



Spring Data GemFire Reference Guide

1.4.0.M1

Costin Leau SpringSource, a division of Pivotal , David Turanski SpringSource, a division of Pivotal , John Blum SpringSource, a division of Pivotal , Oliver Gierke SpringSource, a division of Pivotal , Lyndon Adams Pivotal

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	iv
I. Introduction	1
1. Introduction	2
2. Requirements	3
3. New Features	4
3.1. New in the 1.3.0 Release	4
3.2. New in the 1.2.1 Release	4
3.3. New in the 1.2.0 Release	4
II. Reference Guide	5
4. Document structure	6
5. Bootstrapping GemFire through the Spring Container	7
5.1. Advantages of using Spring over GemFire <code>cache.xml</code>	7
5.2. Using the Core Spring Data GemFire Namespace	7
5.3. Configuring the GemFire Cache	9
Advanced Cache Configuration	10
Enabling PDX Serialization	11
Configuring a GemFire Cache Server	11
Configuring a GemFire Client Cache	12
5.4. Using the GemFire Data Access Namespace	13
An Easy Way to Connect to GemFire	13
5.5. Configuring a GemFire Region	14
Using an externally configured Region	14
Configuring Regions	15
Common Region Attributes	15
Cache Listeners	16
Cache Loaders and Cache Writers	17
Subregions	17
A Word of Caution on Regions, Subregions and Lookups	18
Data Persistence	20
Subscription Interest Policy	21
Data Eviction and Overflowing	21
Data Expiration	22
Local Region	22
Replicated Region	22
Partitioned Region	22
partitioned-region Options	23
Client Region	24
Client Interests	25
JSON Support	25
5.6. Creating an Index	26
5.7. Configuring a Disk Store	26
5.8. Configuring GemFire's Function Service	26
5.9. Configuring WAN Gateways	27
WAN Configuration in GemFire 7.0	27
WAN Configuration in GemFire 6.6	29
6. Working with the GemFire APIs	30
6.1. Exception translation	30

6.2. GemfireTemplate	30
6.3. Support for Spring Cache Abstraction	31
6.4. Transaction Management	31
6.5. GemFire Continuous Query Container	32
Continuous Query Listener Container	32
The ContinuousQueryListenerAdapter and ContinuousQueryListener	32
6.6. Wiring Declarable components	34
Configuration using <i>template</i> definitions	35
Configuration using auto-wiring and annotations	36
7. Working with GemFire Serialization	38
7.1. Wiring deserialized instances	38
7.2. Auto-generating custom Instantiators	38
8. POJO mapping	40
8.1. Entity mapping	40
8.2. Mapping PDX serializer	40
9. GemFire Repositories	41
9.1. Introduction	41
9.2. Spring configuration	41
9.3. Executing OQL queries	41
10. Annotation Support for Function Execution	43
10.1. Introduction	43
10.2. Implementation vs Execution	43
10.3. Implementing a Function	43
Annotations for Function Implementation	44
Batching Results	45
Enabling Annotation Processing	45
10.4. Executing a Function	45
Annotations for Function Execution	46
Enabling Annotation Processing	46
10.5. Programmatic Function Execution	46
11. Sample Applications	48
11.1. Hello World	48
Starting and stopping the sample	48
Using the sample	48
Hello World Sample Explained	49
III. Other Resources	50
12. Useful Links	51
IV. Appendices	52
A. Spring Data GemFire Schema	53

Preface

Spring Data GemFire focuses on integrating Spring Framework's powerful, non-invasive programming model and concepts with vFabric GemFire, simplifying configuration, development and providing high-level abstractions. This document assumes the reader already has a basic familiarity with the Spring Framework and vFabric GemFire concepts and APIs.

While every effort has been made to ensure that this documentation is comprehensive and there are no errors, some topics might require more explanation and some typos might have crept in. If you do spot any mistakes or even more serious errors and you can spare a few cycles during lunch, please do bring the error to the attention of the Spring Data GemFire team by raising an [issue](#). Thank you.

Part I. Introduction

1. Introduction

This reference guide for the Spring Data GemFire project explains how to use Spring framework to configure and develop applications with vFabric GemFire. It presents the basic concepts, semantics and provides numerous examples to help you get started.



Note

Spring Data GemFire started as a top level Spring project called Spring GemFire (SGF) and has since moved under the Spring Data umbrella project and has been renamed accordingly.

2. Requirements

Spring Data GemFire requires JDK level 6.0 and above, Spring [Framework](#) 3 and [vFabric GemFire](#) 6.6 and above (version 7 or above is recommended).

3. New Features



Note

As of the 1.2.0 release, this project, formerly known as Spring GemFire, has been renamed to Spring Data GemFire to reflect that it is now a component of the [Spring Data](#) project.

3.1 New in the 1.3.0 Release

- Annotation support for GemFire functions. It is now possible to declare and register functions written as POJOs with annotations. In addition function executions are defined as annotated interfaces, similar to the way Spring Data repositories work. See Chapter 10, *Annotation Support for Function Execution*
- We have added a `<datasource>` tag to the *gfe-data* XML namespace. This simplifies establishing a basic client [connection](#) to a GemFire data grid.
- To [support JSON](#) features introduced in GemFire 7.0, we have added a `<json-region-autoproxy>` tag to the *gfe-data* XML namespace, enabling Spring AOP to perform the necessary conversions automatically on region operations.
- Upgraded to GemFire 7.0.1 and added namespace support for new AsyncEventQueue attributes
- Added support for setting subscription interest policy on regions

3.2 New in the 1.2.1 Release

- WAN Gateway Support for GemFire 7.0

3.3 New in the 1.2.0 Release

- Full support for GemFire configuration via the *gfe* namespace. Now GemFire components may be configured completely without requiring a native `cache.xml` file.
- WAN Gateway support for GemFire 6.6.x. See Section 5.9, “Configuring WAN Gateways”
- Spring Data Repository support with a dedicated namespace, *gfe-data*. See Chapter 9, *GemFire Repositories*
- Namespace support for registering GemFire functions. See Section 5.8, “Configuring GemFire's Function Service”
- A top level `<disk-store>` element has been added to the *gfe* namespace to allow sharing of persist stores among regions, and other components that support persistent backup. See Section 5.7, “Configuring a Disk Store”



Caution

The `<*-region>` elements no longer allow a nested `<disk-store>`

- GemFire subregions are supported via nested `<*-region>` elements
- A `<local-region>` element has been added to configure a local region

Part II. Reference Guide

4. Document structure

The following chapters explain the core functionality offered by Spring Data GemFire.

Chapter 5, *Bootstrapping GemFire through the Spring Container* describes the configuration support provided for bootstrapping, initializing, configuring, and accessing GemFire caches, cache servers, regions, and related distributed system components

Chapter 6, *Working with the GemFire APIs* explains the integration between the GemFire APIs and the various data access features available in Spring, such as transaction management and exception translation.

Chapter 7, *Working with GemFire Serialization* describes the enhancements for GemFire (de)serialization and management of associated objects.

Chapter 8, *POJO mapping* describes persistence mapping for POJOs stored in GemFire using Spring Data

Chapter 9, *GemFire Repositories* describes how to create and use GemFire Repositories using Spring Data

Chapter 10, *Annotation Support for Function Execution* describes how to create and use GemFire Functions using Spring Data

Chapter 11, *Sample Applications* describes the samples provided with the distribution to illustrate the various features available in Spring Data GemFire.

5. Bootstrapping GemFire through the Spring Container

Spring Data GemFire provides full configuration and initialization of the GemFire data grid through Spring's IoC container and provides several classes that simplify the configuration of GemFire components including caches, regions, WAN gateways, persistence backup, and other distributed system components to support a variety of scenarios with minimal effort.



Note

This section assumes basic familiarity with GemFire. For more information see the [product documentation](#).

5.1 Advantages of using Spring over GemFire `cache.xml`

As of release 1.2.0, Spring Data GemFire's XML namespace supports full configuration of the data grid. In fact, the Spring Data GemFire namespace is considered the preferred way to configure GemFire. GemFire will continue to support native `cache.xml` for legacy reasons, but you can now do everything in Spring XML and take advantage of the many wonderful things Spring has to offer such as modular XML configuration, property placeholders, SpEL, and environment profiles. Behind the namespace, Spring Data GemFire makes extensive use of Spring's `FactoryBean` pattern to simplify the creation and initialization of GemFire components.

For example, GemFire provides several callback interfaces such as `CacheListener`, `CacheWriter`, and `CacheLoader` to allow developers to add custom event handlers. Using the Spring IoC container, these may be configured as normal Spring beans and injected into GemFire components. This is a significant improvement over native `cache.xml` which provides relatively limited configuration options and requires callbacks to implement GemFire's `Declarable` interface (see Section 6.6, "Wiring `Declarable` components" to see how you can still use `Declarables` within Spring's DI container).

In addition, IDEs such as the Spring Tool Suite (STS) provide excellent support for Spring XML namespaces, such as code completion, pop-up annotations, and real time validation, making them easy to use.

5.2 Using the Core Spring Data GemFire Namespace

To simplify configuration, Spring Data GemFire provides a dedicated XML namespace for configuring core GemFire components. It is also possible to configure the beans directly through Spring's standard `<bean>` definition. However, as of Spring Data GemFire 1.2.0, all bean properties are exposed via the namespace so there is little benefit to using raw bean definitions. For more information about XML Schema-based configuration in Spring, see [this](#) appendix in the Spring Framework reference documentation.



Note

Spring Data Repository support uses a separate XML namespace. See Chapter 9, *GemFire Repositories* for more information on how to configure GemFire Repositories.

To use the Spring Data GemFire namespace, simply declare it in your Spring XML configuration meta-data:

```
<?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:gfe="http://www.springframework.org/schema/gemfire"❶
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans http://www.springframework.org/schema/
    beans/spring-beans.xsd
    ❸http://www.springframework.org/schema/gemfire http://www.springframework.org/
    schema/gemfire/spring-gemfire.xsd">

  <bean id ... >

    ❹<gfe:cache ...>

</beans>
```

- ❶ Spring GemFire namespace prefix. Any name will do but through out the reference documentation, `gfe` will be used.
- ❷ The namespace URI.
- ❸ The namespace URI location. Note that even though the location points to an external address (which exists and is valid), Spring will resolve the schema locally as it is included in the Spring Data GemFire library.
- ❹ Declaration example for the GemFire namespace. Notice the prefix usage.

Once declared, the namespace elements can be declared simply by appending the aforementioned prefix.

Note

It is possible to change the default namespace, for example from `beans` to `gfe`. This is useful for configuration composed mainly of GemFire components as it avoids declaring the prefix. To achieve this, simply swap the namespace prefix declaration above:

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

  ❸<beans:bean id ... >

  ❹<cache ...>

</beans>
```

- ❶ The default namespace declaration for this XML file points to the Spring Data GemFire namespace.
- ❷ The beans namespace prefix declaration.

- ③ Bean declaration using the `beans` namespace. Notice the prefix.
- ④ Bean declaration using the `gfe` namespace. Notice the lack of prefix (as the default namespace is used).

5.3 Configuring the GemFire Cache

In order to use GemFire, one needs to either create a new `Cache` or connect to an existing one. In the current version of GemFire, there can be only one opened cache per VM (or per classloader to be technically correct). In most cases the cache is created once.



Note

This section describes the creation and configuration of a full cache member, appropriate for peer to peer cache topologies and cache servers. A full cache is also commonly used for standalone applications, integration tests and proofs of concept. In a typical production system, most application processes will act as cache clients and will create a `ClientCache` instance instead. This is described in the sections the section called “Configuring a GemFire Client Cache” and the section called “Client Region”

A cache with default configuration can be created with a very simple declaration:

```
<gfe:cache/>
```

A Spring application context containing this definition will, upon initialization, will register a `CacheFactoryBean` to create a Spring bean named `gemfireCache` referencing a GemFire `Cache` instance. This will be either an existing cache, or if one does not exist, a newly created one. Since no additional properties were specified, a newly created cache will apply the default cache configuration.

All Spring Data GemFire components which depend on the `Cache` respect this naming convention so that there is no need to explicitly declare the `Cache` dependency. If you prefer, you can make the dependence explicit via the `cache-ref` attribute provided by various namespace elements. Also you can easily override the `Cache`'s bean name:

```
<gfe:cache id="my-cache"/>
```

Starting with Spring Data GemFire 1.2.0, The GemFire `Cache` may be fully configured using Spring. However, GemFire's native XML configuration file (e.g., `cache.xml`) is also supported. For scenarios in which the GemFire cache needs to be configured natively, simply provide a reference the GemFire configuration file using the `cache-xml-location` attribute:

```
<gfe:cache id="cache-with-xml" cache-xml-location="classpath:cache.xml"/>
```

In this example, if the cache needs to be created, it will use the file named `cache.xml` located in the classpath root.



Note

Note that the configuration makes use of Spring's [Resource](#) abstraction to locate the file. This allows various search patterns to be used, depending on the runtime environment or the prefix specified (if any) in the resource location.

In addition to referencing an external configuration file one can specify GemFire [properties](#) using any of Spring's common properties support features. For example, one can use the `properties`

element defined in the `util` namespace to define properties directly or load properties from properties files. The latter is recommended for externalizing environment specific settings outside the application configuration:

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

    <gfe:cache properties-ref="props"/>

    <util:properties id="props" location="file:/vfabric/gemfire/gemfire.properties"/>
</beans>
```



Note

The cache settings apply only if a new cache needs to be created. If an open cache already exists in the JVM, these settings will be ignored.

Advanced Cache Configuration

For advanced cache configuration, the `cache` element provides a number of configuration options exposed as attributes or child elements

```
❶ <gfe:cache
    copy-on-read="true"
    critical-heap-percentage="70"
    eviction-heap-percentage="60"
    lock-lease="120"
    lock-timeout="60"
    pdx-serializer="myPdxSerializer"
    pdx-disk-store="diskStore"
    pdx-ignore-unread-fields="true"
    pdx-persistent="true"
    pdx-read-serialized="false"
    message-sync-interval="1"
    search-timeout="300"
    close="false"
    lazy-init="true"
>
❷ <gfe:transaction-listener ref="myTransactionListener"/>
❸ <gfe:transaction-writer>
    <bean class="org.springframework.data.gemfire.example.TransactionListener"/>
</gfe:transaction-writer>
❹ <gfe:dynamic-region-factory/>
❺ <gfe:jndi-binding jndi-name="myDataSource" type="ManagedDataSource"/>
</gfe:cache>
```

- ❶ Various cache options are supported by attributes. For further information regarding anything shown in this example, please consult the GemFire product [documentation](#)

The `close` attribute determines if the cache should be closed when the Spring application context is closed. The default is `true` however for cases in which multiple application contexts use the cache (common in web applications), set this value to `false`.

The `lazy-init` attribute determines if the cache should be initialized before another bean references it. The default is `true` however in some cases it may be convenient to set this value to `false`.

- ❷ An example of a `TransactionListener` callback declaration using a bean reference. The referenced bean must implement [TransactionListener](#)
- ❸ An example of a `TransactionWriter` callback declaration using an inner bean declaration this time. The bean must implement [TransactionWriter](#)
- ❹ Enable GemFire's [DynamicRegionFactory](#)
- ❺ Declares a JNDI binding to enlist an external datasource in a GemFire transaction



Note

The `use-bean-factory-locator` attribute (not shown) deserves a mention. The factory bean responsible for creating the cache uses an internal Spring type called a `BeanFactoryLocator` to enable user classes declared in GemFire's native `cache.xml` to be registered as Spring beans. The `BeanFactoryLocator` implementation also permits only one bean definition for a cache with a given id. In certain situations, such as running JUnit integration tests from within Eclipse, it is necessary to disable the `BeanFactoryLocator` by setting this value to `false` to prevent an exception. This exception may also arise during JUnit tests running from a build script. In this case the test runner should be configured to fork a new JVM for each test (in maven, set `<forkmode>always</forkmode>`). Generally there is no harm in setting this value to `false`.

Enabling PDX Serialization

The example above includes a number of attributes related to GemFire's enhanced serialization framework, PDX. While a complete discussion of PDX is beyond the scope of this reference guide, it is important to note that PDX is enabled by registering a PDX serializer which is done via the `pdx-serializer` attribute. GemFire provides an implementation class `com.gemstone.gemfire.pdx.ReflectionBasedAutoSerializer`, however it is common for developers to provide their own implementation. The value of the attribute is simply a reference to a Spring bean that implements the required interface. More information on serialization support can be found in Chapter 7, *Working with GemFire Serialization*

Configuring a GemFire Cache Server

In Spring Data GemFire 1.1 dedicated support for configuring a [CacheServer](#) was added, allowing complete configuration through the Spring container:

```
<?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:gfe="http://www.springframework.org/schema/gemfire"
  xmlns:context="http://www.springframework.org/schema/context"
  xsi:schemaLocation="http://www.springframework.org/schema/gemfire http://
www.springframework.org/schema/gemfire/spring-gemfire.xsd
  http://www.springframework.org/schema/beans http://www.springframework.org/schema/beans/
spring-beans.xsd
  http://www.springframework.org/schema/context http://www.springframework.org/schema/
context/spring-context.xsd">

  <gfe:cache />

  <!-- Advanced example depicting various cache server configuration options -->
  <gfe:cache-server id="advanced-config" auto-startup="true"
    bind-address="localhost" port="${gfe.port.6}" host-name-for-clients="localhost"
    load-poll-interval="2000" max-connections="22" max-threads="16"
    max-message-count="1000" max-time-between-pings="30000"
    groups="test-server">

    <gfe:subscription-config eviction-type="ENTRY" capacity="1000" disk-store="file://
${java.io.tmpdir}"/>
  </gfe:cache-server>

  <context:property-placeholder location="classpath:cache-server.properties"/>

</beans>
```

The configuration above illustrates the `cache-server` element and the many options available.



Note

Rather than hard-coding the port, this configuration uses Spring's [context](#) namespace to declare a `property-placeholder`. The [property placeholder](#) reads one or more properties file and then replaces property placeholders with values at runtime. This allows administrators to change such values without having to touch the main application configuration. Spring also provides [SpEL](#) and the [environment abstraction](#) one to support externalization of environment specific properties from the main code base, easing the deployment across multiple machines.



Note

To avoid initialization problems, the `CacheServers` started by Spring Data GemFire will start *after* the container has been fully initialized. This allows potential regions, listeners, writers or instantiators defined declaratively to be fully initialized and registered before the server starts accepting connections. Keep this in mind when programmatically configuring these items as the server might start after your components and thus not be seen by the clients connecting right away.

Configuring a GemFire Client Cache

Another configuration addition in Spring Data GemFire 1.1 is the dedicated support for configuring [ClientCache](#). This is similar to a [cache](#) in both usage and definition and supported by `org.springframework.data.gemfire.clientClientCacheFactoryBean`.


```
<beans>
  <gfe:client-cache />
</beans>
```

`client-cache` supports much of the same options as the `cache` element. However as opposed to a *full* cache, a client cache connects to a remote cache server through a pool. By default a pool is created to connect to a server on `localhost` port `40404`. The the default pool is used by all client regions unless the region is configured to use a different pool.

Pools can be defined through the `pool` element; The client side `pool` can be used to configure connectivity to the server for individual entities or for the entire cache. For example, to customize the default pool used by `client-cache`, one needs to define a pool and wire it to cache definition:

```
<beans>
  <gfe:client-cache id="simple" pool-name="my-pool"/>

  <gfe:pool id="my-pool" subscription-enabled="true">
    <gfe:locator host="${locatorHost}" port="${locatorPort}"/>
  </gfe:pool>
</beans>
```

The `<client-cache>` tag also includes a `ready-for-events` attribute. If set to `true`, the client cache initialization will include [ClientCache.readyForEvents\(\)](#).

Client side configuration is covered in more detail in the section called “Client Region”.

5.4 Using the GemFire Data Access Namespace

In addition to the core `gfe` namespace, Spring Data GemFire provides a `gfe-data` namespace intended primarily to simplify the development of GemFire client applications. This namespace currently supports for GemFire [repositories](#) and function [execution](#) and a `<datasource>` tag that offers a convenient way to connect to the data grid.

An Easy Way to Connect to GemFire

For many applications, A basic connection to a GemFire grid, using default values is sufficient. Spring Data GemFire's `<datasource>` tag provides a simple way to access data. The data source creates a client cache and connection pool. In addition, it will query the member servers for all existing root regions and create a proxy (empty) client region for each one.

```
<gfe-data:datasource>
  <locator host="somehost" port="1234"/>
</gfe-data:datasource>
```

The `datasource` tag is syntactically similar to `<gfe:pool>`. It may be configured with one or more `locator` or `server` tags to connect to an existing data grid. Additionally, all attributes available to configure a pool are supported. This configuration will automatically create `ClientRegion` beans for each region defined on members connected to the locator, so they may be seamlessly referenced by Spring Data mapping annotations, `GemfireTemplate`, and wired into application classes.

Of course, you can explicitly configure client regions. For example, if you want to cache data in local memory:

```
<gfe-data:datasource>
  <locator host="somehost" port="1234"/>
</gfe-data:datasource>

<gfe:client-region id="Customer" shortcut="CACHING_PROXY"/>
```

5.5 Configuring a GemFire Region

A region is required to store and retrieve data from the cache. `Region` is an interface extending `java.util.Map` and enables basic data access using familiar key-value semantics. The `Region` interface is wired into classes that require it so the actual region type is decoupled from the programming model. Typically each region is associated with one domain object, similar to a table in a relational database.

GemFire implements the following types of regions:

- *Replicated* - Data is replicated across all cache members that define the region. This provides very high read performance but writes take longer to perform the replication.
- *Partitioned* - Data is partitioned into buckets among cache members that define the region. This provides high read and write performance and is suitable for very large data sets that are too big for a single node.
- *Local* - Data only exists on the local node.
- *Client* - Technically a client region is a local region that acts as a proxy to a replicated or partitioned region hosted on cache servers. It may hold data created or fetched locally. Alternately, it can be empty. Local updates are synchronized to the cache server. Also, a client region may subscribe to events in order to stay synchronized with changes originating from remote processes that access the same region.

For more information about the various region types and their capabilities as well as configuration options, please refer to the GemFire Developer's [Guide](#) and community [site](#).

Using an externally configured Region

For referencing regions already configured through GemFire native configuration, e.g. a `cache.xml` file, use the `lookup-region` element. Simply declare the target region name with the `name` attribute; for example, to declare a bean definition named `region-bean` for an existing region named `Orders` one can use the following bean definition:

```
<gfe:lookup-region id="region-bean" name="Orders"/>
```

If the `name` is not specified, the bean's `id` will be used. The example above becomes:

```
<!-- lookup for a region called 'Orders' -->
<gfe:lookup-region id="Orders"/>
```



Note

If the region does not exist, an initialization exception will be thrown. For configuring new GemFire regions, proceed to the appropriate sections below.

Note, in the previous examples, since no cache name was defined, the default naming convention (`gemfireCache`) was used. Alternately, one can reference the cache bean through the `cache-ref` attribute:

```
<gfe:cache id="cache"/>
<gfe:lookup-region id="region-bean" name="Orders" cache-ref="cache"/>
```

The `lookup-region` provides a simple way of retrieving existing, pre-configured regions without exposing the region semantics or setup infrastructure.

Configuring Regions

Spring Data GemFire provides comprehensive support for configuring any type of GemFire Region via the following elements:

- Local Region `<local-region>`
- Replicated Region `<replicated-region>`
- Partitioned Region `<partitioned-region>`
- Client Region `<client-region>`

For a comprehensive description of [Region types](#) please consult the GemFire product documentation.

Common Region Attributes

The following table(s) list attributes available for various region types:

Table 5.1. Common Region Attributes

Name	Values	Description
<code>cache-ref</code>	<i>GemFire cache bean name</i>	The name of the bean defining the GemFire cache (by default 'gemfireCache').
<code>close</code>	<i>boolean, default:false (Note: The default was true prior to 1.3.0)</i>	Indicates whether the region should be closed at shutdown
<code>data-policy</code>	<i>See GemFire's Data Policy</i>	The region's data policy. Note not all data policies are supported for every region type
<code>destroy</code>	<i>boolean, default:false</i>	Indicates whether the region should be destroyed at shutdown
<code>disk-store-ref</code>	<i>The name of a configured disk store</i>	A reference to a bean created via the <code>disk-store</code> element. Note: This will automatically enable persistence. If <code>persistent</code> is explicitly set to false, an exception will be thrown.
<code>disk-synchronous</code>	<i>boolean, default:false</i>	Indicates whether disk store writes are synchronous. Note: This will automatically enable persistence. If <code>persistent</code> is explicitly set to false, an exception will be thrown.

Name	Values	Description
enable-gateway	<i>boolean, default:false</i>	Indicates whether the region will synchronize entries over a WAN gateway.
hub-id	<i>The name of the Gateway Hub</i>	This will automatically set enable-gateway to true. If enable-gateway is explicitly set to false, an exception will be thrown.
id	<i>any valid bean name</i>	Will also be the region name by default
ignore-jta	<i>boolean, default:false</i>	Indicates whether the region participates in JTA transactions
index-update-type	<i>synchronous or asynchronous, default:asynchronous</i>	Indicates whether indices will be updated synchronously or asynchronously on entry creation
initial-capacity	<i>integer, default:16</i>	The initial memory allocation for number of entries
key-constraint	<i>any valid java class name</i>	The expected key type
name	<i>any valid region name</i>	The name of the region definition. If no specified, it will assume the value of the id attribute (the bean name).
persistent	<i>boolean, default:false</i>	Indicates whether the region persists entries to a disk store
statistics	<i>boolean, default:false</i>	Indicates whether the region reports statistics
value-constraint	<i>any valid java class name</i>	The expected value type

Cache Listeners

Cache Listeners are registered with a region to handle region events such as entries being created, updated, destroyed, etc. A Cache Listener can be any bean that implements the [CacheListener](#) interface. A region may have multiple listeners, declared using the `cache-listener` element enclosed in a `*-region` element. In the example below, there are two `CacheListeners` declared. The first references a top-level named Spring bean; the second is an anonymous inner bean definition.

```
<gfe:replicated-region id="region-with-listeners">
  <gfe:cache-listener>
    <!-- nested cache listener reference -->
    <ref bean="c-listener"/>
    <!-- nested cache listener declaration -->
    <bean class="some.pkg.AnotherSimpleCacheListener"/>
  </gfe:cache-listener>

  <bean id="c-listener" class="some.pkg.SimpleCacheListener"/>
</gfe:replicated-region>
```

The following example uses an alternate form of the `cache-listener` element with a `ref` attribute. This allows for more concise configuration for a single cache listener. Note that the namespace only allows a single `cache-listener` element so either the style above or below must be used.

Caution

Using `ref` and a nested declaration in a `cache-listener`, or similar element, is illegal. The two options are mutually exclusive and using both on the same element will result in an exception.

```
<beans>
  <gfe:replicated-region id="region-with-one listener">
    <gfe:cache-listener ref="c-listener"/>
  </gfe:replicated-region>

  <bean id="c-listener" class="some.pkg.SimpleCacheListener"/>
</beans>
```

Bean Reference Conventions

The `cache-listener` element is an example of a common pattern used in the namespace anywhere GemFire provides a callback interface to be implemented in order to invoke custom code in response to cache or region events. Using Spring's IoC container, the implementation is a standard Spring bean. In order to simplify the configuration, the schema allows a single occurrence of the `cache-listener` element, but it may contain nested bean references and inner bean definitions in any combination if multiple instances are permitted. The convention is to use the singular form (i.e., `cache-listener` vs `cache-listeners`) reflecting that the most common scenario will in fact be a single instance. We have already seen examples of this pattern in the [advanced cache](#) configuration example.

Cache Loaders and Cache Writers

Similar to `cache-listener`, the namespace provides `cache-loader` and `cache-writer` elements to register these respective components for a region. A `CacheLoader` is invoked on a cache miss to allow an entry to be loaded from an external data source, a database for example. A `CacheWriter` is invoked after an entry is created or updated, intended for synchronizing to an external data source. The difference is GemFire only supports at most a single instance of each for each region. However, either declaration style may be used. See [CacheLoader](#) and [CacheWriter](#) for more details.

Subregions

In Release 1.2.0, Spring Data GemFire added support for subregions, allowing regions to be arranged in a hierarchical relationship. For example, GemFire allows for a `/Customer/Address` region and a different `/Employee/Address` region. Additionally, a subregion may have its own subregions and its own configuration. A subregion does not inherit attributes from the parent region. Regions types may be mixed and matched subject to GemFire constraints. A subregion is naturally declared as a child element of a region. The subregion's name attribute is the simple name. The above example might be configured as:

```
<beans>

  <gfe:replicated-region name="Customer">
    <gfe:replicated-region name="Address"/>
  </gfe:replicated-region>

  <gfe:replicated-region name="Employee">
    <gfe:replicated-region name="Address"/>
  </gfe:replicated-region>

</beans>
```

Note that the `id` attribute is not permitted for a subregion. The subregions will be created with bean names `/Customer/Address` and `/Employee/Address`, respectively. So they may be injected using the full path name into other beans that use them, such as `GemfireTemplate`. The full path should also be used in OQL query strings.

A Word of Caution on Regions, Subregions and Lookups

One of the underlying properties of the high-level `replicated-region`, `partitioned-region`, `local-region` and `client-region` elements in Spring Data GemFire's XML namespace, which correspond to GemFire's Region types based on Data Policy, is that these elements perform a lookup first before attempting to create the region. This is done in case the region already exists, which might be the case if the region was defined in GemFire's native configuration, e.g. `cache.xml`, thereby avoiding any errors. This was by design, though subject to change.



Caution

The Spring team highly recommends that the `replicated-region`, `partitioned-region`, `local-region` and `client-region` elements be strictly used only for defining new regions. One of the problems with these elements doing a lookup first is, if the developer assumed that defining a bean definition for a REPLICATE region would create a new region, however, consequently a region with the same name already exists having different semantics for eviction, expiration, subscription and/or other attributes, this could adversely affect application logic and/or expectations thereby violating application requirements.



Important

Recommended Practice - Only use the `replicated-region`, `partitioned-region`, `local-region` and `client-region` XML namespace elements for defining new regions.

However, because the high-level region elements perform a lookup first, this can cause problems for dependency injected region resources to application code, like DAOs or Repositories.

Take for instance the following native GemFire configuration file (e.g. `cachel.xml`)...

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE cache PUBLIC "-//GemStone Systems, Inc.//GemFire Declarative Caching 7.0//EN"
"http://www.gemstone.com/dtd/cache7_0.dtd">
<cache>
  <region name="Customers" refid="REPLICATE">
    <region name="Accounts" refid="REPLICATE">
      <region name="Orders" refid="REPLICATE">
        <region name="Items" refid="REPLICATE"/>
      </region>
    </region>
  </region>
</cache>
```

Also, consider that you might have defined a DAO as follows...

```
public class CustomerAccountDao extends GemDaoSupport {

    @Resource(name = "Customers/Accounts")
    private Region customersAccounts;

    ...

}
```

Here, we are injecting a reference to the `Customers/Accounts` GemFire Region in our DAO. As such, it is not uncommon for a developer to define beans for all or some of these regions in Spring XML configuration meta-data as follows...

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

  <gfe:cache cache-xml-location="classpath:cache.xml"/>

  <gfe:lookup-region name="Customers/Accounts"/>
  <gfe:lookup-region name="Customers/Accounts/Orders"/>
</beans>
```

Here the `Customers/Accounts` and `Customers/Accounts/Orders` GemFire Regions are referenced as beans in the Spring context as `"Customers/Accounts"` and `"Customers/Accounts/Orders"`, respectively. The nice thing about using the `lookup-region` element and the corresponding syntax above is that it allows a developer to reference a subregion directly without unnecessarily defining a bean for the parent region (e.g. `Customers`).

However, if now the developer changes his/her configuration meta-data syntax to using the nested format, like so...

```
<gfe:lookup-region name="Customers">
  <gfe:lookup-region name="Accounts">
    <gfe:lookup-region name="Orders"/>
  </gfe:lookup-region>
</gfe:lookup-region>
```

Or, perhaps the developer erroneously chooses to use the high-level `replicated-region` element, which will do a lookup first, as in...

```
<gfe:replicated-region name="Customers" persistent="true">
  <gfe:replicated-region name="Accounts" persistent="true">
    <gfe:replicated-region name="Orders" persistent="true"/>
  </gfe:replicated-region>
</gfe:replicated-region>
```

Then the region beans defined in the Spring context will consist of the following: { "Customers", "/Customers/Accounts", "/Customers/Accounts/Orders" }. This means the dependency injected reference (i.e. `@Resource(name = "Customers/Accounts")`) is now broken since no bean with name "Customers/Accounts" is defined.

GemFire is flexible in referencing both parent regions and subregions. The parent can be referenced as "/Customers" or "Customers" and the child as "/Customers/Accounts" or just "Customers/Accounts". However, Spring Data GemFire is very specific when it comes to naming beans after regions, typically always using the forward slash (/) to represents subregions (e.g. "/Customers/Accounts").

Therefore, it is recommended that users use either the nested `lookup-region` syntax as illustrated above, or define direct references with a leading forward slash (/) like so...

```
<gfe:lookup-region name="/Customers/Accounts"/>
<gfe:lookup-region name="/Customers/Accounts/Orders"/>
```

The example above where the nested `replicated-region` elements were used to reference the subregions serves to illustrate the problem stated earlier. Are the Customers, Accounts and Orders Regions/Subregions persistent or not? Not, since the regions were defined in native GemFire configuration (i.e. `cache.xml`) and will exist by the time the cache is initialized, or once the `<gfe:cache>` bean is created. Since the high-level region XML namespace abstractions, like `replicated-region`, perform the lookup first, it uses the regions as defined in the `cache.xml` configuration file.

Data Persistence

Regions can be made persistent. GemFire ensures that all the data you put into a region that is configured for persistence will be written to disk in a way that it can be recovered the next time you create the region. This allows data to be recovered after a machine or process failure or after an orderly shutdown and restart of GemFire.

To enable persistence with Spring Data GemFire, simply set the `persistent` attribute to `true`:

```
<gfe:partitioned-region id="persitent-partition" persistent="true"/>
```




Important

Persistence for partitioned regions is supported from GemFire 6.5 onwards - configuring this option on a previous release will trigger an initialization exception.

Persistence may also be configured using the `data-policy` attribute, set to one of [GemFire's data policy settings](#). For instance...

```
<gfe:partitioned-region id="persitent-partition" data-policy="PERSISTENT_PARTITION"/>
```

The data policy must match the region type and must also agree with the `persistent` attribute if explicitly set. An initialization exception will be thrown if, for instance, the `persistent` attribute is set to `false`, yet a persistent data policy was specified.

When persisting regions, it is recommended to configure the storage through the `disk-store` element for maximum efficiency. The diskstore is referenced using the `disk-store-ref` attribute. Additionally, the region may perform disk writes synchronously or asynchronously:

```
<gfe:partitioned-region id="persitent-partition" persistent="true" disk-store-ref="myDiskStore" disk-synchronous="true"/>
```

This is discussed further in Section 5.7, “Configuring a Disk Store”

Subscription Interest Policy

GemFire allows configuration of subscriptions to control [peer to peer event handling](#). Spring Data GemFire provides a `<gfe:subscription/>` to set the interest policy on replicated and partitioned regions to either `ALL` or `CACHE_CONTENT`.

```
<gfe:partitioned-region id="subscription-partition">
  <gfe:subscription type="CACHE_CONTENT"/>
</gfe:partitioned-region>
```

Data Eviction and Overflowing

Based on various constraints, each region can have an eviction policy in place for evicting data from memory. Currently, in GemFire, eviction applies to the least recently used entry (also known as [LRU](#)). Evicted entries are either destroyed or paged to disk (also known as *overflow*).

Spring Data GemFire supports all eviction policies (entry count, memory and heap usage) for both `partitioned-region` and `replicated-region` as well as `client-region`, through the nested `eviction` element. For example, to configure a partition to overflow to disk if its size is more than 512 MB, one could use the following configuration:

```
<gfe:partitioned-region id="overflow-partition">
  <gfe:eviction type="MEMORY_SIZE" threshold="512" action="OVERFLOW_TO_DISK"/>
</gfe:partitioned-region>
```



Important

Replicas cannot use a `local destroy` eviction since that would invalidate them. See the GemFire docs for more information.

When configuring regions for overflow, it is recommended to configure the storage through the `disk-store` element for maximum efficiency.

For a detailed description of eviction policies, see the GemFire documentation (such as [this](#) page).

Data Expiration

GemFire allows you to control how long entries exist in the cache. Eviction is driven by elapsed time, as opposed to eviction which is driven by memory usage. Once an entry expires it may no longer be accessed from the cache. GemFire supports the following expiration types:

- *Time to live (TTL)* - The amount of time, in seconds, the object may remain in the cache after the last creation or update. For entries, the counter is set to zero for create and put operations. Region counters are reset when the region is created and when an entry has its counter reset.
- *Idle timeout* - The amount of time, in seconds, the object may remain in the cache after the last access. The idle timeout counter for an object is reset any time its TTL counter is reset. In addition, an entry's idle timeout counter is reset any time the entry is accessed through a get operation or a netSearch . The idle timeout counter for a region is reset whenever the idle timeout is reset for one of its entries.

Each of these may be applied to the region itself or entries in the region. Spring Data GemFire provides `<region-ttl>`, `<region-tti>`, `<entry-ttl>` and `<entry-tti>` region child elements to specify timeout values and expiration actions.

Local Region

Spring Data GemFire offers a dedicated `local-region` element for creating local regions. Local regions, as the name implies, are standalone meaning they do not share data with any other distributed system member. Other than that, all common region configuration options are supported. A minimal declaration looks as follows (again, the example relies on the Spring Data GemFire namespace naming conventions to wire the cache):

```
<gfe:local-region id="myLocalRegion" />
```

Here, a local region is created (if one doesn't exist already). The name of the region is the same as the bean id (myLocalRegion) and the bean assumes the existence of a GemFire cache named `gemfireCache`.

Replicated Region

One of the common region types is a *replicated region* or *replica*. In short, when a region is configured to be a replicated region, every member that hosts that region stores a copy of the region's entries locally. Any update to a replicated region is distributed to all copies of the region. When a replica is created, it goes through an initialization stage in which it discovers other replicas and automatically copies all the entries. While one replica is initializing you can still continue to use the other rep

Spring Data GemFire offers a `replicated-region` element. A minimal declaration looks as follows. All common configuration options are available for replicated regions.

```
<gfe:replicated-region id="simpleReplica" />
```

Partitioned Region

Another region type supported out of the box by the Spring Data GemFire namespace is the partitioned region. To quote the GemFire docs:

"A partitioned region is a region where data is divided between peer servers hosting the region so that each peer stores a subset of the data. When using a partitioned region, applications are presented with

a logical view of the region that looks like a single map containing all of the data in the region. Reads or writes to this map are transparently routed to the peer that hosts the entry that is the target of the operation. [...] GemFire divides the domain of hashcodes into buckets. Each bucket is assigned to a specific peer, but may be relocated at any time to another peer in order to improve the utilization of resources across the cluster."

A partition is created using the `partitioned-region` element. Its configuration options are similar to that of the `replicated-region` plus the partition specific features such as the number of redundant copies, total maximum memory, number of buckets, partition resolver and so on. Below is a quick example on setting up a partition region with 2 redundant copies:

```
<!-- bean definition named 'distributed-partition' backed by a region named 'redundant'
with 2 copies
and a nested resolver declaration -->
<gfe:partitioned-region id="distributed-partition" copies="2" total-
buckets="4" name="redundant">
  <gfe:partition-resolver>
    <bean class="some.pkg.SimplePartitionResolver"/>
  </gfe:partition-resolver>
</gfe:partitioned-region>
```

partitioned-region Options

The following table offers a quick overview of configuration options specific to partitioned regions. These are in addition to the common region configuration options described above.

Table 5.2. *partitioned-region options*

Name	Values	Description
partition-resolver	<i>bean name</i>	The name of the partitioned resolver used by this region, for custom partitioning.
partition-listener	<i>bean name</i>	The name of the partitioned listener used by this region, for handling partition events.
copies	0..4	The number of copies for each partition for high-availability. By default, no copies are created meaning there is no redundancy. Each copy provides extra backup at the expense of extra storage.
colocated-with	<i>valid region name</i>	The name of the partitioned region with which this newly created partitioned region is colocated.
local-max-memory	<i>positive integer</i>	The maximum amount of memory, in megabytes, to be used by the region in <i>this</i> process.
total-max-memory	<i>any integer value</i>	The maximum amount of memory, in megabytes, to be used by the region in <i>all</i> processes.

Name	Values	Description
recovery-delay	<i>any long value</i>	The delay in milliseconds that existing members will wait before satisfying redundancy after another member crashes. -1 (the default) indicates that redundancy will not be recovered after a failure.
startup-recovery-delay	<i>any long value</i>	The delay in milliseconds that new members will wait before satisfying redundancy. -1 indicates that adding new members will not trigger redundancy recovery. The default is to recover redundancy immediately when a new member is added.

Client Region

GemFire supports various deployment topologies for managing and distributing data. The topic is outside the scope of this documentation however to quickly recap, they can be classified in short in: peer-to-peer (p2p), client-server, and wide area cache network (or WAN). In the last two scenarios, it is common to declare *client* regions which connect to a cache server. Spring Data GemFire offers dedicated support for such configuration through the section called “Configuring a GemFire Client Cache”, `client-region` and `pool` elements. As the names imply, the former defines a client region while the latter defines connection pools to be used/shared by the various client regions.

Below is a typical client region configuration:

```
<!-- client region using the default client-cache pool -->
<gfe:client-region id="simple">
  <gfe:cache-listener ref="c-listener"/>
</gfe:client-region>

<!-- region using its own dedicated pool -->
<gfe:client-region id="complex" pool-name="gemfire-pool">
  <gfe:cache-listener ref="c-listener"/>
</gfe:client-region>

<bean id="c-listener" class="some.pkg.SimpleCacheListener"/>

<!-- pool declaration -->
<gfe:pool id="gemfire-pool" subscription-enabled="true">
  <gfe:locator host="someHost" port="40403"/>
</gfe:pool>
```

As with the other region types, `client-region` supports `CacheListeners` as well as a single `CacheLoader` or `CacheWriter`. It also requires a connection `pool` for connecting to a server. Each client can have its own pool or they can share the same one.



Note

In the above example, the pool is configured with a `locator`. The `locator` is a separate process used to discover cache servers in the distributed system and are recommended for production

systems. It is also possible to configure the pool to connect directly to one or more cache servers using the `server` element.

For a full list of options to set on the client and especially on the pool, please refer to the Spring Data GemFire schema (Appendix A, *Spring Data GemFire Schema*) and the GemFire documentation.

Client Interests

To minimize network traffic, each client can define its own 'interest', pointing out to GemFire, the data it actually needs. In Spring Data GemFire, interests can be defined for each client, both key-based and regular-expression-based types being supported; for example:

```
<gfe:client-region id="complex" pool-name="gemfire-pool">
  <gfe:key-interest durable="true" result-policy="KEYS">
    <bean id="key" class="java.lang.String">
      <constructor-arg value="someKey" />
    </bean>
  </gfe:key-interest>
  <gfe:regex-interest pattern=".*" receive-values="false"/>
</gfe:client-region>
```

A special key `ALL_KEYS` means interest is registered for all keys (identical to a regex interest of `.*`). The `receive-values` attribute indicates whether or not the values are received for create and update events. If true, values are received; if false, only invalidation events are received - refer to the GemFire documentation for more details.

JSON Support

Gemfire 7.0 introduced support for caching JSON documents with OQL query support. These are stored internally as [PdxInstance](#) types using the [JSONFormatter](#) to perform conversion to and from JSON strings. Spring Data GemFire provides a `<gfe-data:json-region-autoproxy/>` tag to enable a [Spring AOP](#) component to advise appropriate region operations, effectively encapsulating the `JSONFormatter`, allowing your application to work directly with JSON strings. In addition, Java objects written to JSON configured regions will be automatically converted to JSON using the Jackson `ObjectMapper`. Reading these values will return a JSON string.

By default, `<gfe-data:json-region-autoproxy/>` will perform the conversion on all regions. To apply this feature to selected regions, provide a comma delimited list of their ids via the `region-refs` attribute. Other attributes include a `pretty-print` flag (false by default) and `convert-returned-collections`. By default the results of region operations `getAll()` and `values()` will be converted for configured regions. This is done by creating a parallel structure in local memory. This can incur significant overhead for large collections. Set this flag to false to disable automatic conversion for these operation.



Note

Certain region operations, specifically those that use GemFire's proprietary `Region.Entry` such as `entries(boolean)`, `entrySet(boolean)` and `getEntry()` type are not targeted for AOP advice. In addition, the `entrySet()` method which returns a `Set<java.util.Map.Entry<?,?>>` is not affected.

```
<gfe-data:json-region-autoproxy pretty-print="true" region-
refs="myJsonRegion" convert-returned-collections="true"/>
```

This feature also works with seamlessly with GemfireTemplate operations, provided that the template is declared as a Spring bean. Currently native QueryService operations are not supported.

5.6 Creating an Index

GemFire allows creation on indexes (or indices) to improve the performance of (common) queries. Spring Data GemFire allows indecies to be declared through the `index` element:

```
<gfe:index id="myIndex" expression="someField" from="/someRegion"/>
```

Before creating an index, Spring Data GemFire will verify whether one with the same name already exists. If it does, it will compare the properties and if they don't match, will remove the old one to create a new one. If the properties match, Spring Data GemFire will simply return the index (in case it does not exist it will simply create one). To prevent the update of the index, even if the properties do not match, set the property `override` to `false`.

Note that index declaration are not bound to a region but rather are top-level elements (just like `gfe:cache`). This allows one to declare any number of indecies on any region whether they are just created or already exist - an improvement versus the GemFire `cache.xml`. By default the index relies on the default cache declaration but one can customize it accordingly or use a pool (if need be) - see the namespace schema for the full set of options.

5.7 Configuring a Disk Store

As of Release 1.2.0, Spring Data GemFire supports disk store configuration via a top level `disk-store` element.



Note

Prior to Release 1.2.0, `disk-store` was a child element of `*-region`. If you have regions configured with disk storage using a prior release of Spring Data GemFire and want to upgrade to the latest release, move the `disk-store` element to the top level, assign an id and use the region's `disk-store-ref` attribute. Also, `disk-synchronous` is now a region level attribute.

```
<gfe:disk-store id="diskStore1" queue-size="50" auto-compact="true"
  max-oplog-size="10" time-interval="9999">
  <gfe:disk-dir location="/gemfire/store1/" max-size="20"/>
  <gfe:disk-dir location="/gemfire/store2/" max-size="20"/>
</gfe:disk-store>
```

Disk stores are used by regions for file system persistent backup or overflow storage of evicted entries, and persistent backup of WAN gateways. Note that multiple components may share the same disk store. Also multiple directories may be defined for a single disk store. Please refer to the GemFire documentation for an explanation of the configuration options.

5.8 Configuring GemFire's Function Service

As of Release 1.3.0, Spring Data GemFire provides [annotation](#) support for implementing and registering functions. Spring Data GemFire also provides namespace support for registering GemFire [Functions](#) for remote function execution. Please refer to the GemFire documentation for more information on the function execution framework. Functions are declared as Spring beans and

must implement the `com.gemstone.gemfire.cache.execute.Function` interface or extend `com.gemstone.gemfire.cache.execute.FunctionAdapter`. The namespace uses a familiar pattern to declare functions:

```
<gfe:function-service>
  <gfe:function>
    <bean class="com.company.example.Function1"/>
    <ref bean="function2"/>
  </gfe:function>
</gfe:function-service>

<bean id="function2" class="com.company.example.Function2"/>
```

5.9 Configuring WAN Gateways

WAN gateways provide a way to synchronize GemFire distributed systems across geographic distributed areas. As of Release 1.2.0, Spring Data GemFire provides namespace support for configuring WAN gateways as illustrated in the following examples:

WAN Configuration in GemFire 7.0

GemFire 7.0 introduces new APIs for WAN configuration. While the original APIs provided in GemFire 6 are still supported, it is recommended that you use the new APIs if you are using GemFire 7.0. The Spring Data GemFire namespace supports either. In the example below, `GatewaySenders` are configured for a partitioned region by adding child elements to the region (`gateway-sender` and `gateway-sender-ref`). The `GatewaySender` may register `EventFilters` and `TransportFilters`. Also shown below is an example configuration of an `AsyncEventQueue` which must also be wired into a region (not shown).

```

<gfe:partitioned-region id="region-inner-gateway-sender" >
  <gfe:gateway-sender
    remote-distributed-system-id="1">
    <gfe:event-filter>
      <bean class="org.springframework.data.gemfire.example.SomeEventFilter"/>
    </gfe:event-filter>
    <gfe:transport-filter>

    <bean class="org.springframework.data.gemfire.example.SomeTransportFilter"/>
    </gfe:transport-filter>
  </gfe:gateway-sender>
  <gfe:gateway-sender-ref bean="gateway-sender"/>
</gfe:partitioned-region>

<gfe:async-event-queue id="async-event-queue" batch-size="10" persistent="true" disk-
store-ref="diskstore"
  maximum-queue-memory="50">
  <gfe:async-event-listener>
    <bean class="org.springframework.data.gemfire.example.SomeAsyncEventListener"/>
  </gfe:async-event-listener>
</gfe:async-event-queue>

<gfe:gateway-sender id="gateway-sender" remote-distributed-system-id="2">
  <gfe:event-filter>
    <ref bean="event-filter"/>
    <bean class="org.springframework.data.gemfire.example.SomeEventFilter"/>
  </gfe:event-filter>
  <gfe:transport-filter>
    <ref bean="transport-filter"/>
    <bean class="org.springframework.data.gemfire.example.SomeTransportFilter"/>
  </gfe:transport-filter>
</gfe:gateway-sender>

<bean id="event-
filter" class="org.springframework.data.gemfire.example.AnotherEventFilter"/>
<bean id="transport-
filter" class="org.springframework.data.gemfire.example.AnotherTransportFilter"/>

```

On the other end of a GatewaySender is a corresponding GatewayReceiver to receive gateway events. The GatewayReceiver may also be configured with EventFilters and TransportFilters.

```

<gfe:gateway-receiver id="gateway-receiver"
  start-port="12345" end-port="23456" bind-address="192.168.0.1">
  <gfe:transport-filter>
    <bean class="org.springframework.data.gemfire.example.SomeTransportFilter"/>
  </gfe:transport-filter>
</gfe:gateway-receiver>

```

Please refer to the GemFire product document for a detailed explanation of all the configuration options.

WAN Configuration in GemFire 6.6

```
<gfe:cache/>

<gfe:replicated-region id="region-with-gateway" enable-gateway="true" hub-id="gateway-
hub"/>

<gfe:gateway-hub id="gateway-hub" manual-start="true">
  <gfe:gateway gateway-id="gateway">
    <gfe:gateway-listener>
      <bean class="com.company.example.MyGatewayListener"/>
    </gfe:gateway-listener>
    <gfe:gateway-queue maximum-queue-memory="5" batch-size="3"
      batch-time-interval="10" />
  </gfe:gateway>

  <gfe:gateway gateway-id="gateway2">
    <gfe:gateway-endpoint port="1234" host="host1" endpoint-id="endpoint1"/>
    <gfe:gateway-endpoint port="2345" host="host2" endpoint-id="endpoint2"/>
  </gfe:gateway>
</gfe:gateway-hub>
```

A region may synchronize all or part of its contents to a gateway hub used to access one or more remote systems. The region must set `enable-gateway` to `true` and specify the `hub-id`.



Note

If just a `hub-id` is specified, Spring Data GemFire automatically assumes that the gateway should be enabled.

Please refer to the GemFire product document for a detailed explanation of all the configuration options.

6. Working with the GemFire APIs

Once the GemFire cache and regions have been configured they can be injected and used inside application objects. This chapter describes the integration with Spring's transaction management functionality and `DaoException` hierarchy. It also covers support for dependency injection of GemFire managed objects.

6.1 Exception translation

Using a new data access technology requires not just accommodating to a new API but also handling exceptions specific to that technology. To accommodate this case, Spring Framework provides a technology agnostic, consistent exception [hierarchy](#) that abstracts one from proprietary (and usually checked) exceptions to a set of focused runtime exceptions. As mentioned in the Spring Framework documentation, [exception translation](#) can be applied transparently to your data access objects through the use of the `@Repository` annotation and AOP by defining a `PersistenceExceptionTranslationPostProcessor` bean. The same exception translation functionality is enabled when using GemFire as long as at least a `CacheFactoryBean` is declared, e.g., using a `<gfe:cache/>` declaration) as it acts as an exception translator which is automatically detected by the Spring infrastructure and used accordingly.

6.2 GemfireTemplate

As with many other high-level abstractions provided by the Spring projects, Spring Data GemFire provides a *template* that simplifies GemFire data access. The class provides several *one-line* methods, for common region operations but also the ability to *execute* code against the native GemFire API without having to deal with GemFire checked exceptions for example through the `GemfireCallback`.

The template class requires a GemFire `Region` instance and once configured is thread-safe and should be reused across multiple classes:

```
<bean id="gemfireTemplate" class="org.springframework.data.gemfire.GemfireTemplate" p:region-ref="someRegion"/>
```

Once the template is configured, one can use it alongside `GemfireCallback` to work directly with the GemFire `Region`, without having to deal with checked exceptions, threading or resource management concerns:

```
template.execute(new GemfireCallback<Iterable<String>>() {
    public Iterable<String> doInGemfire(Region reg) throws GemFireCheckedException,
        GemFireException {
        // working against a Region of String
        Region<String, String> region = reg;

        region.put("1", "one");
        region.put("3", "three");

        return region.query("length < 5");
    }
});
```

For accessing the full power of the GemFire query language, one can use the `find` and `findUnique` which, as opposed to the `query` method, can execute queries across multiple regions, execute projections, and the like. The `find` method should be used when the query selects multiple items

(through `SelectResults`) and the latter, `findUnique`, as the name suggests, when only one object is returned.

6.3 Support for Spring Cache Abstraction

Since 1.1, Spring GemFire provides an implementation for Spring 3.1 [cache abstraction](#). To use GemFire as a backing implementation, simply add `GemfireCacheManager` to your configuration:

```
<beans xmlns="http://www.springframework.org/schema/beans" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance"
    xmlns:cache="http://www.springframework.org/schema/cache"
    xmlns:gfe="http://www.springframework.org/schema/gemfire"
    xmlns:p="http://www.springframework.org/schema/p"
    xsi:schemaLocation="http://www.springframework.org/schema/beans http://
www.springframework.org/schema/beans/spring-beans.xsd
        http://www.springframework.org/schema/gemfire http://www.springframework.org/schema/
gemfire/spring-gemfire.xsd
        http://www.springframework.org/schema/cache http://www.springframework.org/schema/
cache/spring-cache.xsd">
    <!-- turn on declarative caching -->
    <cache:annotation-driven />

    <gfe:cache id="gemfire-cache" />
    <!-- declare GemFire Cache Manager -->
    <bean id="cacheManager" class="org.springframework.data.gemfire.support.GemfireCacheManager" p:cache-
ref="gemfire-cache">
</beans>
```

6.4 Transaction Management

One of the most popular features of Spring Framework is [transaction](#) management. If you are not familiar with it, we strongly recommend [looking](#) into it as it offers a consistent programming model that works transparently across multiple APIs that can be configured either programmatically or declaratively (the most popular choice).

For GemFire, Spring Data GemFire provides a dedicated, per-cache, transaction manager that once declared, allows region operations to be executed atomically through Spring:

```
<gfe:transaction-manager id="tx-manager" cache-ref="cache"/>
```



Note

The example above can be simplified even more by eliminating the `cache-ref` attribute if the GemFire cache is defined under the default name `gemfireCache`. As with the other Spring Data GemFire namespace elements, if the cache name is not configured, the aforementioned naming convention will be used. Additionally, the transaction manager name, if not specified, is `gemfireTransactionManager`.

Note that currently GemFire supports optimistic transactions with *read committed* isolation. Furthermore, to guarantee this isolation, developers should avoid making *in-place* changes, that is manually modifying the values present in the cache. To prevent this from happening, the transaction manager configured the cache to use *copy on read* semantics, meaning a clone of the actual value is created, each time a read is performed. This behavior can be disabled if needed through the `copyOnRead` property. For

more information on the semantics of the underlying GemFire transaction manager, see the [GemFire documentation](#).

6.5 GemFire Continuous Query Container

A powerful functionality offered by GemFire is [continuous querying](#) (or CQ). In short, CQ allows one to create a query and automatically be notified when new data that gets added to GemFire matches the query. Spring GemFire provides dedicated support for CQs through the `org.springframework.data.gemfire.listener` package and its *listener container*; very similar in functionality and naming to the JMS integration in Spring Framework; in fact, users familiar with the JMS support in Spring, should feel right at home. Basically Spring Data GemFire allows methods on POJOs to become end-points for CQ - simply define the query and indicate the method that should be notified when there is a match - Spring Data GemFire takes care of the rest. This is similar Java EE's message-driven bean style, but without any requirement for base class or interface implementations, based on GemFire.



Note

Currently, continuous queries are supported by GemFire only in client/server topologies. Additionally the pool used is required to have the `subscription` property enabled. Please refer to the documentation for more information.

Continuous Query Listener Container

Spring Data GemFire simplifies the creation, registration, life-cycle and dispatch of CQs by taking care of the infrastructure around them through `ContinuousQueryListenerContainer` which does all the heavy lifting on behalf of the user - users familiar with EJB and JMS should find the concepts familiar as it is designed as close as possible to the support in Spring Framework and its message-driven POJOs (MDPs)

`ContinuousQueryListenerContainer` acts as an event (or message) listener container; it is used to receive the events from the registered CQs and drive the POJOs that are injected into it. The listener container is responsible for all threading of message reception and dispatches into the listener for processing. It acts as the intermediary between an EDP (Event Driven POJO) and the event provider and takes care of creation and registration of CQs (to receive events), resource acquisition and release, exception conversion and the like. This allows you as an application developer to write the (possibly complex) business logic associated with receiving an event (and reacting to it), and delegates boilerplate GemFire infrastructure concerns to the framework.

The container is fully customizable - one can chose either to use the CQ thread to perform the dispatch (synchronous delivery) or a new thread (from an existing pool for examples) for an asynchronous approach by defining the suitable `java.util.concurrent.Executor` (or Spring's `TaskExecutor`). Depending on the load, the number of listeners or the runtime environment, one should change or tweak the executor to better serve her needs - in particular in managed environments (such as app servers), it is highly recommended to pick a proper `TaskExecutor` to take advantage of its runtime.

The `ContinuousQueryListenerAdapter` and `ContinuousQueryListener`

The `ContinuousQueryListenerAdapter` class is the final component in Spring Data GemFire CQ support: in a nutshell, it allows you to expose almost *any* class as a EDP (there are of course

some constraints) - it implements `ContinuousQueryListener`, a simpler listener interface similar to GemFire [CqListener](#).

Consider the following interface definition. Notice the various event handling methods and their parameters:

```
public interface EventDelegate {
    void handleEvent(CqEvent event);
    void handleEvent(Operation baseOp);
    void handleEvent(Object key);
    void handleEvent(Object key, Object newValue);
    void handleEvent(Throwable th);
    void handleQuery(CqQuery cq);
    void handleEvent(CqEvent event, Operation baseOp, byte[] deltaValue);
    void handleEvent(CqEvent event, Operation baseOp, Operation queryOp, Object key,
        Object newValue);
}
```

```
public class DefaultEventDelegate implements EventDelegate {
    // implementation elided for clarity...
}
```

In particular, note how the above implementation of the `EventDelegate` interface (the above `DefaultEventDelegate` class) has *no* GemFire dependencies at all. It truly is a POJO that we will make into an EDP via the following configuration (note that the class doesn't have to implement an interface, one is present only to better show case the decoupling between contract and implementation).

```
<?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:gfe="http://www.springframework.org/schema/gemfire"
    xsi:schemaLocation="http://www.springframework.org/schema/beans http://
www.springframework.org/schema/beans/spring-beans.xsd
    http://www.springframework.org/schema/gemfire http://www.springframework.org/
schema/gemfire/spring-gemfire.xsd">

    <gfe:client-cache pool-name="client"/>

    <gfe:pool id="client" subscription-enabled="true">
        <gfe:server host="localhost" port="40404"/>
    </gfe:pool>

    <gfe:cq-listener-container>
        <!-- default handle method -->
        <gfe:listener ref="listener" query="SELECT * from /region" />
        <gfe:listener ref="another-listener" query="SELECT * from /another-region" name="my-
query" method="handleQuery"/>
    </gfe:cq-listener-container>

    <bean id="listener" class="gemfireexample.DefaultMessageDelegate"/>
    <bean id="another-listener" class="gemfireexample.DefaultMessageDelegate"/>
    ...
</beans>
```



Note

The example above shows some of the various forms that a listener can have; at its minimum the listener reference and the actual query definition are required. It's possible however to specify a

name for the resulting continuous query (useful for monitoring) but also the name of the method (the default is `handleEvent`). The specified method can have various argument types, the `EventDelegate` interface lists the allowed types.

The example above uses the Spring Data GemFire namespace to declare the event listener container and automatically register the listeners. The full blown, *beans* definition is displayed below:

```
<!-- this is the Event Driven POJO (MDP) -->
<bean id="eventListener"
  class="org.springframework.data.gemfire.listener.adapter.ContinuousQueryListenerAdapter">
  <constructor-arg>
    <bean class="gemfireexample.DefaultEventDelegate"/>
  </constructor-arg>
</bean>

<!-- and this is the event listener container... -->
<bean id="gemfireListenerContainer" class="org.springframework.data.gemfire.listener.ContinuousQueryListenerContainer">
  <property name="cache" ref="gemfireCache"/>
  <property name="queryListeners">
    <!-- set of listeners -->
    <set>
      <bean class="org.springframework.data.gemfire.listener.ContinuousQueryDefinition">
        <constructor-arg value="SELECT * from /region" />
        <constructor-arg ref="eventListener" />
      </bean>
    </set>
  </property>
</bean>
```

Each time an event is received, the adapter automatically performs type translation between the GemFire event and the required method argument(s) transparently. Any exception caused by the method invocation is caught and handled by the container (by default, being logged).

6.6 Wiring Declarable components

GemFire XML configuration (usually named `cache.xml`) allows *user* objects to be declared as part of the configuration. Usually these objects are `CacheLoaders` or other pluggable callback components supported by GemFire. Using native GemFire configuration, each user type declared through XML must implement the `Declarable` interface which allows arbitrary parameters to be passed to the declared class through a `Properties` instance.

In this section we describe how you can configure these pluggable components defined in `cache.xml` using Spring while keeping your `Cache/Region` configuration defined in `cache.xml`. This allows your pluggable components to focus on the application logic and not the location or creation of `DataSources` or other collaboration objects.

However, if you are starting a green field project, it is recommended that you configure `Cache`, `Region`, and other pluggable components directly in Spring. This avoids inheriting from the `Declarable` interface or the base class presented in this section. See the following sidebar for more information on this approach.

Eliminate Declarable components

One can configure custom types entirely through Spring as mentioned in Section 5.5, “Configuring a GemFire Region”. That way, one does not have to implement the `Declarable` interface and also benefits from all the features of the Spring IoC container (not just dependency injection but also life-cycle and instance management).

As an example of configuring a `Declarable` component using Spring, consider the following declaration (taken from the `Declarable` javadoc):

```
<cache-loader>
  <class-name>com.company.app.DBLoader</class-name>
  <parameter name="URL">
    <string>jdbc://12.34.56.78/mydb</string>
  </parameter>
</cache-loader>
```

To simplify the task of parsing, converting the parameters and initializing the object, Spring Data GemFire offers a base class (`WiringDeclarableSupport`) that allows GemFire user objects to be wired through a *template* bean definition or, in case that is missing, perform autowiring through the Spring container. To take advantage of this feature, the user objects need to extend `WiringDeclarableSupport` which automatically locates the declaring `BeanFactory` and performs wiring as part of the initialization process.

Why is a base class needed?

In the current GemFire release there is no concept of an *object factory* and the types declared are instantiated and used as is. In other words, there is no easy way to manage object creation outside GemFire.

Configuration using *template* definitions

When used, `WiringDeclarableSupport` tries to first locate an existing bean definition and use that as wiring template. Unless specified, the component class name will be used as an implicit bean definition name. Let's see how our `DBLoader` declaration would look in that case:

```
public class DBLoader extends WiringDeclarableSupport implements CacheLoader {
  private DataSource dataSource;

  public void setDataSource(DataSource ds){
    this.dataSource = ds;
  }

  public Object load(LoaderHelper helper) { ... }
}
```

```
<cache-loader>
  <class-name>com.company.app.DBLoader</class-name>
  <!-- no parameter is passed (use the bean implicit name
       that is the class name) -->
</cache-loader>
```

```
<?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:p="http://www.springframework.org/schema/p"
       xsi:schemaLocation="http://www.springframework.org/schema/beans
                           http://www.springframework.org/schema/beans/spring-beans.xsd">

    <bean id="dataSource" ... />

    <!-- template bean definition -->
    <bean id="com.company.app.DBLoader" abstract="true" p:dataSource-ref="dataSource"/>
</beans>
```

In the scenario above, as no parameter was specified, a bean with the id/name `com.company.app.DBLoader` was used as a template for wiring the instance created by GemFire. For cases where the bean name uses a different convention, one can pass in the `bean-name` parameter in the GemFire configuration:

```
<cache-loader>
    <class-name>com.company.app.DBLoader</class-name>
    <!-- pass the bean definition template name
         as parameter -->
    <parameter name="bean-name">
        <string>template-bean</string>
    </parameter>
</cache-loader>
```

```
<?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:p="http://www.springframework.org/schema/p"
       xsi:schemaLocation="http://www.springframework.org/schema/beans
                           http://www.springframework.org/schema/beans/spring-beans.xsd">

    <bean id="dataSource" ... />

    <!-- template bean definition -->
    <bean id="template-bean" abstract="true" p:dataSource-ref="dataSource"/>

</beans>
```



Note

The *template* bean definitions do not have to be declared in XML - any format is allowed (Groovy, annotations, etc..).

Configuration using auto-wiring and annotations

If no bean definition is found, by default, `WiringDeclarableSupport` will [autowire](#) the declaring instance. This means that unless any dependency injection *metadata* is offered by the instance, the container will find the object setters and try to automatically satisfy these dependencies. However, one can also use JDK 5 annotations to provide additional information to the auto-wiring process. We strongly recommend reading the dedicated [chapter](#) in the Spring documentation for more information on the supported annotations and enabling factors.

For example, the hypothetical `DBLoader` declaration above can be injected with a Spring-configured `DataSource` in the following way:


```
public class DBLoader extends WiringDeclarableSupport implements CacheLoader {  
    // use annotations to 'mark' the needed dependencies  
    @javax.inject.Inject  
    private DataSource dataSource;  
  
    public Object load(LoaderHelper helper) { ... }  
}
```

```
<cache-loader>  
    <class-name>com.company.app.DBLoader</class-name>  
    <!-- no need to declare any parameters anymore  
         since the class is auto-wired -->  
</cache-loader>
```

```
<?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"  
    xsi:schemaLocation="http://www.springframework.org/schema/beans  
        http://www.springframework.org/schema/beans/spring-beans.xsd  
        http://www.springframework.org/schema/context  
        http://www.springframework.org/schema/context/spring-context.xsd">  
  
    <!-- enable annotation processing -->  
    <context:annotation-config/>  
  
</beans>
```

By using the JSR-330 annotations, the cache loader code has been simplified since the location and creation of the `DataSource` has been externalized and the user code is concerned only with the loading process. The `DataSource` might be transactional, created lazily, shared between multiple objects or retrieved from JNDI - these aspects can be easily configured and changed through the Spring container without touching the `DBLoader` code.

7. Working with GemFire Serialization

To improve overall performance of the data grid, GemFire supports a dedicated serialization protocol (PDX) that is both faster and offers more compact results over the standard Java serialization and works transparently across various language [platforms](#) (such as [Java](#), [.NET](#) and C++). This chapter discusses the various ways in which Spring Data GemFire simplifies and improves GemFire custom serialization in Java.

7.1 Wiring deserialized instances

It is fairly common for serialized objects to have transient data. Transient data is often dependent on the node or environment where it lives at a certain point in time, for example a `DataSource`. Serializing such information is useless (and potentially even dangerous) since it is local to a certain VM/machine. For such cases, Spring Data GemFire offers a special [Instantiator](#) that performs wiring for each new instance created by GemFire during deserialization.

Through such a mechanism, one can rely on the Spring container to inject (and manage) certain dependencies making it easy to split transient from persistent data and have *rich domain objects* in a transparent manner (Spring users might find this approach similar to that of [@Configurable](#)). The `WiringInstantiator` works just like `WiringDeclarableSupport`, trying to first locate a bean definition as a wiring template and following to autowiring otherwise. Please refer to the previous section (Section 6.6, “Wiring Declarable components”) for more details on wiring functionality.

To use this `Instantiator`, simply declare it as a usual bean:

```
<bean id="instantiator" class="org.springframework.data.gemfire.serialization.WiringInstantiator">
  <!-- DataSerializable type -->
  <constructor-arg>org.pkg.SomeDataSerializableClass</constructor-arg>
  <!-- type id -->
  <constructor-arg>95</constructor-arg>
</bean>
```

During the container startup, once it is being initialized, the `instantiator` will, by default, register itself with the GemFire system and perform wiring on all instances of `SomeDataSerializableClass` created by GemFire during deserialization.

7.2 Auto-generating custom Instantiators

For data intensive applications, a large number of instances might be created on each machine as data flows in. Out of the box, GemFire uses reflection to create new types but for some scenarios, this might prove to be expensive. As always, it is good to perform profiling to quantify whether this is the case or not. For such cases, Spring Data GemFire allows the automatic generation of `Instantiator` classes which instantiate a new type (using the default constructor) without the use of reflection:

```
<bean id="instantiator-
factory" class="org.springframework.data.gemfire.serialization.InstantiatorFactoryBean">
  <property name="customTypes">
    <map>
      <entry key="org.pkg.CustomTypeA" value="1025"/>
      <entry key="org.pkg.CustomTypeB" value="1026"/>
    </map>
  </property>
</bean>
```

The definition above, automatically generated two `Instantiators` for two classes, namely `CustomTypeA` and `CustomTypeB` and registers them with `GemFire`, under user id 1025 and 1026. The two instantiators avoid the use of reflection and create the instances directly through Java code.

8. POJO mapping

8.1 Entity mapping

Spring Data GemFire provides support to map entities to be stored in a GemFire grid. The mapping metadata is define by using annotations at the domain classes just like this:

```
@Region("myRegion")
public class Person {

    @Id Long id;
    String firstname;
    String lastname;

    @PersistenceConstructor
    public Person(String firstname, String lastname) {
        // ...
    }

    ...
}
```

Example 8.1 Mapping a domain class to GemFire

The first thing you see here is the `@Region` annotation that can be used to customize the region instances of the `Person` class are stored in. The `@Id` annotation can be used to annotate the property that shall be used as cache key. The `@PersistenceConstructor` annotation actually helps disambiguating multiple potentially available constructors taking parameters and explicitly marking the one annotated as the one to be used to create entities. With none or only a single constructor you can omit the annotation.

8.2 Mapping PDX serializer

Spring Data GemFire provides a custom `PDXSerializer` implementation that uses the mapping information to customize entity serialization. Beyond that it allows customizing the entity instantiation by using the Spring Data `EntityInstantiator` abstraction. By default the serializer uses a `ReflectionEntityInstantiator` that will use the persistence constructor of the mapped entity (either the single declared one or explicitly annotated with `@PersistenceConstructor`). To provide values for constructor parameters it will read fields with name of the constructor parameters from the `PDXReader` supplied.

```
public class Person {

    public Person(@Value("#root.foo") String firstname, @Value("bean") String lastname) {
        // ...
    }

}
```

Example 8.2 Using @Value on entity constructor parameters

The entity annotated as such will get the field `foo` read from the `PDXReader` and handed as constructor parameter value for `firstname`. The value for `lastname` will be the Spring bean with name `bean`.

9. GemFire Repositories

9.1 Introduction

Spring Data GemFire provides support to use the Spring Data repository abstraction to easily persist entities into GemFire and execute queries. A general introduction into the repository programming model is been provided [here](#).

9.2 Spring configuration

To bootstrap Spring Data repositories you use the `<repositories />` element from the GemFire namespace:

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

    <gfe-data:repositories base-package="com.acme.repository" />

</beans>
```

Example 9.1 Bootstrap GemFire repositories

This configuration snippet will look for interfaces below the configured base package and create repository instances for those interfaces backed by a `SimpleGemFireRepository`. Note that you have to have your domain classes correctly mapped to configured regions as the bootstrap process will fail otherwise.

9.3 Executing OQL queries

The GemFire repositories allow the definition of query methods to easily execute OQL queries against the Region the managed entity is mapped to.

```
@Region("myRegion")
public class Person { ... }
```

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

    Person findByEmailAddress(String emailAddress);

    Collection<Person> findByFirstname(String firstname);

    @Query("SELECT * FROM /Person p WHERE p.firstname = $1")
    Collection<Person> findByFirstnameAnnotated(String firstname);

    @Query("SELECT * FROM /Person p WHERE p.firstname IN SET $1")
    Collection<Person> findByFirstnamesAnnotated(Collection<String> firstnames);

}
```

Example 9.2 Sample repository

The first method listed here will cause the following query to be derived: `SELECT x FROM /myRegion x WHERE x.emailAddress = $1`. The second method works the same way except it's returning all entities found whereas the first one expects a single result value. In case the supported keywords are not sufficient to declare your query or the method name gets too verbose you can annotate the query methods with `@Query` as seen for methods 3 and 4.

Table 9.1. Supported keywords for query methods

Keyword	Sample	Logical result
GreaterThan	<code>findByAgeGreaterThan(int age)</code>	<code>x.age > \$1</code>
GreaterThanEqual	<code>findByAgeGreaterThanEqual(int age)</code>	<code>x.age >= \$1</code>
LessThan	<code>findByAgeLessThan(int age)</code>	<code>x.age < \$1</code>
LessThanEqual	<code>findByAgeLessThanEqual(int age)</code>	<code>x.age <= \$1</code>
IsNotNull, NotNull	<code>findByFirstnameNotNull()</code>	<code>x.firstname != NULL</code>
IsNull, Null	<code>findByFirstnameNull()</code>	<code>x.firstname = NULL</code>
In	<code>findByFirstnameIn(Collection<String> x)</code>	<code>x.firstname IN SET \$1</code>
NotIn	<code>findByFirstnameNotIn(Collection<String> x)</code>	<code>x.firstname NOT IN SET \$1</code>
(No keyword)	<code>findByFirstname(String name)</code>	<code>x.firstname = \$1</code>
Like	<code>findByFirstnameLike(String name)</code>	<code>x.firstname LIKE \$1</code>
Not	<code>findByFirstnameNot(String name)</code>	<code>x.firstname != \$1</code>
IsTrue, True	<code>findByActiveIsTrue()</code>	<code>x.active = true</code>
IsFalse, False	<code>findByActiveIsFalse()</code>	<code>x.active = false</code>

10. Annotation Support for Function Execution

10.1 Introduction

Spring Data GemFire 1.3.0 introduces annotation support to simplify working with [GemFire function execution](#). The GemFire API provides classes to implement and register [Functions](#) deployed to cache servers that may be invoked remotely by member applications, typically cache clients. Functions may execute in parallel, distributed among multiple servers, combining results in a map-reduce pattern, or may be targeted to a single server. A Function execution may be also be targeted to a specific region.

GemFire's also provides APIs to support remote execution of functions targeted to various defined scopes (region, member groups, servers, etc.) and the ability to aggregate results. The API also provides certain runtime options. The implementation and execution of remote functions, as with any RPC protocol, requires some boilerplate code. Spring Data GemFire, true to Spring's core value proposition, aims to hide the mechanics of remote function execution and allow developers to focus on POJO programming and business logic. To this end, Spring Data GemFire introduces annotations to declaratively register public methods as functions, and the ability to invoke registered functions remotely via annotated interfaces.

10.2 Implementation vs Execution

There are two separate concerns to address. First is the function implementation (server) which must interact with the [FunctionContext](#) to obtain the invocation arguments, the [ResultsSender](#) and other execution context information. The function implementation typically accesses the Cache and or Region and is typically registered with the [FunctionService](#) under a unique Id. The application invoking a function (the client) does not depend on the implementation. To invoke a function remotely, the application instantiates an [Execution](#) providing the function ID, invocation arguments, the function target or scope (region, server, servers, member, members). If the function produces a result, the invoker uses a [ResultCollector](#) to aggregate and acquire the execution results. In certain scenarios, a custom ResultCollector implementation is required and may be registered with the Execution.

Note

'Client' and 'Server' are used here in the context of function execution which may have a different meaning than client and server in a client-server cache topology. While it is common for a member with a Client Cache to invoke a function on one or more Cache Server members it is also possible to execute functions in a peer-to-peer configuration

10.3 Implementing a Function

Using GemFire APIs, the FunctionContext provides a runtime invocation context including the client's calling arguments and a ResultSender interface to send results back to the client. Additionally, if the function is executed on a Region, the FunctionContext is an instance of RegionFunctionContext which provides additional context such as the target Region and any Filter (set of specific keys) associated with the Execution. If the Region is a Partition Region, the function should use the PartitionRegionHelper to extract only the local data.

Using Spring, one can write a simple POJO and enable the Spring container bind one or more of its public methods to a Function. The signature for a POJO method intended to be used as a function must

generally conform to the the client's execution arguments. However, in the case of a region execution, the region data must also be provided (presumably the data held in the local partition if the region is a partition region). Additionally the function may require the filter that was applied, if any. This suggests that the client and server may share a contract for the calling arguments but that the method signature may include additional parameters to pass values provided by the FunctionContext. One possibility is that the client and server share a common interface, but this is not required. The only constraint is that the method signature includes the same sequence of calling arguments with which the function was invoked after the additional parameters are resolved. For example, suppose the client provides a String and int as the calling arguments. These are provided by the FunctionContext as an array:

```
Object[] args = new Object[]{"hello", 123}
```

Then the Spring container should be able to bind to any method signature similar to the following. Let's ignore the return type for the moment:

```
public Object method1(String s1, int i2) {...}
public Object method2(Map<?,?> data, String s1, int i2) {...}
public Object method3(String s1, Map<?,?>data, int i2) {...}
public Object method4(String s1, Map<?,?> data, Set<?> filter, int i2) {...}
public void method4(String s1, Set<?> filter, int i2, Region<?,?> data) {...}
public void method5(String s1, ResultSender rs, int i2);
public void method6(FunctionContext fc);
```

The general rule is that once any additional arguments, i.e., region data and filter, are resolved the remaining arguments must correspond exactly, in order and type, to the expected calling parameters. The method's return type must be void or a type that may be serialized (either `java.io.Serializable`, `DataSerializable`, or PDX serializable). The latter is also a requirement for the calling arguments. The Region data should normally be defined as a Map, to facilitate unit testing, but may also be of type Region if necessary. As shown in the example above, it is also valid to pass the FunctionContext itself, or the ResultSender, if you need to control how the results are returned to the client.

Annotations for Function Implementation

The following example illustrates how annotations are used to expose a POJO as a GemFire function:

```
@Component
public class MyFunctions {
    @GemfireFunction
    public String function1(String s1, @RegionData Map<?,?> data, int i2) { ... }

    @GemfireFunction("myFunction", HA=true, optimizedForWrite=true, batchSize=100)
    public List<String> function2(String s1, @RegionData Map<?,?> data, int i2, @Filter
    Set<?> keys) { ... }

    @GemfireFunction(hasResult=true)
    public void functionWithContext(FunctionContext functionContext) { ... }
}
```

Note that the class itself must be registered as a Spring bean. Here the `@Component` annotation is used, but you may register the bean by any method provided by Spring (e.g. XML configuration or Java configuration class). This allows the Spring container to create an instance of this class and wrap it in a [PojoFunctionWrapper](#) (PFW). Spring creates one PFW instance for each method annotated with `@GemfireFunction`. Each will all share the same target object instance to invoke the corresponding method.



Note

The fact that the function class is a Spring bean may offer other benefits since it shares the application context with GemFire components such as a Cache and Regions. These may be injected into the class if necessary.

Spring creates the wrapper class, and registers the function with GemFire's Function Service. The function id used to register the functions must be unique. By convention it defaults to the simple (unqualified) method name. Note that this annotation also provides configuration attributes, `HA` and `optimizedForWrite` which correspond to properties defined by GemFire's Function interface. If the method's return type is void, then the `hasResult` property is automatically set to `false`; otherwise it is `true`.

For void return types, the annotation provides a `hasResult` attribute that can be set to true to override this convention, as shown in the `functionWithContext` method above. Presumably, the intention is to use the `ResultSender` directly to send results to the caller.

The PFW implements GemFire's Function interface, binds the method parameters, and invokes the target method in its `execute()` method. It also sends the method's return value using the `ResultSender`.

Batching Results

If the return type is a Collection or Array, then some consideration must be given to how the results are returned. By default, the PFW returns the entire collection at once. If the number of items is large, this may incur a performance penalty. To divide the payload into small sections (sometimes called chunking), you can set the `batchSize` attribute, as illustrated in `function2`, above.



Note

If you need more control of the `ResultSender`, especially if the method itself would use too much memory to create the collection, you can pass the `ResultSender`, or access it via the `FunctionContext`, to use it directly within the method.

Enabling Annotation Processing

In accordance with Spring standards, you must explicitly activate annotation processing for `@GemfireFunction` using XML:

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

or by annotating a Java configuration class:

```
@EnableGemfireFunctions
```

10.4 Executing a Function

A process invoking a remote function needs to provide calling arguments, a function id, the execution target (onRegion, onServers, onServer, onMember, onMembers) and optionally a Filter set. All you need to do is define an interface supported by annotations. Spring will create a dynamic proxy the interface which will use the `FunctionService` to create an `Execution`, invoke the execution and coerce the results to a defined return type, if necessary. This technique is very similar to the way Spring Data repositories work, thus some of the configuration and concepts should be familiar. Generally a single

interface definition maps to multiple function executions, one corresponding to each method defined in the interface.

Annotations for Function Execution

To support client side function execution, the following annotations are provided: `@OnRegion`, `@OnServer`, `@OnServers`, `@OnMember`, `@OnMembers`. These correspond to the Execution implementations GemFire's `FunctionService` provides. Each annotation exposes the appropriate attributes. These annotations also provide an optional `resultCollector` attribute whose value is the name of a Spring bean implementing [ResultCollector](#) to use for the execution.



Note

The proxy interface binds all declared methods to the same execution configuration. Although it is expected that single method interfaces will be common, all methods in the interface are backed by the same proxy instance and therefore all share the same configuration.

Here are some examples:

```
@OnRegion(region="someRegion", resultCollector="myCollector")
public interface FunctionExecution {
    @FunctionId("function1")
    public String doIt(String s1, int i2);
    public String getString(Object arg1, @Filter Set<Object> keys) ;
}
```

By default, the function id is the simple (unqualified) method name. `@FunctionId` is used to bind this invocation to a different function id.

Enabling Annotation Processing

The client side uses Spring's component scanning capability to discover annotated interfaces. To enable function execution annotation processing, you can use XML:

```
<gfe-data:function-executions base-package="org.example.myapp.functions" />
```

Note that the `function-executions` tag is provided in the `gfe-data` namespace. The `base-package` attribute is required to avoid scanning the entire class path. Additional filters are provided as described in the Spring [reference](#).

Or annotate your Java configuration class:

```
@EnableGemfireFunctionExecutions(basePackages = "org.example.myapp.functions")
```

10.5 Programmatic Function Execution

Using the annotated interface as described in the previous section, simply wire your interface into a bean that will invoke the function:

```
@Component
public class MyApp {

    @Autowired FunctionExecution functionExecution;

    public void doSomething() {
        functionExecution.doIt("hello", 123);
    }

}
```

Alternately, you can use a Function Execution template directly. For example `GemfireOnRegionFunctionTemplate` creates an onRegion execution. For example:

```
Set<?,?> myFilter = getFilter();
Region<?,?> myRegion = getRegion();
GemfireOnRegionOperations template = new GemfireOnRegionFunctionTemplate(myRegion);
String result = template.executeAndExtract("someFunction", myFilter, "hello", "world", 1234);
```

Internally, function executions always return a List. `executeAndExtract` assumes a singleton list containing the result and will attempt to coerce that value into the requested type. There is also an `execute` method that returns the List itself. The first parameter is the function id. The filter argument is optional. The following arguments are a variable argument list.

11. Sample Applications



Note

Sample applications are now maintained in the [Spring Data GemFire Examples](#) repository.

The Spring Data GemFire project also includes one sample application. Named "Hello World", the sample demonstrates how to configure and use GemFire inside a Spring application. At runtime, the sample offers a *shell* to the user allowing him to run various commands against the grid. It provides an excellent starting point for users unfamiliar with the essential components or the Spring and GemFire concepts.

The sample is bundled with the distribution and is Maven-based. One can easily import them into any Maven-aware IDE (such as SpringSource [Tool Suite](#)) or run them from the command-line.

11.1 Hello World

The Hello World sample demonstrates the core functionality of the Spring GemFire project. It bootstraps GemFire, configures it, executes arbitrary commands against it and shuts it down when the application exits. Multiple instances can be started at the same time as they will work with each other sharing data without any user intervention.



Running under Linux

If you experience networking problems when starting GemFire or the samples, try adding the following system property `java.net.preferIPv4Stack=true` to the command line (insert `-Djava.net.preferIPv4Stack=true`). For an alternative (global) fix especially on Ubuntu see this [link](#)

Starting and stopping the sample

Hello World is designed as a stand-alone java application. It features a `Main` class which can be started either from your IDE of choice (in Eclipse/STS through `Run As/Java Application`) or from the command line through Maven using `mvn exec:java`. One can also use `java` directly on the resulting artifact if the classpath is properly set.

To stop the sample, simply type `exit` at the command line or press `Ctrl+C` to stop the VM and shutdown the Spring container.

Using the sample

Once started, the sample will create a shared data grid and allow the user to issue commands against it. The output will likely look as follows:

```
INFO: Created GemFire Cache [Spring GemFire World] v. X.Y.Z
INFO: Created new cache region [myWorld]
INFO: Member xxxxxx:50694/51611 connecting to region [myWorld]
Hello World!
Want to interact with the world ? ...
Supported commands are:

get <key> - retrieves an entry (by key) from the grid
put <key> <value> - puts a new entry into the grid
remove <key> - removes an entry (by key) from the grid
...
```

For example to add new items to the grid one can use:

```
-> put 1 unu
INFO: Added [1=unu] to the cache
null
-> put 1 one
INFO: Updated [1] from [unu] to [one]
unu
-> size
1
-> put 2 two
INFO: Added [2=two] to the cache
null
-> size
2
```

Multiple instances can be created at the same time. Once started, the new VMs automatically see the existing region and its information:

```
INFO: Connected to Distributed System ['Spring GemFire World'=xxxx:56218/49320@yyyyy]
Hello World!
...

-> size
2
-> map
[2=two] [1=one]
-> query length = 3
[one, two]
```

Experiment with the example, start (and stop) as many instances as you want, run various commands in one instance and see how the others react. To preserve data, at least one instance needs to be alive all times - if all instances are shutdown, the grid data is completely destroyed (in this example - to preserve data between runs, see the GemFire documentations).

Hello World Sample Explained

Hello World uses both Spring XML and annotations for its configuration. The initial bootstrapping configuration is `app-context.xml` which includes the cache configuration, defined under `cache-context.xml` file and performs classpath [scanning](#) for Spring [components](#). The cache configuration defines the GemFire cache, region and for illustrative purposes a simple cache listener that acts as a logger.

The main *beans* are `HelloWorld` and `CommandProcessor` which rely on the `GemfireTemplate` to interact with the distributed fabric. Both classes use annotations to define their dependency and life-cycle callbacks.

Part III. Other Resources

In addition to this reference documentation, there are a number of other resources that may help you learn how to use GemFire and Spring framework. These additional, third-party resources are enumerated in this section.

12. Useful Links

- [Spring Data GemFire Home Page](#)
- [vFabric GemFire Home Page](#)
- [vFabric GemFire Documentation](#)
- [GemFire Community Home Page](#)
- [Spring Data GemFire Forum](#)

Part IV. Appendices

Appendix A. Spring Data GemFire Schema

Spring Data GemFire Core Schema (gfe)

```

<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<xsd:schema xmlns="http://www.springframework.org/schema/gemfire"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:beans="http://www.springframework.org/
schema/beans"
  xmlns:tool="http://www.springframework.org/schema/tool" xmlns:context="http://
www.springframework.org/schema/context"
  xmlns:repository="http://www.springframework.org/schema/data/repository"
  targetNamespace="http://www.springframework.org/schema/gemfire"
  elementFormDefault="qualified" attributeFormDefault="unqualified"
  version="1.3">
  <xsd:import namespace="http://www.springframework.org/schema/beans" />
  <xsd:import namespace="http://www.springframework.org/schema/tool" />
  <xsd:import namespace="http://www.springframework.org/schema/context" />
  <!-- -->
  <xsd:annotation>
    <xsd:documentation><![CDATA[
      Namespace support for the Spring GemFire project.
    ]]></xsd:documentation>
  </xsd:annotation>
  <!-- -->
  <xsd:complexType name="cacheBaseType">
    <xsd:sequence>
      <xsd:element name="transaction-listener" type="beanDeclarationType"
        minOccurs="0" maxOccurs="unbounded">
        <xsd:annotation>
          <xsd:documentation><![CDATA[
Registers a bean as a TransactionListener with the CacheTransactionManager. The bean must
implement com.gemstone.gemfire.cache.TransactionListener
and may be nested or referenced.
          ]]></xsd:documentation>
        </xsd:annotation>
      </xsd:element>
      <xsd:element name="transaction-writer" type="beanDeclarationType"
        minOccurs="0" maxOccurs="1">
        <xsd:annotation>
          <xsd:documentation><![CDATA[
Registers a bean as a TransactionWriter with the CacheTransactionManager. The bean must
implement com.gemstone.gemfire.cache.TransactionWriter
and may be nested or referenced.
          ]]></xsd:documentation>
        </xsd:annotation>
      </xsd:element>
      <xsd:element name="gateway-conflict-resolver" minOccurs="0"
        maxOccurs="1">
        <xsd:annotation>
          <xsd:documentation
            source="com.gemstone.gemfire.cache.util.GatewayConflictResolver"><![CDATA[
A gateway conflict resolver for this cache. A gateway conflict resolver handles conflicts
in the case of concurrent updates using a WAN gateway. The bean
must implement com.gemstone.gemfire.cache.util.GatewayConflictResolver. Requires Gemfire
version 7.0 or higher.
          ]]></xsd:documentation>
        </xsd:annotation>
        <xsd:appinfo>
          <tool:annotation>
            <tool:exports
              type="com.gemstone.gemfire.cache.util.GatewayConflictResolver" />
            </tool:annotation>
          </xsd:appinfo>
        </xsd:annotation>
      </xsd:complexType>
      <xsd:sequence>
        <xsd:any namespace="##other" processContents="skip"
          minOccurs="0" maxOccurs="unbounded">
          <xsd:annotation>
            <xsd:documentation><![CDATA[
Inner bean definition of the gateway conflict resolver.
          ]]></xsd:documentation>
          </xsd:annotation>
        </xsd:any>
      </xsd:sequence>
    </xsd:sequence>
  </xsd:complexType>

```

Spring Data GemFire Data Access Schema (gfe-data)

```

<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<xsd:schema xmlns="http://www.springframework.org/schema/data/gemfire" xmlns:xsd="http://
www.w3.org/2001/XMLSchema" xmlns:beans="http://www.springframework.org/schema/beans"
xmlns:tool="http://www.springframework.org/schema/tool"
xmlns:repository="http://www.springframework.org/schema/data/repository"
xmlns:context="http://www.springframework.org/schema/context"
xmlns:gfe="http://www.springframework.org/schema/gemfire"
targetNamespace="http://www.springframework.org/schema/data/
gemfire" elementFormDefault="qualified" attributeFormDefault="unqualified" version="1.3">
  <xsd:import namespace="http://www.springframework.org/schema/beans"/>
  <xsd:import namespace="http://www.springframework.org/schema/tool"/>
  <xsd:import namespace="http://www.springframework.org/schema/data/repository"
    schemaLocation="http://www.springframework.org/schema/data/repository/spring-
repository.xsd"/>
  <xsd:import namespace="http://www.springframework.org/schema/gemfire"
    schemaLocation="http://www.springframework.org/schema/gemfire/spring-gemfire.xsd"/>
  <xsd:import namespace="http://www.springframework.org/schema/context"
    schemaLocation="http://www.springframework.org/schema/context/spring-context.xsd" />
  <!-- -->
  <xsd:annotation>
    <xsd:documentation><![CDATA[
      Namespace support for the Spring Data GemFire Client side data access.
    ]]></xsd:documentation>
  </xsd:annotation>
  <!-- -->
  <!-- Repositories -->
  <xsd:element name="repositories">
    <xsd:complexType>
      <xsd:complexContent>
        <xsd:extension base="repository:repositories">
          <xsd:attributeGroup ref="gemfire-repository-attributes"/>
          <xsd:attributeGroup ref="repository:repository-attributes"/>
        </xsd:extension>
      </xsd:complexContent>
    </xsd:complexType>
  </xsd:element>
  <!-- -->
  <xsd:element name="function-executions">
    <xsd:annotation>
      <xsd:documentation><![CDATA[
        Enables component scanning for annotated function execution interfaces.
      ]]></xsd:documentation>
    </xsd:annotation>
    <xsd:complexType>
      <xsd:sequence>
        <xsd:element name="include-
filter" type="context:filterType" minOccurs="0" maxOccurs="unbounded">
          <xsd:annotation>
            <xsd:documentation><![CDATA[
              Controls which eligible types to include for component scanning.
            ]]></xsd:documentation>
          </xsd:annotation>
        </xsd:element>
        <xsd:element name="exclude-
filter" type="context:filterType" minOccurs="0" maxOccurs="unbounded">
          <xsd:annotation>
            <xsd:documentation><![CDATA[
              Controls which eligible types to exclude for component scanning.
            ]]></xsd:documentation>
          </xsd:annotation>
        </xsd:element>
      </xsd:sequence>
    </xsd:complexType>
    <xsd:attribute name="base-package" type="xsd:string" use="required">
      <xsd:annotation>
        <xsd:documentation><![CDATA[
          Defines the base package where function execution interfaces will be tried to be
          detected.
        ]]></xsd:documentation>
      </xsd:annotation>
    </xsd:attribute>
  </xsd:element>

```