List<Person> people = mongoOps.findAll(Person.class);
        log.info("Number of people = : " + people.size());

        mongoOps.dropCollection(Person.class);
    }
}
```

This would produce the following log output (including debug messages from MongoTemplate itself)

```
DEBUG apping.MongoPersistentEntityIndexCreator: 80 - Analyzing class class
org.springframework.example.Person for index information.
DEBUG work.data.mongodb.core.MongoTemplate: 632 - insertDBObject containing fields:
[_class, age, name] in collection: person
INFO org.springframework.example.MongoApp: 30 - Insert: Person
[id=4ddc6e784ce5b1eba3ceaf5c, name=Joe, age=34]
DEBUG work.data.mongodb.core.MongoTemplate:1246 - findOne using query: { "_id" :
{ "$oid" : "4ddc6e784ce5b1eba3ceaf5c"}} in db.collection: database.person
INFO org.springframework.example.MongoApp: 34 - Found: Person
[id=4ddc6e784ce5b1eba3ceaf5c, name=Joe, age=34]
DEBUG work.data.mongodb.core.MongoTemplate: 778 - calling update using query: { "name" :
"Joe"} and update: { "$set" : { "age" : 35}} in collection: person
DEBUG work.data.mongodb.core.MongoTemplate:1246 - findOne using query: { "name" : "Joe"}
in db.collection: database.person
INFO org.springframework.example.MongoApp: 39 - Updated: Person
[id=4ddc6e784ce5b1eba3ceaf5c, name=Joe, age=35]
DEBUG work.data.mongodb.core.MongoTemplate: 823 - remove using query: { "id" :
"4ddc6e784ce5b1eba3ceaf5c"} in collection: person
INFO org.springframework.example.MongoApp: 46 - Number of people = : 0
DEBUG work.data.mongodb.core.MongoTemplate: 376 - Dropped collection [database.person]
```

There was implicit conversion using the `MongoConverter` between a `String` and `ObjectId` as stored in the database and recognizing a convention of the property "id" name.

Note

This example is meant to show the use of save, update and remove operations on `MongoTemplate` and not to show complex mapping functionality

The query syntax used in the example is explained in more detail in the section [Querying Documents](#).

How the '_id' field is handled in the mapping layer

MongoDB requires that you have an '_id' field for all documents. If you don't provide one the driver will assign a `ObjectId` with a generated value. When using the `MongoMappingConverter` there are certain rules that govern how properties from the Java class is mapped to this '_id' field.

The following outlines what property will be mapped to the '_id' document field:

- A property or field annotated with `@Id` (`org.springframework.data.annotation.Id`) will be mapped to the '_id' field.
- A property or field without an annotation but named `id` will be mapped to the '_id' field.

The following outlines what type conversion, if any, will be done on the property mapped to the '_id' document field when using the `MappingMongoConverter`, the default for `MongoTemplate`.

- An id property or field declared as a `String` in the Java class will be converted to and stored as an `ObjectId` if possible using a `Spring Converter<String, ObjectId>`. Valid conversion rules are delegated to the MongoDB Java driver. If it cannot be converted to an `ObjectId`, then the value will be stored as a string in the database.
- An id property or field declared as `BigInteger` in the Java class will be converted to and stored as an `ObjectId` using a `Spring Converter<BigInteger, ObjectId>`.

If no field or property specified above is present in the Java class then an implicit '_id' file will be generated by the driver but not mapped to a property or field of the Java class.

When querying and updating `MongoTemplate` will use the converter to handle conversions of the `Query` and `Update` objects that correspond to the above rules for saving documents so field names and types used in your queries will be able to match what is in your domain classes.

Type mapping

As MongoDB collections can contain documents that represent instances of a variety of types. A great example here is if you store a hierarchy of classes or simply have a class with a property of type `Object`. In the latter case the values held inside that property have to be read in correctly when retrieving the object. Thus we need a mechanism to store type information alongside the actual document.

To achieve that the `MappingMongoConverter` uses a `MongoTypeMapper` abstraction with `DefaultMongoTypeMapper` as it's main implementation. It's default behaviour is storing the fully qualified classname under `_class` inside the document for the top-level document as well as for every value if it's a complex type and a subtype of the property type declared.

```
public class Sample {
    Contact value;
}

public abstract class Contact { ... }

public class Person extends Contact { ... }

Sample sample = new Sample();
sample.value = new Person();

mongoTemplate.save(sample);

{ "_class" : "com.acme.Sample",
  "value" : { "_class" : "com.acme.Person" }
}
```

Example 5.7 Type mapping

As you can see we store the type information for the actual root class persisted as well as for the nested type as it is complex and a subtype of `Contact`. So if you're now using `mongoTemplate.findAll(Object.class, "sample")` we are able to find out that the document stored shall be a `Sample` instance. We are also able to find out that the value property shall be a `Person` actually.

Customizing type mapping

In case you want to avoid writing the entire Java class name as type information but rather like to use some key you can use the `@TypeAlias` annotation at the entity class being persisted. If you need to customize the mapping even more have a look at the `TypeInformationMapper` interface. An instance of that interface can be configured at the `DefaultMongoTypeMapper` which can be configured in turn on `MappingMongoConverter`.

Methods for saving and inserting documents

There are several convenient methods on `MongoTemplate` for saving and inserting your objects. To have more fine grained control over the conversion process you can register Spring converters with the `MappingMongoConverter`, for example `Converter<Person,DBObject>` and `Converter<DBObject, Person>`.

Note

The difference between insert and save operations is that a save operation will perform an insert if the object is not already present.

The simple case of using the save operation is to save a POJO. In this case the collection name will be determined by name (not fully qualified) of the class. You may also call the save operation with a specific collection name. The collection to store the object can be overridden using mapping metadata.

When inserting or saving, if the `Id` property is not set, the assumption is that its value will be autogenerated by the database. As such, for autogeneration of an `ObjectId` to succeed the type of the `Id` property/field in your class must be either a `String`, `ObjectId`, or `BigInteger`.

Here is a basic example of using the save operation and retrieving its contents.

```
import static org.springframework.data.mongodb.core.query.Criteria.where;
import static org.springframework.data.mongodb.core.query.Criteria.query;

...

Person p = new Person("Bob", 33);
mongoTemplate.insert(p);

Person qp = mongoTemplate.findOne(query(where("age").is(33)), Person.class);
```

Example 5.8 Inserting and retrieving documents using the MongoTemplate

The insert/save operations available to you are listed below.

- **void save** (Object objectToSave) Save the object to the default collection.
- **void save** (Object objectToSave, String collectionName) Save the object to the specified collection.

A similar set of insert operations is listed below

- **void insert** (Object objectToSave) Insert the object to the default collection.
- **void insert** (Object objectToSave, String collectionName) Insert the object to the specified collection.

Which collection will my documents be saved into?

There are two ways to manage the collection name that is used for operating on the documents. The default collection name that is used is the class name changed to start with a lower-case letter. So a `com.test.Person` class would be stored in the "person" collection. You can customize this by providing a different collection name using the `@Document` annotation. You can also override the collection name by providing your own collection name as the last parameter for the selected `MongoTemplate` method calls.

Inserting or saving individual objects

The MongoDB driver supports inserting a collection of documents in one operation. The methods in the `MongoOperations` interface that support this functionality are listed below

- **insert** Insert an object. If there is an existing document with the same id then an error is generated.

- **insertAll** Takes a `Collection` of objects as the first parameter. This method inspects each object and inserts it to the appropriate collection based on the rules specified above.
- **save** Save the object overwriting any object that might exist with the same id.

Inserting several objects in a batch

The MongoDB driver supports inserting a collection of documents in one operation. The methods in the `MongoOperations` interface that support this functionality are listed below

- **insert** methods that take a `Collection` as the first argument. This inserts a list of objects in a single batch write to the database.

Updating documents in a collection

For updates we can elect to update the first document found using `MongoOperation`'s method `updateFirst` or we can update all documents that were found to match the query using the method `updateMulti`. Here is an example of an update of all SAVINGS accounts where we are adding a one time \$50.00 bonus to the balance using the `$inc` operator.

```
import static org.springframework.data.mongodb.core.query.Criteria.where;
import static org.springframework.data.mongodb.core.query.Query;
import static org.springframework.data.mongodb.core.query.Update;

...

WriteResult wr = mongoTemplate.updateMulti(new
Query(where("accounts.accountType").is(Account.Type.SAVINGS)),
new Update().inc("accounts.
$.balance", 50.00),
Account.class);
```

Example 5.9 Updating documents using the MongoTemplate

In addition to the `Query` discussed above we provide the update definition using an `Update` object. The `Update` class has methods that match the update modifiers available for MongoDB.

As you can see most methods return the `Update` object to provide a fluent style for the API.

Methods for executing updates for documents

- **updateFirst** Updates the first document that matches the query document criteria with the provided updated document.
- **updateMulti** Updates all objects that match the query document criteria with the provided updated document.

Methods for the Update class

The `Update` class can be used with a little 'syntax sugar' as its methods are meant to be chained together and you can kickstart the creation of a new `Update` instance via the static method `public static Update update(String key, Object value)` and using static imports.

Here is a listing of methods on the `Update` class

- Update **addToSet** (String key, Object value) Update using the \$addToSet update modifier
- Update **inc** (String key, Number inc) Update using the \$inc update modifier
- Update **pop** (String key, Update.Position pos) Update using the \$pop update modifier
- Update **pull** (String key, Object value) Update using the \$pull update modifier
- Update **pullAll** (String key, Object[] values) Update using the \$pullAll update modifier
- Update **push** (String key, Object value) Update using the \$push update modifier
- Update **pushAll** (String key, Object[] values) Update using the \$pushAll update modifier
- Update **rename** (String oldName, String newName) Update using the \$rename update modifier
- Update **set** (String key, Object value) Update using the \$set update modifier
- Update **unset** (String key) Update using the \$unset update modifier

Upserting documents in a collection

Related to performing an updateFirst operations, you can also perform an upsert operation which will perform an insert if no document is found that matches the query. The document that is inserted is a combination of the query document and the update document. Here is an example

```
template.upsert(query(where("ssn").is(1111).and("firstName").is("Joe").and("Fraizer").is("Update")),
    update("address", addr), Person.class);
```

Finding and Upserting documents in a collection

The findAndModify(...) method on DBCollection can update a document and return either the old or newly updated document in a single operation. MongoTemplate provides a findAndModify method that takes Query and Update classes and converts from DBObject to your POJOs. Here are the methods

```
<T> T findAndModify(Query query, Update update, Class<T> entityClass);

<T> T findAndModify(Query query, Update update, Class<T> entityClass, String
    collectionName);

<T> T findAndModify(Query query, Update update, FindAndModifyOptions options, Class<T>
    entityClass);

<T> T findAndModify(Query query, Update update, FindAndModifyOptions options, Class<T>
    entityClass, String collectionName);
```

As an example usage, we will insert of few Person objects into the container and perform a simple findAndUpdate operation

```
mongoTemplate.insert(new Person("Tom", 21));
mongoTemplate.insert(new Person("Dick", 22));
mongoTemplate.insert(new Person("Harry", 23));

Query query = new Query(Criteria.where("firstName").is("Harry"));
Update update = new Update().inc("age", 1);
Person p = mongoTemplate.findAndModify(query, update, Person.class); // return's old
    person object

assertThat(p.getFirstName(), is("Harry"));
assertThat(p.getAge(), is(23));
p = mongoTemplate.findOne(query, Person.class);
assertThat(p.getAge(), is(24));

// Now return the newly updated document when updating
p = template.findAndModify(query, update, new FindAndModifyOptions().returnNew(true),
    Person.class);
assertThat(p.getAge(), is(25));
```

The `FindAndModifyOptions` lets you set the options of `returnNew`, `upsert`, and `remove`. An example extending off the previous code snippet is shown below

```
Query query2 = new Query(Criteria.where("firstName").is("Mary"));
p = mongoTemplate.findAndModify(query2, update, new
    FindAndModifyOptions().returnNew(true).upsert(true), Person.class);
assertThat(p.getFirstName(), is("Mary"));
assertThat(p.getAge(), is(1));
```

Methods for removing documents

You can use several overloaded methods to remove an object from the database.

- **remove** Remove the given document based on one of the following: a specific object instance, a query document criteria combined with a class or a query document criteria combined with a specific collection name.

5.6 Querying Documents

You can express your queries using the `Query` and `Criteria` classes which have method names that mirror the native MongoDB operator names such as `lt`, `lte`, `is`, and others. The `Query` and `Criteria` classes follow a fluent API style so that you can easily chain together multiple method criteria and queries while having easy to understand code. Static imports in Java are used to help remove the need to see the 'new' keyword for creating `Query` and `Criteria` instances so as to improve readability. If you like to create `Query` instances from a plain JSON String use `BasicQuery`.

```
BasicQuery query = new BasicQuery("{ age : { $lt : 50 }, accounts.balance : { $gt :
    1000.00 }}");
List<Person> result = mongoTemplate.find(query, Person.class);
```

Example 5.10 Creating a Query instance from a plain JSON String

GeoSpatial queries are also supported and are described more in the section [GeoSpatial Queries](#).

Map-Reduce operations are also supported and are described more in the section [Map-Reduce](#).

Querying documents in a collection

We saw how to retrieve a single document using the `findOne` and `findById` methods on `MongoTemplate` in previous sections which return a single domain object. We can also query for a collection of documents to be returned as a list of domain objects. Assuming that we have a number of `Person` objects with name and age stored as documents in a collection and that each person has an embedded account document with a balance. We can now run a query using the following code.

```
import static org.springframework.data.mongodb.core.query.Criteria.where;
import static org.springframework.data.mongodb.core.query.Query.query;

...

List<Person> result = mongoTemplate.find(query(where("age").lt(50)
                                              .and("accounts.balance").gt(1000.00d)),
    Person.class);
```

Example 5.11 Querying for documents using the MongoTemplate

All find methods take a `Query` object as a parameter. This object defines the criteria and options used to perform the query. The criteria is specified using a `Criteria` object that has a static factory method named `where` used to instantiate a new `Criteria` object. We recommend using a static import for `org.springframework.data.mongodb.core.query.Criteria.where` and `Query.query` to make the query more readable.

This query should return a list of `Person` objects that meet the specified criteria. The `Criteria` class has the following methods that correspond to the operators provided in MongoDB.

As you can see most methods return the `Criteria` object to provide a fluent style for the API.

Methods for the Criteria class

- `Criteria all (Object o)` Creates a criterion using the `$all` operator
- `Criteria and (String key)` Adds a chained `Criteria` with the specified key to the current `Criteria` and returns the newly created one
- `Criteria andOperator (Criteria... criteria)` Creates an and query using the `$and` operator for all of the provided criteria (requires MongoDB 2.0 or later)
- `Criteria elemMatch (Criteria c)` Creates a criterion using the `$elemMatch` operator
- `Criteria exists (boolean b)` Creates a criterion using the `$exists` operator
- `Criteria gt (Object o)` Creates a criterion using the `$gt` operator
- `Criteria gte (Object o)` Creates a criterion using the `$gte` operator
- `Criteria in (Object... o)` Creates a criterion using the `$in` operator for a varargs argument.
- `Criteria in (Collection<?> collection)` Creates a criterion using the `$in` operator using a collection
- `Criteria is (Object o)` Creates a criterion using the `$is` operator
- `Criteria lt (Object o)` Creates a criterion using the `$lt` operator
- `Criteria lte (Object o)` Creates a criterion using the `$lte` operator

- Criteria **mod** (Number value, Number remainder) Creates a criterion using the `$mod` operator
- Criteria **ne** (Object o) Creates a criterion using the `$ne` operator
- Criteria **nin** (Object... o) Creates a criterion using the `$nin` operator
- Criteria **norOperator** (Criteria... criteria) Creates an `nor` query using the `$nor` operator for all of the provided criteria
- Criteria **not** () Creates a criterion using the `$not` meta operator which affects the clause directly following
- Criteria **orOperator** (Criteria... criteria) Creates an `or` query using the `$or` operator for all of the provided criteria
- Criteria **regex** (String re) Creates a criterion using a `$regex`
- Criteria **size** (int s) Creates a criterion using the `$size` operator
- Criteria **type** (int t) Creates a criterion using the `$type` operator

There are also methods on the Criteria class for geospatial queries. Here is a listing but look at the section on [GeoSpatial Queries](#) to see them in action.

- Criteria **withinCenter** (Circle circle) Creates a geospatial criterion using `$within` `$center` operators
- Criteria **withinCenterSphere** (Circle circle) Creates a geospatial criterion using `$within` `$center` operators. This is only available for MongoDB 1.7 and higher.
- Criteria **withinBox** (Box box) Creates a geospatial criterion using a `$within` `$box` operation
- Criteria **near** (Point point) Creates a geospatial criterion using a `$near` operation
- Criteria **nearSphere** (Point point) Creates a geospatial criterion using `$nearSphere` `$center` operations. This is only available for MongoDB 1.7 and higher.
- Criteria **maxDistance** (double maxDistance) Creates a geospatial criterion using the `$maxDistance` operation, for use with `$near`.

The Query class has some additional methods used to provide options for the query.

Methods for the Query class

- Query **addCriteria** (Criteria criteria) used to add additional criteria to the query
- Field **fields** () used to define fields to be included in the query results
- Query **limit** (int limit) used to limit the size of the returned results to the provided limit (used for paging)
- Query **skip** (int skip) used to skip the provided number of documents in the results (used for paging)
- Sort **sort** () used to provide sort definition for the results

Methods for querying for documents

The query methods need to specify the target type `T` that will be returned and they are also overloaded with an explicit collection name for queries that should operate on a collection other than the one indicated by the return type.

- **findAll** Query for a list of objects of type `T` from the collection.
- **findOne** Map the results of an ad-hoc query on the collection to a single instance of an object of the specified type.
- **findById** Return an object of the given id and target class.
- **find** Map the results of an ad-hoc query on the collection to a List of the specified type.
- **findAndRemove** Map the results of an ad-hoc query on the collection to a single instance of an object of the specified type. The first document that matches the query is returned and also removed from the collection in the database.

GeoSpatial Queries

MongoDB supports GeoSpatial queries through the use of operators such as `$near`, `$within`, and `$nearSphere`. Methods specific to geospatial queries are available on the `Criteria` class. There are also a few shape classes, `Box`, `Circle`, and `Point` that are used in conjunction with geospatial related `Criteria` methods.

To understand how to perform GeoSpatial queries we will use the following `Venue` class taken from the integration tests, which relies on using the rich `MappingMongoConverter`.


```
@Document(collection="newyork")
public class Venue {

    @Id
    private String id;
    private String name;
    private double[] location;

    @PersistenceConstructor
    Venue(String name, double[] location) {
        super();
        this.name = name;
        this.location = location;
    }

    public Venue(String name, double x, double y) {
        super();
        this.name = name;
        this.location = new double[] { x, y };
    }

    public String getName() {
        return name;
    }

    public double[] getLocation() {
        return location;
    }

    @Override
    public String toString() {
        return "Venue [id=" + id + ", name=" + name + ", location="
            + Arrays.toString(location) + "]";
    }
}
```

To find locations within a Circle, the following query can be used.

```
Circle circle = new Circle(-73.99171, 40.738868, 0.01);
List<Venue> venues =
    template.find(new Query(Criteria.where("location").withinCenter(circle)),
        Venue.class);
```

To find venues within a Circle using spherical coordinates the following query can be used

```
Circle circle = new Circle(-73.99171, 40.738868, 0.003712240453784);
List<Venue> venues =
    template.find(new Query(Criteria.where("location").withinCenterSphere(circle)),
        Venue.class);
```

To find venues within a Box the following query can be used

```
//lower-left then upper-right
Box box = new Box(new Point(-73.99756, 40.73083), new Point(-73.988135, 40.741404));
List<Venue> venues =
    template.find(new Query(Criteria.where("location").withinBox(box)), Venue.class);
```

To find venues near a Point, the following query can be used

```
Point point = new Point(-73.99171, 40.738868);
List<Venue> venues =
    template.find(new Query(Criteria.where("location").near(point).maxDistance(0.01)),
        Venue.class);
```

To find venues near a Point using spherical coordinates the following query can be used

```
Point point = new Point(-73.99171, 40.738868);
List<Venue> venues =
    template.find(new Query(
        Criteria.where("location").nearSphere(point).maxDistance(0.003712240453784)),
        Venue.class);
```

Geo near queries

MongoDB supports querying the database for geo locations and calculation the distance from a given origin at the very same time. With geo-near queries it's possible to express queries like: "find all restaurants in the surrounding 10 miles". To do so `MongoOperations` provides `geoNear(...)` methods taking a `NearQuery` as argument as well as the already familiar entity type and collection

```
Point location = new Point(-73.99171, 40.738868);
NearQuery query = NearQuery.near(location).maxDistance(new Distance(10, Metrics.MILES));

GeoResults<Restaurant> = operations.geoNear(query, Restaurant.class);
```

As you can see we use the `NearQuery` builder API to set up a query to return all `Restaurant` instances surrounding the given `Point` by 10 miles maximum. The `Metrics` enum used here actually implements an interface so that other metrics could be plugged into a distance as well. A `Metric` is backed by a multiplier to transform the distance value of the given metric into native distances. The sample shown here would consider the 10 to be miles. Using one of the pre-built in metrics (miles and kilometers) will automatically trigger the spherical flag to be set on the query. If you want to avoid that, simply hand in plain double values into `maxDistance(...)`. For more information see the JavaDoc of `NearQuery` and `Distance`.

The geo near operations return a `GeoResults` wrapper object that encapsulates `GeoResult` instances. The wrapping `GeoResults` allows to access the average distance of all results. A single `GeoResult` object simply carries the entity found plus its distance from the origin.

5.7 Map-Reduce Operations

You can query MongoDB using Map-Reduce which is useful for batch processing, data aggregation, and for when the query language doesn't fulfill your needs.

Spring provides integration with MongoDB's map reduce by providing methods on `MongoOperations` to simplify the creation and execution of Map-Reduce operations. It can convert the results of a Map-Reduce operation to a POJO also integrates with Spring's [Resource abstraction](#). This will let you place your JavaScript files on the file system, classpath, http server or any other Spring Resource implementation and then reference the JavaScript resources via an easy URI style syntax, e.g. 'classpath:reduce.js'. Externalizing JavaScript code in files is often preferable to embedding them as Java strings in your code. Note that you can still pass JavaScript code as Java strings if you prefer.

Example Usage

To understand how to perform Map-Reduce operations an example from the book 'MongoDB - The definitive guide' is used. In this example we will create three documents that have the values [a,b], [b,c], and [c,d] respectfully. The values in each document are associated with the key 'x' as shown below. For this example assume these documents are in the collection named "jmr1".

```
{ "_id" : ObjectId("4e5ff893c0277826074ec533"), "x" : [ "a", "b" ] }
{ "_id" : ObjectId("4e5ff893c0277826074ec534"), "x" : [ "b", "c" ] }
{ "_id" : ObjectId("4e5ff893c0277826074ec535"), "x" : [ "c", "d" ] }
```

A map function that will count the occurrence of each letter in the array for each document is shown below

```
function () {
    for (var i = 0; i < this.x.length; i++) {
        emit(this.x[i], 1);
    }
}
```

The reduce function that will sum up the occurrence of each letter across all the documents is shown below

```
function (key, values) {
    var sum = 0;
    for (var i = 0; i < values.length; i++)
        sum += values[i];
    return sum;
}
```

Executing this will result in a collection as shown below.

```
{ "_id" : "a", "value" : 1 }
{ "_id" : "b", "value" : 2 }
{ "_id" : "c", "value" : 2 }
{ "_id" : "d", "value" : 1 }
```

Assuming that the map and reduce functions are located in map.js and reduce.js and bundled in your jar so they are available on the classpath, you can execute a map-reduce operation and obtain the results as shown below

```
MapReduceResults<ValueObject> results =
    mongoOperations.mapReduce("jmr1", "classpath:map.js", "classpath:reduce.js",
        ValueObject.class);
for (ValueObject valueObject : results) {
    System.out.println(valueObject);
}
```

The output of the above code is

```
ValueObject [id=a, value=1.0]
ValueObject [id=b, value=2.0]
ValueObject [id=c, value=2.0]
ValueObject [id=d, value=1.0]
```

The MapReduceResults class implements Iterable and provides access to the raw output, as well as timing and count statistics. The ValueObject class is simply

```
public class ValueObject {

    private String id;
    private float value;

    public String getId() {
        return id;
    }

    public float getValue() {
        return value;
    }

    public void setValue(float value) {
        this.value = value;
    }

    @Override
    public String toString() {
        return "ValueObject [id=" + id + ", value=" + value + "]";
    }
}
```

By default the output type of `INLINE` is used so you don't have to specify an output collection. To specify additional map-reduce options use an overloaded method that takes an additional `MapReduceOptions` argument. The class `MapReduceOptions` has a fluent API so adding additional options can be done in a very compact syntax. Here an example that sets the output collection to `"jmr1_out"`. Note that setting only the output collection assumes a default output type of `REPLACE`.

```
MapReduceResults<ValueObject> results =
    mongoOperations.mapReduce("jmr1", "classpath:map.js", "classpath:reduce.js",
                               new
    MapReduceOptions().outputCollection("jmr1_out"), ValueObject.class);
```

There is also a static import `import org.springframework.data.mongodb.core.mapreduce.MapReduceOptions.options;` that can be used to make the syntax slightly more compact

```
MapReduceResults<ValueObject> results =
    mongoOperations.mapReduce("jmr1", "classpath:map.js", "classpath:reduce.js",
                               options().outputCollection("jmr1_out"), ValueObject.class);
```

You can also specify a query to reduce the set of data that will be used to feed into the map-reduce operation. This will remove the document that contains [a,b] from consideration for map-reduce operations.

```
Query query = new Query(where("x").ne(new String[] { "a", "b" }));
MapReduceResults<ValueObject> results =
    mongoOperations.mapReduce(query, "jmr1", "classpath:map.js", "classpath:reduce.js",
                               options().outputCollection("jmr1_out"), ValueObject.class);
```

Note that you can specify additional limit and sort values as well on the query but not skip values.

5.8 Group Operations

As an alternative to using Map-Reduce to perform data aggregation, you can use the [group operation](#) which feels similar to using SQL's group by query style, so it may feel more approachable vs. using

Map-Reduce. Using the group operations does have some limitations, for example it is not supported in a shared environment and it returns the full result set in a single BSON object, so the result should be small, less than 10,000 keys.

Spring provides integration with MongoDB's group operation by providing methods on `MongoOperations` to simplify the creation and execution of group operations. It can convert the results of the group operation to a POJO and also integrates with Spring's [Resource abstraction](#) abstraction. This will let you place your JavaScript files on the file system, classpath, http server or any other Spring Resource implementation and then reference the JavaScript resources via an easy URI style syntax, e.g. 'classpath:reduce.js;'. Externalizing JavaScript code in files is often preferable to embedding them as Java strings in your code. Note that you can still pass JavaScript code as Java strings if you prefer.

Example Usage

In order to understand how group operations work the following example is used, which is somewhat artificial. For a more realistic example consult the book 'MongoDB - The definitive guide'. A collection named "group_test_collection" created with the following rows.

```
{ "_id" : ObjectId("4ec1d25d41421e2015da64f1"), "x" : 1 }
{ "_id" : ObjectId("4ec1d25d41421e2015da64f2"), "x" : 1 }
{ "_id" : ObjectId("4ec1d25d41421e2015da64f3"), "x" : 2 }
{ "_id" : ObjectId("4ec1d25d41421e2015da64f4"), "x" : 3 }
{ "_id" : ObjectId("4ec1d25d41421e2015da64f5"), "x" : 3 }
{ "_id" : ObjectId("4ec1d25d41421e2015da64f6"), "x" : 3 }
```

We would like to group by the only field in each row, the 'x' field and aggregate the number of times each specific value of 'x' occurs. To do this we need to create an initial document that contains our count variable and also a reduce function which will increment it each time it is encountered. The Java code to execute the group operation is shown below

```
GroupByResults<XObject> results = mongoTemplate.group("group_test_collection",
    GroupBy.key("x").initialDocument("{ count: 0 }").reduceFunction("function(doc, prev)
    { prev.count += 1 }"),
    XObject.class);
```

The first argument is the name of the collection to run the group operation over, the second is a fluent API that specifies properties of the group operation via a `GroupBy` class. In this example we are using just the `initialDocument` and `reduceFunction` methods. You can also specify a key-function, as well as a finalizer as part of the fluent API. If you have multiple keys to group by, you can pass in a comma separated list of keys.

The raw results of the group operation is a JSON document that looks like this

```
{
  "retval" : [ { "x" : 1.0 , "count" : 2.0 } ,
               { "x" : 2.0 , "count" : 1.0 } ,
               { "x" : 3.0 , "count" : 3.0 } ] ,
  "count" : 6.0 ,
  "keys" : 3 ,
  "ok" : 1.0
}
```

The document under the "retval" field is mapped onto the third argument in the group method, in this case `XObject` which is shown below.

```
public class XObject {

    private float x;

    private float count;

    public float getX() {
        return x;
    }

    public void setX(float x) {
        this.x = x;
    }

    public float getCount() {
        return count;
    }

    public void setCount(float count) {
        this.count = count;
    }

    @Override
    public String toString() {
        return "XObject [x=" + x + " count = " + count + "]";
    }
}
```

You can also obtain the raw result as a `DBObject` by calling the method `getRawResults` on the `GroupByResults` class.

There is an additional method overload of the `group` method on `MongoOperations` which lets you specify a `Criteria` object for selecting a subset of the rows. An example which uses a `Criteria` object, with some syntax sugar using static imports, as well as referencing a key-function and reduce function javascript files via a Spring Resource string is shown below.

```
import static org.springframework.data.mongodb.core.mapreduce.GroupBy.keyFunction;
import static org.springframework.data.mongodb.core.query.Criteria.where;

GroupByResults<XObject> results = mongoTemplate.group(where("x").gt(0),
    "group_test_collection",

    keyFunction("classpath:keyFunction.js").initialDocument("{ count:
0 }").reduceFunction("classpath:groupReduce.js"), XObject.class);
```

5.9 Overriding default mapping with custom converters

In order to have more fine grained control over the mapping process you can register Spring converters with the `MongoConverter` implementations such as the `MappingMongoConverter`.

The `MappingMongoConverter` checks to see if there are any Spring converters that can handle a specific class before attempting to map the object itself. To 'hijack' the normal mapping strategies of the `MappingMongoConverter`, perhaps for increased performance or other custom mapping needs, you first need to create an implementation of the Spring `Converter` interface and then register it with the `MappingConverter`.

Note

For more information on the Spring type conversion service see the reference docs [here](#).

Saving using a registered Spring Converter

An example implementation of the `Converter` that converts from a `Person` object to a `com.mongodb.DBOBJECT` is shown below

```
import org.springframework.core.convert.converter.Converter;

import com.mongodb.BasicDBObject;
import com.mongodb.DBOBJECT;

public class PersonWriteConverter implements Converter<Person, DBOBJECT> {

    public DBOBJECT convert(Person source) {
        DBOBJECT dbo = new BasicDBObject();
        dbo.put("_id", source.getId());
        dbo.put("name", source.getFirstName());
        dbo.put("age", source.getAge());
        return dbo;
    }
}
```

Reading using a Spring Converter

An example implementation of a `Converter` that converts from a `DBObject` to a `Person` object is shown below

```
public class PersonReadConverter implements Converter<DBObject, Person> {

    public Person convert(DBOBJECT source) {
        Person p = new Person((ObjectId) source.get("_id"), (String) source.get("name"));
        p.setAge((Integer) source.get("age"));
        return p;
    }
}
```

Registering Spring Converters with the MongoConverter

The Mongo Spring namespace provides a convenience way to register Spring Converters with the `MappingMongoConverter`. The configuration snippet below shows how to manually register converter beans as well as configuring the wrapping `MappingMongoConverter` into a `MongoTemplate`.

```
<mongo:db-factory dbname="database"/>

<mongo:mapping-converter>
  <mongo:custom-converters>
    <mongo:converter ref="readConverter"/>
    <mongo:converter>
      <bean class="org.springframework.data.mongodb.test.PersonWriteConverter"/>
    </mongo:converter>
  </mongo:custom-converters>
</mongo:mapping-converter>

<bean id="readConverter" class="org.springframework.data.mongodb.test.PersonReadConverter"/>

<bean id="mongoTemplate" class="org.springframework.data.mongodb.core.MongoTemplate">
  <constructor-arg name="mongoDbFactory" ref="mongoDbFactory"/>
  <constructor-arg name="mongoConverter" ref="mappingConverter"/>
</bean>
```

You can also use the `base-package` attribute of the `custom-converters` element to enable classpath scanning for all `Converter` and `GenericConverter` implementations below the given package.

```
<mongo:mapping-converter>
  <mongo:custom-converters base-package="com.acme.**.converters" />
</mongo:mapping-converter>
```

Converter disambiguation

Generally we inspect the `Converter` implementations for the source and target types they convert from and to. Depending on whether one of those is a type MongoDB can handle natively we will register the converter instance as reading or writing one. Have a look at the following samples:

```
// Write converter as only the target type is one Mongo can handle natively
class MyConverter implements Converter<Person, String> { ... }

// Read converter as only the source type is one Mongo can handle natively
class MyConverter implements Converter<String, Person> { ... }
```

In case you write a `Converter` whose source and target type are native Mongo types there's no way for us to determine whether we should consider it as reading or writing converter. Registering the converter instance as both might lead to unwanted results then. E.g. a `Converter<String, Long>` is ambiguous although it probably does not make sense to try to convert all `Strings` into `Longs` when writing. To be generally able to force the infrastructure to register a converter for one way only we provide `@ReadingConverter` as well as `@WritingConverter` to be used at the converter implementation.

5.10 Index and Collection management

`MongoTemplate` provides a few methods for managing indexes and collections. These are collected into a helper interface called `IndexOperations`. You access these operations by calling the method `indexOps` and pass in either the collection name or the `java.lang.Class` of your entity (the collection name will be derived from the `.class` either by name or via annotation metadata).

The `IndexOperations` interface is shown below


```
public interface IndexOperations {

    void ensureIndex(IndexDefinition indexDefinition);

    void dropIndex(String name);

    void dropAllIndexes();

    void resetIndexCache();

    List<IndexInfo> getIndexInfo();
}
```

Methods for creating an Index

We can create an index on a collection to improve query performance.

```
mongoTemplate.indexOps(Person.class).ensureIndex(new Index().on("name", Order.ASCENDING));
```

Example 5.12 Creating an index using the MongoTemplate

- **ensureIndex** Ensure that an index for the provided IndexDefinition exists for the collection.

You can create both standard indexes and geospatial indexes using the classes IndexDefinition and GeoSpatialIndex respectfully. For example, given the Venue class defined in a previous section, you would declare a geospatial query as shown below

```
mongoTemplate.indexOps(Venue.class).ensureIndex(new GeospatialIndex("location"));
```

Accessing index information

The IndexOperations interface has the method getIndexInfo that returns a list of IndexInfo objects. This contains all the indexes defined on the collection. Here is an example that defines an index on the Person class that has age property.

```
template.indexOps(Person.class).ensureIndex(new Index().on("age",
    Order.DESCENDING).unique(Duplicates.DROP));

List<IndexInfo> indexInfoList = template.indexOps(Person.class).getIndexInfo();

// Contains
// [IndexInfo {fieldSpec={_id=ASCENDING}, name=_id, unique=false, dropDuplicates=false,
// sparse=false},
// IndexInfo {fieldSpec={age=DESCENDING}, name=age_-1, unique=true, dropDuplicates=true,
// sparse=false}]
```

Methods for working with a Collection

It's time to look at some code examples showing how to use the MongoTemplate. First we look at creating our first collection.

```
DBCollection collection = null;
if (!mongoTemplate.getCollectionNames().contains("MyNewCollection")) {
    collection = mongoTemplate.createCollection("MyNewCollection");
}

mongoTemplate.dropCollection("MyNewCollection");
```

Example 5.13 Working with collections using the MongoTemplate

- **getCollectionNames** Returns a set of collection names.
- **collectionExists** Check to see if a collection with a given name exists.
- **createCollection** Create an uncapped collection
- **dropCollection** Drop the collection
- **getCollection** Get a collection by name, creating it if it doesn't exist.

5.11 Executing Commands

You can also get at the MongoDB driver's `DB.command()` method using the `executeCommand(...)` methods on `MongoTemplate`. These will also perform exception translation into Spring's `DataAccessException` hierarchy.

Methods for executing commands

- `CommandResult executeCommand (DBObject command)` Execute a MongoDB command.
- `CommandResult executeCommand (String jsonCommand)` Execute the a MongoDB command expressed as a JSON string.

5.12 Lifecycle Events

Built into the MongoDB mapping framework are several `org.springframework.context.ApplicationEvent` events that your application can respond to by registering special beans in the `ApplicationContext`. By being based off Spring's `ApplicationContext` event infrastructure this enables other products, such as Spring Integration, to easily receive these events as they are a well known eventing mechanism in Spring based applications.

To intercept an object before it goes through the conversion process (which turns your domain object into a `com.mongodb.DBObject`), you'd register a subclass of `AbstractMongoEventListener` that overrides the `onBeforeConvert` method. When the event is dispatched, your listener will be called and passed the domain object before it goes into the converter.

```
public class BeforeConvertListener extends AbstractMongoEventListener<Person> {
    @Override
    public void onBeforeConvert(Person p) {
        ... does some auditing manipulation, set timestamps, whatever ...
    }
}
```

Example 5.14

To intercept an object before it goes into the database, you'd register a subclass of `org.springframework.data.mongodb.core.mapping.event.AbstractMongoEventListener`

that overrides the `onBeforeSave` method. When the event is dispatched, your listener will be called and passed the domain object and the converted `com.mongodb.DBObject`.

```
public class BeforeSaveListener extends AbstractMongoEventListener<Person> {
    @Override
    public void onBeforeSave(Person p, DBObject dbo) {
        ... change values, delete them, whatever ...
    }
}
```

Example 5.15

Simply declaring these beans in your Spring ApplicationContext will cause them to be invoked whenever the event is dispatched.

The list of callback methods that are present in `AbstractMappingEventListener` are

- `onBeforeConvert` - called in `MongoTemplate` `insert`, `insertList` and `save` operations before the object is converted to a `DBObject` using a `MongoConverter`.
- `onBeforeSave` - called in `MongoTemplate` `insert`, `insertList` and `save` operations *before* inserting/saving the `DBObject` in the database.
- `onAfterSave` - called in `MongoTemplate` `insert`, `insertList` and `save` operations *after* inserting/saving the `DBObject` in the database.
- `onAfterLoad` - called in `MongoTemplate` `find`, `findAndRemove`, `findOne` and `getCollection` methods after the `DBObject` is retrieved from the database.
- `onAfterConvert` - called in `MongoTemplate` `find`, `findAndRemove`, `findOne` and `getCollection` methods after the `DBObject` retrieved from the database was converted to a POJO.

5.13 Exception Translation

The Spring framework provides exception translation for a wide variety of database and mapping technologies. This has traditionally been for JDBC and JPA. The Spring support for MongoDB extends this feature to the MongoDB Database by providing an implementation of the `org.springframework.dao.support.PersistenceExceptionTranslator` interface.

The motivation behind mapping to Spring's [consistent data access exception hierarchy](#) is that you are then able to write portable and descriptive exception handling code without resorting to coding against [MongoDB error codes](#). All of Spring's data access exceptions are inherited from the root `DataAccessException` class so you can be sure that you will be able to catch all database related exception within a single try-catch block. Note, that not all exceptions thrown by the MongoDB driver inherit from the `MongoException` class. The inner exception and message are preserved so no information is lost.

Some of the mappings performed by the `MongoExceptionTranslator` are: `com.mongodb.Network` to `DataAccessResourceFailureException` and `MongoException` error codes 1003, 12001, 12010, 12011, 12012 to `InvalidDataAccessApiUsageException`. Look into the implementation for more details on the mapping.

5.14 Execution callbacks

One common design feature of all Spring template classes is that all functionality is routed into one of the templates execute callback methods. This helps ensure that exceptions and any resource management that maybe required are performed consistency. While this was of much greater need in the case of JDBC and JMS than with MongoDB, it still offers a single spot for exception translation and logging to occur. As such, using these execute callback is the preferred way to access the MongoDB driver's `DB` and `DBCollection` objects to perform uncommon operations that were not exposed as methods on `MongoTemplate`.

Here is a list of execute callback methods.

- `<T> T execute (Class<?> entityClass, CollectionCallback<T> action)` Executes the given `CollectionCallback` for the entity collection of the specified class.
- `<T> T execute (String collectionName, CollectionCallback<T> action)` Executes the given `CollectionCallback` on the collection of the given name.
- `<T> T execute (DbCallback<T> action)` Executes a `DbCallback` translating any exceptions as necessary.
- `<T> T execute (String collectionName, DbCallback<T> action)` Executes a `DbCallback` on the collection of the given name translating any exceptions as necessary.
- `<T> T executeInSession (DbCallback<T> action)` Executes the given `DbCallback` within the same connection to the database so as to ensure consistency in a write heavy environment where you may read the data that you wrote.

Here is an example that uses the `CollectionCallback` to return information about an index

```
boolean hasIndex = template.execute("geolocation", new CollectionCallbackBoolean<>() {
    public Boolean doInCollection(Venue.class, DBCollection collection) throws
        MongoException, DataAccessException {
        List<DBObject> indexes = collection.getIndexInfo();
        for (DBObject dbo : indexes) {
            if ("location_2d".equals(dbo.get("name"))) {
                return true;
            }
        }
        return false;
    }
});
```

5.15 GridFS support

MongoDB supports storing binary files inside it's filesystem GridFS. Spring Data MongoDB provides a `GridFsOperations` interface as well as the according implementation `GridFsTemplate` to easily interact with the filesystem. You can setup a `GridFsTemplate` instance by handing it a `MongoDbFactory` as well as a `MongoConverter`:

```
class GridFsConfiguration extends AbstractMongoConfiguration {

    // ... further configuration omitted

    @Bean
    public GridFsTemplate gridFsTemplate() {
        return new GridFsTemplate(mongoDbFactory(), mappingMongoConverter());
    }
}
```

Example 5.16 JavaConfig setup for a GridFsTemplate

An according XML configuration looks like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:mongo="http://www.springframework.org/schema/data/mongo"
       xsi:schemaLocation="http://www.springframework.org/schema/data/mongo
                           http://www.springframework.org/schema/data/mongo/spring-mongo.xsd
                           http://www.springframework.org/schema/beans
                           http://www.springframework.org/schema/beans/spring-beans.xsd">

    <mongo:db-factory id="mongoDbFactory" dbname="database" />
    <mongo:mapping-converter id="converter" />

    <bean class="org.springframework.data.mongodb.gridfs.GridFsTemplate">
        <constructor-arg ref="mongoDbFactory" />
        <constructor-arg ref="converter" />
    </bean>

</beans>
```

Example 5.17 XML configuration for a GridFsTemplate

You can now get the template injected and perform storing and retrieving operations to it.

```
class GridFsClient {

    @Autowired
    GridFsOperations operations;

    @Test
    public void storeFileToGridFs {

        FileMetadata metadata = new FileMetadata();
        // populate metadata
        Resource file = ... // lookup File or Resource

        operations.store(file.getInputStream(), "filename.txt", metadata);
    }
}
```

Example 5.18 Using GridFsTemplate to store files

The `store(...)` operations take an `InputStream`, a filename and optionally metadata information about the file to store. The metadata can be an arbitrary object which will be marshalled by the `MongoConverter` configured with the `GridFsTemplate`. Alternatively you can also provide a `DBObject` as well.

Reading files from the filesystem can either be achieved through the `find(...)` or `getResources(...)` methods. Let's have a look at the `find(...)` methods first. You can either find a single file matching a

Query or multiple ones. To easily define file queries we provide the `GridFsCriteria` helper class. It provides static factory methods to encapsulate default metadata fields (e.g. `whereFilename()`, `whereContentType()`) or the custom one through `whereMetaData()`.

```
class GridFsClient {

    @Autowired
    GridFsOperations operations;

    @Test
    public void findFilesInGridFs {
        List<GridFSDBFile> result = operations.find(query(whereFilename().is("filename.txt")))
    }
}
```

Example 5.19 Using GridFsTemplate to query for files

Note

Currently MongoDB does not support defining sort criterias when retrieving files from GridFS. Thus any sort criterias defined on the `Query` instance handed into the `find(...)` method will be disregarded.

The other option to read files from the GridFs is using the methods introduced by the `ResourcePatternResolver` interface. They allow handing an Ant path into the method and thus retrieve files matching the given pattern.

```
class GridFsClient {

    @Autowired
    GridFsOperations operations;

    @Test
    public void readFilesFromGridFs {
        GridFsResources[] txtFiles = operations.getResources("*.txt");
    }
}
```

Example 5.20 Using GridFsTemplate to read files

`GridFsOperations` extending `ResourcePatternResolver` allows the `GridFsTemplate` e.g. to be plugged into an `ApplicationContext` to read Spring Config files from a MongoDB.

6. MongoDB repositories

6.1 Introduction

This chapter will point out the specialties for repository support for MongoDB. This builds on the core repository support explained in Chapter 4, *Repositories*. So make sure you've got a sound understanding of the basic concepts explained there.

6.2 Usage

To access domain entities stored in a MongoDB you can leverage our sophisticated repository support that eases implementing those quite significantly. To do so, simply create an interface for your repository:

```
public class Person {

    @Id
    private String id;
    private String firstname;
    private String lastname;
    private Address address;

    // ... getters and setters omitted
}
```

Example 6.1 Sample Person entity

We have a quite simple domain object here. Note that it has a property named `id` of type `ObjectId`. The default serialization mechanism used in `MongoTemplate` (which is backing the repository support) regards properties named `id` as document id. Currently we support `String`, `ObjectId` and `BigInteger` as id-types.

```
public interface PersonRepository extends PagingAndSortingRepository<Person, Long> {

    // additional custom finder methods go here
}
```

Example 6.2 Basic repository interface to persist Person entities

Right now this interface simply serves typing purposes but we will add additional methods to it later. In your Spring configuration simply add

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:mongo="http://www.springframework.org/schema/data/mongo"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans-3.0.xsd
    http://www.springframework.org/schema/data/mongo
    http://www.springframework.org/schema/data/mongo/spring-mongo-1.0.xsd">

  <mongo:mongo id="mongo" />

  <bean id="mongoTemplate" class="org.springframework.data.mongodb.core.MongoTemplate">
    <constructor-arg ref="mongo" />
    <constructor-arg value="databaseName" />
  </bean>

  <mongo:repositories base-package="com.acme.*.repositories" />

</beans>
```

Example 6.3 General MongoDB repository Spring configuration

This namespace element will cause the base packages to be scanned for interfaces extending `MongoRepository` and create Spring beans for each of them found. By default the repositories will get a `MongoTemplate` Spring bean wired that is called `mongoTemplate`, so you only need to configure `mongo-template-ref` explicitly if you deviate from this convention.

If you'd rather like to go with JavaConfig use the `@EnableMongoRepositories` annotation. The annotation carries the very same attributes like the namespace element. If no base package is configured the infrastructure will scan the package of the annotated configuration class.

```
@Configuration
@EnableMongoRepositories
class ApplicationConfig extends AbstractMongoConfiguration {

  @Override
  protected String getDatabaseName() {
    return "e-store";
  }

  @Override
  public Mongo mongo() throws Exception {
    return new Mongo();
  }

  @Override
  protected String getMappingBasePackage() {
    return "com.oreilly.springdata.mongodb"
  }
}
```

Example 6.4 JavaConfig for repositories

As our domain repository extends `PagingAndSortingRepository` it provides you with CRUD operations as well as methods for paginated and sorted access to the entities. Working with the repository instance is just a matter of dependency injecting it into a client. So accessing the second page of `Persons` at a page size of 10 would simply look something like this:


```

@RunWith(SpringJUnit4ClassRunner.class)
@ContextConfiguration
public class PersonRepositoryTests {

    @Autowired PersonRepository repository;

    @Test
    public void readsFirstPageCorrectly() {

        Page<Person> persons = repository.findAll(new PageRequest(0, 10));
        assertThat(persons.isFirstPage(), is(true));
    }
}

```

Example 6.5 Paging access to Person entities

The sample creates an application context with Spring's unit test support which will perform annotation based dependency injection into test cases. Inside the test method we simply use the repository to query the datastore. We hand the repository a `PageRequest` instance that requests the first page of persons at a page size of 10.

6.3 Query methods

Most of the data access operations you usually trigger on a repository result a query being executed against the MongoDB databases. Defining such a query is just a matter of declaring a method on the repository interface

```

public interface PersonRepository extends PagingAndSortingRepository<Person, String> {

    List<Person> findByLastname(String lastname);

    Page<Person> findByFirstname(String firstname, Pageable pageable);

    Person findByShippingAddresses(Address address);

}

```

Example 6.6 PersonRepository with query methods

The first method shows a query for all people with the given lastname. The query will be derived parsing the method name for constraints which can be concatenated with `And` and `Or`. Thus the method name will result in a query expression of `{"lastname" : lastname}`. The second example shows how pagination is applied to a query. Just equip your method signature with a `Pageable` parameter and let the method return a `Page` instance and we will automatically page the query accordingly. The third examples shows that you can query based on properties which are not a primitive type.

Note

Note that for version 1.0 we currently don't support referring to parameters that are mapped as `DBRef` in the domain class.

Table 6.1. Supported keywords for query methods

| Keyword | Sample | Logical result |
|-------------|--|---------------------------------------|
| GreaterThan | <code>findByAgeGreaterThan(int age)</code> | <code>{"age" : {"\$gt" : age}}</code> |
| LessThan | <code>findByAgeLessThan(int age)</code> | <code>{"age" : {"\$lt" : age}}</code> |

| Keyword | Sample | Logical result |
|--------------------|---|---|
| Between | <code>findByAgeBetween(int from, int to)</code> | <code>{"age" : { "\$gt" : from, "\$lt" : to }}</code> |
| IsNotNull, NotNull | <code>findByFirstnameNotNull()</code> | <code>{"age" : { "\$ne" : null }}</code> |
| IsNull, Null | <code>findByFirstnameNull()</code> | <code>{"age" : null}</code> |
| Like | <code>findByFirstnameLike(String name)</code> | <code>{"age" : age} (age as regex)</code> |
| Regex | <code>findByFirstnameRegex(String firstname)</code> | <code>{"firstname" : { "\$regex" : firstname }}</code> |
| (No keyword) | <code>findByFirstname(String name)</code> | <code>{"age" : name}</code> |
| Not | <code>findByFirstnameNot(String name)</code> | <code>{"age" : { "\$ne" : name }}</code> |
| Near | <code>findByLocationNear(Point point)</code> | <code>{"location" : { "\$near" : [x,y] }}</code> |
| Within | <code>findByLocationWithin(Circle circle)</code> | <code>{"location" : { "\$within" : { "\$center" : [x, y], "distance" : distance } }}</code> |
| Within | <code>findByLocationWithin(Box box)</code> | <code>{"location" : { "\$within" : { "\$box" : [[x1, y1], [x2, y2]] } } True</code> |
| IsTrue, True | <code>findByActiveIsTrue()</code> | <code>{"active" : true}</code> |
| IsFalse, False | <code>findByActiveIsFalse()</code> | <code>{"active" : false}</code> |
| Exists | <code>findByLocationExists(boolean exists)</code> | <code>{"location" : { "\$exists" : exists }}</code> |

Geo-spatial repository queries

As you've just seen there are a few keywords triggering geo-spatial operations within a MongoDB query. The Near keyword allows some further modification. Let's have look at some examples:

```
public interface PersonRepository extends MongoRepository<Person, String>

    // { 'location' : { '$near' : [point.x, point.y], '$maxDistance' : distance } }
    List<Person> findByLocationNear(Point location, Distance distance);
}
```

Example 6.7 Advanced Near queries

Adding a Distance parameter to the query method allows restricting results to those within the given distance. If the Distance was set up containing a Metric we will transparently use \$nearSphere instead of \$code.

```
Point point = new Point(43.7, 48.8);
Distance distance = new Distance(200, Metrics.KILOMETERS);
... = repository.findByLocationNear(point, distance);
// {'location' : {'$nearSphere' : [43.7, 48.8], '$maxDistance' : 0.03135711885774796}}
```

Example 6.8 Using Distance with Metrics

As you can see using a `Distance` equipped with a `Metric` causes `$nearSphere` clause to be added instead of a plain `$near`. Beyond that the actual distance gets calculated according to the `Metrics` used.

Geo-near queries

```
public interface PersonRepository extends MongoRepository<Person, String>

    // {'geoNear' : 'location', 'near' : [x, y] }
    GeoResults<Person> findByLocationNear(Point location);

    // No metric: {'geoNear' : 'person', 'near' : [x, y], maxDistance : distance }
    // Metric: {'geoNear' : 'person', 'near' : [x, y], 'maxDistance' : distance,
    //          'distanceMultiplier' : metric.multiplier, 'spherical' : true }
    GeoResults<Person> findByLocationNear(Point location, Distance distance);

    // {'geoNear' : 'location', 'near' : [x, y] }
    GeoResults<Person> findByLocationNear(Point location);
}
```

MongoDB JSON based query methods and field restriction

By adding the annotation `org.springframework.data.mongodb.repository.Query` repository finder methods you can specify a MongoDB JSON query string to use instead of having the query derived from the method name. For example

```
public interface PersonRepository extends MongoRepository<Person, String>

    @Query("{ 'firstname' : ?0 }")
    List<Person> findByThePersonsFirstname(String firstname);

}
```

The placeholder `?0` lets you substitute the value from the method arguments into the JSON query string.

You can also use the `filter` property to restrict the set of properties that will be mapped into the Java object. For example,

```
public interface PersonRepository extends MongoRepository<Person, String>

    @Query(value="{ 'firstname' : ?0 }", fields="{ 'firstname' : 1, 'lastname' : 1}")
    List<Person> findByThePersonsFirstname(String firstname);

}
```

This will return only the `firstname`, `lastname` and `Id` properties of the `Person` objects. The `age` property, a `java.lang.Integer`, will not be set and its value will therefore be `null`.

Type-safe Query methods

MongoDB repository support integrates with the [QueryDSL](#) project which provides a means to perform type-safe queries in Java. To quote from the project description, "Instead of writing queries as inline strings or externalizing them into XML files they are constructed via a fluent API." It provides the following features

- Code completion in IDE (all properties, methods and operations can be expanded in your favorite Java IDE)
- Almost no syntactically invalid queries allowed (type-safe on all levels)
- Domain types and properties can be referenced safely (no Strings involved!)
- Adopts better to refactoring changes in domain types
- Incremental query definition is easier

Please refer to the QueryDSL documentation which describes how to bootstrap your environment for APT based code generation [using Maven](#) or [using Ant](#).

Using QueryDSL you will be able to write queries as shown below

```
QPerson person = new QPerson("person");
List<Person> result = repository.findAll(person.address.zipCode.eq("C0123"));

Page<Person> page = repository.findAll(person.lastname.contains("a"),
                                     new PageRequest(0, 2, Direction.ASC, "lastname"));
```

QPerson is a class that is generated (via the Java annotation post processing tool) which is a Predicate that allows you to write type safe queries. Notice that there are no strings in the query other than the value "C0123".

You can use the generated Predicate class via the interface QueryDslPredicateExecutor which is shown below

```
public interface QueryDslPredicateExecutor<T> {

    T findOne(Predicate predicate);

    List<T> findAll(Predicate predicate);

    List<T> findAll(Predicate predicate, OrderSpecifier<?>... orders);

    Page<T> findAll(Predicate predicate, Pageable pageable);

    Long count(Predicate predicate);

}
```

To use this in your repository implementation, simply inherit from it in addition to other repository interfaces. This is shown below

```
public interface PersonRepository extends MongoRepository<Person, String>,
    QueryDslPredicateExecutor<Person> {

    // additional finder methods go here

}
```

We think you will find this an extremely powerful tool for writing MongoDB queries.

7. Mapping

Rich mapping support is provided by the `MongoMappingConverter`. `MongoMappingConverter` has a rich metadata model that provides a full feature set of functionality to map domain objects to MongoDB documents. The mapping metadata model is populated using annotations on your domain objects. However, the infrastructure is not limited to using annotations as the only source of metadata information. The `MongoMappingConverter` also allows you to map objects to documents without providing any additional metadata, by following a set of conventions.

In this section we will describe the features of the `MongoMappingConverter`. How to use conventions for mapping objects to documents and how to override those conventions with annotation based mapping metadata.

Note

`SimpleMongoConverter` has been deprecated in Spring Data MongoDB M3 as all of its functionality has been subsumed into `MappingMongoConverter`.

7.1 Convention based Mapping

`MongoMappingConverter` has a few conventions for mapping objects to documents when no additional mapping metadata is provided. The conventions are:

- The short Java class name is mapped to the collection name in the following manner. The class `'com.bigbank.SavingsAccount'` maps to `'savingsAccount'` collection name.
- All nested objects are stored as nested objects in the document and **not** as DBRefs
- The converter will use any Spring Converters registered with it to override the default mapping of object properties to document field/values.
- The fields of an object are used to convert to and from fields in the document. Public JavaBean properties are not used.
- You can have a single non-zero argument constructor whose constructor argument names match top level field names of document, that constructor will be used. Otherwise the zero arg constructor will be used. if there is more than one non-zero argument constructor an exception will be thrown.

How the '_id' field is handled in the mapping layer

MongoDB requires that you have an '_id' field for all documents. If you don't provide one the driver will assign a `ObjectId` with a generated value. The "_id" field can be of any type the, other than arrays, so long as it is unique. The driver naturally supports all primitive types and Dates. When using the `MongoMappingConverter` there are certain rules that govern how properties from the Java class is mapped to this '_id' field.

The following outlines what field will be mapped to the '_id' document field:

- A field annotated with `@Id (org.springframework.data.annotation.Id)` will be mapped to the '_id' field.
- A field without an annotation but named `id` will be mapped to the '_id' field.

The following outlines what type conversion, if any, will be done on the property mapped to the `_id` document field.

- If a field named 'id' is declared as a `String` or `BigInteger` in the Java class it will be converted to and stored as an `ObjectId` if possible. `ObjectId` as a field type is also valid. If you specify a value for 'id' in your application, the conversion to an `ObjectId` is delegated to the `MongoDBDriver`. If the specified 'id' value cannot be converted to an `ObjectId`, then the value will be stored as is in the document's `_id` field.
- If a field named 'id' is not declared as a `String`, `BigInteger`, or `ObjectId` in the Java class then you should assign it a value in your application so it can be stored 'as-is' in the document's `_id` field.
- If no field named 'id' is present in the Java class then an implicit '_id' field will be generated by the driver but not mapped to a property or field of the Java class.

When querying and updating `MongoTemplate` will use the converter to handle conversions of the `Query` and `Update` objects that correspond to the above rules for saving documents so field names and types used in your queries will be able to match what is in your domain classes.

7.2 Mapping Configuration

Unless explicitly configured, an instance of `MongoMappingConverter` is created by default when creating a `MongoTemplate`. You can create your own instance of the `MappingMongoConverter` so as to tell it where to scan the classpath at startup your domain classes in order to extract metadata and construct indexes. Also, by creating your own instance you can register Spring converters to use for mapping specific classes to and from the database.

You can configure the `MongoMappingConverter` as well as `com.mongodb.Mongo` and `MongoTemplate` either using Java or XML based metadata. Here is an example using Spring's Java based configuration

```
@Configuration
public class GeoSpatialAppConfig extends AbstractMongoConfiguration {

    @Bean
    public Mongo mongo() throws Exception {
        return new Mongo("localhost");
    }

    @Override
    public String getDatabaseName() {
        return "database";
    }

    @Override
    public String getMappingBasePackage() {
        return "com.bigbank.domain";
    }

    // the following are optional

    @Override
    protected void afterMappingMongoConverterCreation(MappingMongoConverter converter) {
        Set<Converter<?, ?>> converterList = new HashSet<Converter<?, ?>>();
        converterList.add(new org.springframework.data.mongodb.test.PersonReadConverter());
        converterList.add(new org.springframework.data.mongodb.test.PersonWriteConverter());
        converter.setCustomConverters(converterList);
    }

    @Bean
    public LoggingEventListener<MongoMappingEvent> mappingEventsListener() {
        return new LoggingEventListener<MongoMappingEvent>();
    }
}
```

Example 7.1 @Configuration class to configure MongoDB mapping support

AbstractMongoConfiguration requires you to implement methods that define a com.mongodb.Mongo as well as provide a database name. AbstractMongoConfiguration also has a method you can override named 'getMappingBasePackage' which tells the converter where to scan for classes annotated with the @org.springframework.data.mongodb.core.mapping.Document annotation.

You can add additional converters to the converter by overriding the method afterMappingMongoConverterCreation. Also shown in the above example is a LoggingEventListener which logs MongoMappingEvents that are posted onto Spring's ApplicationContextEvent infrastructure.

Note

AbstractMongoConfiguration will create a MongoClient instance and registered with the container under the name 'mongoTemplate'.

You can also override the method UserCredentials getUserCredentials() to provide the username and password information to connect to the database.

Spring's MongoDB namespace enables you to easily enable mapping functionality in XML


```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:context="http://www.springframework.org/schema/context"
       xmlns:mongo="http://www.springframework.org/schema/data/mongo"
       xsi:schemaLocation="http://www.springframework.org/schema/context http://
www.springframework.org/schema/context/spring-context-3.0.xsd
http://www.springframework.org/schema/data/mongo http://
www.springframework.org/schema/data/mongo/spring-mongo-1.0.xsd
http://www.springframework.org/schema/beans http://www.springframework.org/
schema/beans/spring-beans-3.0.xsd">

  <!-- Default bean name is 'mongo' -->
  <mongo:mongo host="localhost" port="27017"/>

  <mongo:db-factory dbname="database" mongo-ref="mongo"/>

  <!-- by default look for a Mongo object named 'mongo' - default name used for the
converter is 'mappingConverter' -->
  <mongo:mapping-converter base-package="com.bigbank.domain">
    <mongo:custom-converters>
      <mongo:converter ref="readConverter"/>
      <mongo:converter>
        <bean class="org.springframework.data.mongodb.test.PersonWriteConverter"/>
      </mongo:converter>
    </mongo:custom-converters>
  </mongo:mapping-converter>

  <bean id="readConverter" class="org.springframework.data.mongodb.test.PersonReadConverter"/
>

  <!-- set the mapping converter to be used by the MongoTemplate -->
  <bean id="mongoTemplate" class="org.springframework.data.mongodb.core.MongoTemplate">
    <constructor-arg name="mongoDbFactory" ref="mongoDbFactory"/>
    <constructor-arg name="mongoConverter" ref="mappingConverter"/>
  </bean>

  <bean class="org.springframework.data.mongodb.core.mapping.event.LoggingEventListener"/>

</beans>
```

Example 7.2 XML schema to configure MongoDB mapping support

The `base-package` property tells it where to scan for classes annotated with the `@org.springframework.data.mongodb.core.mapping.Document` annotation.

7.3 Metadata based Mapping

To take full advantage of the object mapping functionality inside the Spring Data/MongoDB support, you should annotate your mapped objects with the `@org.springframework.data.mongodb.core.mapping.Document` annotation. Although it is not necessary for the mapping framework to have this annotation (your POJOs will be mapped correctly, even without any annotations), it allows the classpath scanner to find and pre-process your domain objects to extract the necessary metadata. If you don't use this annotation, your application will take a slight performance hit the first time you store a domain object because the mapping framework needs to build up its internal metadata model so it knows about the properties of your domain object and how to persist them.

```
package com.mycompany.domain;

@Document
public class Person {

    @Id
    private ObjectId id;

    @Indexed
    private Integer ssn;

    private String firstName;

    @Indexed
    private String lastName;

}
```

Example 7.3 Example domain object

Important

The `@Id` annotation tells the mapper which property you want to use for the MongoDB `_id` property and the `@Indexed` annotation tells the mapping framework to call `ensureIndex` on that property of your document, making searches faster.

Mapping annotation overview

The `MappingMongoConverter` can use metadata to drive the mapping of objects to documents. An overview of the annotations is provided below

- `@Id` - applied at the field level to mark the field used for identity purpose.
- `@Document` - applied at the class level to indicate this class is a candidate for mapping to the database. You can specify the name of the collection where the database will be stored.
- `@DBRef` - applied at the field to indicate it is to be stored using a `com.mongodb.DBRef`.
- `@Indexed` - applied at the field level to describe how to index the field.
- `@CompoundIndex` - applied at the type level to declare Compound Indexes
- `@GeoSpatialIndexed` - applied at the field level to describe how to geospatial index the field.
- `@Transient` - by default all private fields are mapped to the document, this annotation excludes the field where it is applied from being stored in the database
- `@PersistenceConstructor` - marks a given constructor - even a package protected one - to use when instantiating the object from the database. Constructor arguments are mapped by name to the key values in the retrieved `DBObject`.
- `@Value` - this annotation is part of the Spring Framework . Within the mapping framework it can be applied to constructor arguments. This lets you use a Spring Expression Language statement to transform a key's value retrieved in the database before it is used to construct a domain object.
- `@Field` - applied at the field level and describes the name of the field as it will be represented in the MongoDB BSON document thus allowing the name to be different than the fieldname of the class.

The mapping metadata infrastructure is defined in a separate spring-data-commons project that is technology agnostic. Specific subclasses are using in the MongoDB support to support annotation based metadata. Other strategies are also possible to put in place if there is demand.

Here is an example of a more complex mapping.

```
@Document
@CompoundIndexes({
    @CompoundIndex(name = "age_idx", def = "{ 'lastName': 1, 'age': -1}")
})
public class Person<T extends Address> {

    @Id
    private String id;

    @Indexed(unique = true)
    private Integer ssn;

    @Field("fName")
    private String firstName;

    @Indexed
    private String lastName;

    private Integer age;

    @Transient
    private Integer accountTotal;

    @DBRef
    private List<Account> accounts;

    private T address;

    public Person(Integer ssn) {
        this.ssn = ssn;
    }

    @PersistenceConstructor
    public Person(Integer ssn, String firstName, String lastName, Integer age, T address) {
        this.ssn = ssn;
        this.firstName = firstName;
        this.lastName = lastName;
        this.age = age;
        this.address = address;
    }

    public String getId() {
        return id;
    }

    // no setter for Id. (getter is only exposed for some unit testing)

    public Integer getSsn() {
        return ssn;
    }

    // other getters/setters omitted
}
```

Compound Indexes

Compound indexes are also supported. They are defined at the class level, rather than on individual properties.

Note

Compound indexes are very important to improve the performance of queries that involve criteria on multiple fields

Here's an example that creates a compound index of `lastName` in ascending order and `age` in descending order:

```
package com.mycompany.domain;

@Document
@CompoundIndexes({
    @CompoundIndex(name = "age_idx", def = "{ 'lastName': 1, 'age': -1 }")
})
public class Person {

    @Id
    private ObjectId id;
    private Integer age;
    private String firstName;
    private String lastName;

}
```

Example 7.4 Example Compound Index Usage

Using DBRefs

The mapping framework doesn't have to store child objects embedded within the document. You can also store them separately and use a `DBRef` to refer to that document. When the object is loaded from MongoDB, those references will be eagerly resolved and you will get back a mapped object that looks the same as if it had been stored embedded within your master document.

Here's an example of using a `DBRef` to refer to a specific document that exists independently of the object in which it is referenced (both classes are shown in-line for brevity's sake):

```
@Document
public class Account {

    @Id
    private ObjectId id;
    private Float total;

}

@Document
public class Person {

    @Id
    private ObjectId id;
    @Indexed
    private Integer ssn;
    @DBRef
    private List<Account> accounts;

}
```

Example 7.5

There's no need to use something like `@OneToMany` because the mapping framework sees that you're wanting a one-to-many relationship because there is a `List` of objects. When the object is stored in MongoDB, there will be a list of `DBRefs` rather than the `Account` objects themselves.

Important

The mapping framework does not handle cascading saves. If you change an `Account` object that is referenced by a `Person` object, you must save the `Account` object separately. Calling `save` on the `Person` object will not automatically save the `Account` objects in the property `accounts`.

Mapping Framework Events

Events are fired throughout the lifecycle of the mapping process. This is described in the [Lifecycle Events](#) section.

Simply declaring these beans in your Spring `ApplicationContext` will cause them to be invoked whenever the event is dispatched.

Overriding Mapping with explicit Converters

When storing and querying your objects it is convenient to have a `MongoConverter` instance handle the mapping of all Java types to `DBObject`s. However, sometimes you may want the `MongoConverter`'s do most of the work but allow you to selectively handle the conversion for a particular type or to optimize performance.

To selectively handle the conversion yourself, register one or more `org.springframework.core.convert.converter.Converter` instances with the `MongoConverter`.

Note

Spring 3.0 introduced a `core.convert` package that provides a general type conversion system. This is described in detail in the Spring reference documentation section entitled [Spring 3 Type Conversion](#).

The `setConverters` method on `SimpleMongoConverter` and `MappingMongoConverter` should be used for this purpose. The method `afterMappingMongoConverterCreation` in `AbstractMongoConfiguration` can be overridden to configure a `MappingMongoConverter`. The examples [here](#) at the beginning of this chapter show how to perform the configuration using Java and XML.

Below is an example of a Spring Converter implementation that converts from a `DBObject` to a `Person` POJO.

```
public class PersonReadConverter implements Converter<DBObject, Person> {

    public Person convert(DBObject source) {
        Person p = new Person((ObjectId) source.get("_id"), (String) source.get("name"));
        p.setAge((Integer) source.get("age"));
        return p;
    }
}
```

Here is an example that converts from a `Person` to a `DBObject`.

```
public class PersonWriteConverter implements Converter<Person, DBObject> {

    public DBObject convert(Person source) {
        DBObject dbo = new BasicDBObject();
        dbo.put("_id", source.getId());
        dbo.put("name", source.getFirstName());
        dbo.put("age", source.getAge());
        return dbo;
    }
}
```

8. Cross Store support

Sometimes you need to store data in multiple data stores and these data stores can be of different types. One might be relational while the other a document store. For this use case we have created a separate module in the MongoDB support that handles what we call cross-store support. The current implementation is based on JPA as the driver for the relational database and we allow select fields in the Entities to be stored in a Mongo database. In addition to allowing you to store your data in two stores we also coordinate persistence operations for the non-transactional MongoDB store with the transaction life-cycle for the relational database.

8.1 Cross Store Configuration

Assuming that you have a working JPA application and would like to add some cross-store persistence for MongoDB. What do you have to add to your configuration?

First of all you need to add a dependency on the `spring-data-mongodb-cross-store` module. Using Maven this is done by adding a dependency to your pom:

```
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/
XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/
maven-4.0.0.xsd">
  <modelVersion>4.0.0</modelVersion>

  ...

  <!-- Spring Data -->
  <dependency>
    <groupId>org.springframework.data</groupId>
    <artifactId>spring-data-mongodb-cross-store</artifactId>
    <version>${spring.data.mongo.version}</version>
  </dependency>

  ...

</project>
```

Example 8.1 Example Maven pom.xml with spring-data-mongodb-cross-store dependency

Once this is done we need to enable AspectJ for the project. The cross-store support is implemented using AspectJ aspects so by enabling compile time AspectJ support the cross-store features will become available to your project. In Maven you would add an additional plugin to the `<build>` section of the pom:

```

<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/
XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/
maven-4.0.0.xsd">
  <modelVersion>4.0.0</modelVersion>

  ...

  <build>
    <plugins>

      ...

      <plugin>
        <groupId>org.codehaus.mojo</groupId>
        <artifactId>aspectj-maven-plugin</artifactId>
        <version>1.0</version>
        <dependencies>
          <!-- NB: You must use Maven 2.0.9 or above or these are ignored (see MNG-2972)
-->
          <dependency>
            <groupId>org.aspectj</groupId>
            <artifactId>aspectjrt</artifactId>
            <version>${aspectj.version}</version>
          </dependency>
          <dependency>
            <groupId>org.aspectj</groupId>
            <artifactId>aspectjtools</artifactId>
            <version>${aspectj.version}</version>
          </dependency>
        </dependencies>
        <executions>
          <execution>
            <goals>
              <goal>compile</goal>
              <goal>test-compile</goal>
            </goals>
          </execution>
        </executions>
        <configuration>
          <outxml>true</outxml>
          <aspectLibraries>
            <aspectLibrary>
              <groupId>org.springframework</groupId>
              <artifactId>spring-aspects</artifactId>
            </aspectLibrary>
            <aspectLibrary>
              <groupId>org.springframework.data</groupId>
              <artifactId>spring-data-mongodb-cross-store</artifactId>
            </aspectLibrary>
          </aspectLibraries>
          <source>1.6</source>
          <target>1.6</target>
        </configuration>
      </plugin>

      ...

    </plugins>
  </build>

  ...
</project>

```


Finally, you need to configure your project to use MongoDB and also configure the aspects that are used. The following XML snippet should be added to your application context:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:jdbc="http://www.springframework.org/schema/jdbc"
  xmlns:jpa="http://www.springframework.org/schema/data/jpa"
  xmlns:mongo="http://www.springframework.org/schema/data/mongo"
  xsi:schemaLocation="http://www.springframework.org/schema/data/mongo
    http://www.springframework.org/schema/data/mongo/spring-mongo.xsd
    http://www.springframework.org/schema/jdbc
    http://www.springframework.org/schema/jdbc/spring-jdbc-3.0.xsd
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans-3.0.xsd
    http://www.springframework.org/schema/data/jpa
    http://www.springframework.org/schema/data/jpa/spring-jpa-1.0.xsd">

  ...

  <!-- Mongo config -->
  <mongo:mongo host="localhost" port="27017"/>

  <bean id="mongoTemplate" class="org.springframework.data.mongodb.core.MongoTemplate">
    <constructor-arg name="mongo" ref="mongo"/>
    <constructor-arg name="databaseName" value="test"/>
    <constructor-arg name="defaultCollectionName" value="cross-store"/>
  </bean>

  <bean class="org.springframework.data.mongodb.core.MongoExceptionTranslator"/>

  <!-- Mongo cross-store aspect config -->
  <bean class="org.springframework.data.persistence.document.mongo.MongoDocumentBacking"
    factory-method="aspectOf">
    <property name="changeSetPersister" ref="mongoChangeSetPersister"/>
  </bean>
  <bean id="mongoChangeSetPersister"
    class="org.springframework.data.persistence.document.mongo.MongoChangeSetPersister">
    <property name="mongoTemplate" ref="mongoTemplate"/>
    <property name="entityManagerFactory" ref="entityManagerFactory"/>
  </bean>

  ...

</beans>
```

Example 8.3 Example application context with MongoDB and cross-store aspect support

8.2 Writing the Cross Store Application

We are assuming that you have a working JPA application so we will only cover the additional steps needed to persist part of your Entity in your Mongo database. First you need to identify the field you want persisted. It should be a domain class and follow the general rules for the Mongo mapping support covered in previous chapters. The field you want persisted in MongoDB should be annotated using the `@RelatedDocument` annotation. That is really all you need to do!. The cross-store aspects take care of the rest. This includes marking the field with `@Transient` so it won't be persisted using JPA, keeping track of any changes made to the field value and writing them to the database on succesfull transaction

completion, loading the document from MongoDB the first time the value is used in your application. Here is an example of a simple Entity that has a field annotated with `@RelatedEntity`.

```
@Entity
public class Customer {

    @Id
    @GeneratedValue(strategy = GenerationType.IDENTITY)
    private Long id;

    private String firstName;

    private String lastName;

    @RelatedDocument
    private SurveyInfo surveyInfo;

    // getters and setters omitted
}
```

Example 8.4 Example of Entity with `@RelatedDocument`

```
public class SurveyInfo {

    private Map<String, String> questionsAndAnswers;

    public SurveyInfo() {
        this.questionsAndAnswers = new HashMap<String, String>();
    }

    public SurveyInfo(Map<String, String> questionsAndAnswers) {
        this.questionsAndAnswers = questionsAndAnswers;
    }

    public Map<String, String> getQuestionsAndAnswers() {
        return questionsAndAnswers;
    }

    public void setQuestionsAndAnswers(Map<String, String> questionsAndAnswers) {
        this.questionsAndAnswers = questionsAndAnswers;
    }

    public SurveyInfo addQuestionAndAnswer(String question, String answer) {
        this.questionsAndAnswers.put(question, answer);
        return this;
    }
}
```

Example 8.5 Example of domain class to be stored as document

Once the `SurveyInfo` has been set on the `Customer` object above the `MongoTemplate` that was configured above is used to save the `SurveyInfo` along with some metadata about the JPA Entity is stored in a MongoDB collection named after the fully qualified name of the JPA Entity class. The following code:

```
Customer customer = new Customer();
customer.setFirstName("Sven");
customer.setLastName("Olafsen");
SurveyInfo surveyInfo = new SurveyInfo()
    .addQuestionAndAnswer("age", "22")
    .addQuestionAndAnswer("married", "Yes")
    .addQuestionAndAnswer("citizenship", "Norwegian");
customer.setSurveyInfo(surveyInfo);
customerRepository.save(customer);
```

Example 8.6 Example of code using the JPA Entity configured for cross-store persistence

Executing the code above results in the following JSON document stored in MongoDB.

```
{ "_id" : ObjectId( "4d9e8b6e3c55287f87d4b79e" ),
  "_entity_id" : 1,
  "_entity_class" : "org.springframework.data.mongodb.examples.custsvc.domain.Customer",
  "_entity_field_name" : "surveyInfo",
  "questionsAndAnswers" : { "married" : "Yes",
    "age" : "22",
    "citizenship" : "Norwegian" },
  "_entity_field_class"
: "org.springframework.data.mongodb.examples.custsvc.domain.SurveyInfo" }
```

Example 8.7 Example of JSON document stored in MongoDB

9. Logging support

An appender for Log4j is provided in the maven module "spring-data-mongodb-log4j". Note, there is no dependency on other Spring Mongo modules, only the MongoDB driver.

9.1 MongoDB Log4j Configuration

Here is an example configuration

```
log4j.rootCategory=INFO, stdout

log4j.appender.stdout=org.springframework.data.document.mongodb.log4j.MongoLog4jAppender
log4j.appender.stdout.layout=org.apache.log4j.PatternLayout
log4j.appender.stdout.layout.ConversionPattern=%d %p [%c] - <%m>%n
log4j.appender.stdout.host = localhost
log4j.appender.stdout.port = 27017
log4j.appender.stdout.database = logs
log4j.appender.stdout.collectionPattern = %X{year}%X{month}
log4j.appender.stdout.applicationId = my.application
log4j.appender.stdout.warnOrHigherWriteConcern = FSYNC_SAFE

log4j.category.org.apache.activemq=ERROR
log4j.category.org.springframework.batch=DEBUG
log4j.category.org.springframework.data.document.mongodb=DEBUG
log4j.category.org.springframework.transaction=INFO
```

The important configuration to look at aside from host and port is the database and collectionPattern. The variables year, month, day and hour are available for you to use in forming a collection name. This is to support the common convention of grouping log information in a collection that corresponds to a specific time period, for example a collection per day.

There is also an applicationId which is put into the stored message. The document stored from logging as the following keys: level, name, applicationId, timestamp, properties, traceback, and message.

10. JMX support

The JMX support for MongoDB exposes the results of executing the 'serverStatus' command on the admin database for a single MongoDB server instance. It also exposes an administrative MBean, MongoAdmin which will let you perform administrative operations such as drop or create a database. The JMX features build upon the JMX feature set available in the Spring Framework. See [here](#) for more details.

10.1 MongoDB JMX Configuration

Spring's Mongo namespace enables you to easily enable JMX functionality

```
<?xml version="1.0" encoding="UTF-8"?>
  <beans xmlns="http://www.springframework.org/schema/beans"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xmlns:context="http://www.springframework.org/schema/context"
    xmlns:mongo="http://www.springframework.org/schema/data/mongo"
    xsi:schemaLocation=
      "http://www.springframework.org/schema/context
      http://www.springframework.org/schema/context/spring-context-3.0.xsd
      http://www.springframework.org/schema/data/mongo
      http://www.springframework.org/schema/data/mongo/spring-mongo-1.0.xsd
      http://www.springframework.org/schema/beans http://www.springframework.org/schema/
beans/spring-beans-3.0.xsd">

    <beans>

      <!-- Default bean name is 'mongo' -->
      <mongo:mongo host="localhost" port="27017"/>

      <!-- by default look for a Mongo object named 'mongo' -->
      <mongo:jmx/>

      <context:mbean-export/>

      <!-- To translate any MongoExceptions thrown in @Repository annotated classes -->
      <context:annotation-config/>

      <bean id="registry" class="org.springframework.remoting.rmi.RmiRegistryFactoryBean" p:port="1099" />

      <!-- Expose JMX over RMI -->

      <bean id="serverConnector" class="org.springframework.jmx.support.ConnectorServerFactoryBean"
        depends-on="registry"
        p:objectName="connector:name=rmi"
        p:serviceUrl="service:jmx:rmi://localhost/jndi/rmi://localhost:1099/myconnector" /
      >

    </beans>
```

Example 10.1 XML schmea to configure MongoDB

This will expose several MBeans

- AssertMetrics
- BackgroundFlushingMetrics

- BtreeIndexCounters
- ConnectionMetrics
- GlobalLoclMetrics
- MemoryMetrics
- OperationCounters
- ServerInfo
- MongoAdmin

This is shown below in a screenshot from JConsole

Part III. Appendix

Appendix A. Namespace reference

A.1 The `<repositories />` element

The `<repositories />` triggers the setup of the Spring Data repository infrastructure. The most important attribute is `base-package` which defines the package to scan for Spring Data repository interfaces.¹

Table A.1. Attributes

| Name | Description |
|--------------------------------------|---|
| <code>base-package</code> | Defines the package to be used to be scanned for repository interfaces extending <code>*Repository</code> (actual interface is determined by specific Spring Data module) in auto detection mode. All packages below the configured package will be scanned, too. Wildcards are also allowed. |
| <code>repository-impl-postfix</code> | Defines the postfix to autodetect custom repository implementations. Classes whose names end with the configured postfix will be considered as candidates. Defaults to <code>Impl</code> . |
| <code>query-lookup-strategy</code> | Determines the strategy to be used to create finder queries. See the section called “Query lookup strategies” for details. Defaults to <code>create-if-not-found</code> . |

¹see the section called “XML Configuration”

Appendix B. Repository query keywords

B.1 Supported query keywords

The following table lists the keywords generally supported by the Spring data repository query derivation mechanism. However consult the store specific documentation for the exact list of supported keywords as some of the ones listed here might not be supported in a particular store.

Table B.1. Query keywords

| Logical keyword | Keyword expressions |
|--------------------|--|
| AFTER | After, IsAfter |
| BEFORE | Before, IsBefore |
| CONTAINING | Containing, IsContaining, Contains |
| BETWEEN | Between, IsBetween |
| ENDING_WITH | EndingWith, IsEndingWith, EndsWith |
| EXISTS | Exists |
| FALSE | False, IsFalse |
| GREATER_THAN | GreaterThan, IsGreaterThan |
| GREATER_THAN_EQUAL | GreaterThanEqual, IsGreaterThanEqual |
| IN | In, IsIn |
| IS | Is, Equals, (or no keyword) |
| IS_NOT_NULL | NotNull, IsNotNull |
| IS_NULL | Null, IsNull |
| LESS_THAN | LessThan, IsLessThan |
| LESS_THAN_EQUAL | LessThanEqual, IsLessThanEqual |
| LIKE | Like, IsLike |
| NEAR | Near, IsNear |
| NOT | Not, IsNot |
| NOT_IN | NotIn, IsNotIn |
| NOT_LIKE | NotLike, IsNotLike |
| REGEX | Regex, MatchesRegex, Matches |
| STARTING_WITH | StartingWith, IsStartingWith, StartsWith |

| Logical keyword | Keyword expressions |
|-----------------|---------------------|
| TRUE | True, IsTrue |
| WITHIN | Within, IsWithin |