

Upgrading to {brandname} 9.3

The {brandname} community

Table of Contents

1. Upgrading from 9.0 to 9.1	2
1.1. Kubernetes Ping changes	2
1.2. Stat Changes	2
1.3. (FineGrained)AtomicMap reimplemented	2
1.4. RemoteCache keySet/entrySet/values	2
1.5. DeltaAware deprecated	2
1.6. {brandname} Query Configuration	3
1.7. Store Batch Size Changes	3
1.8. Partition Handling changes	3
2. Upgrading from 8.x to 9.0	4
2.1. Default transaction mode changed	4
2.2. Removed eagerLocking and eagerLockingSingleNode configuration settings	4
2.3. Removed async transaction support	4
2.4. Deprecated all the dummy related transaction classes.	4
2.5. Clustering configuration changes	5
2.6. Default Cache changes	5
2.7. Marshalling Enhancements and Store Compatibility	5
2.8. New Cloud module for library mode	5
2.9. Entry Retriever is now removed	5
2.10. Map / Reduce is now removed	5
2.11. Spring 3 support is now removed	5
2.12. Function classes have moved packages	6
2.13. SegmentCompletionListener interface has moved	6
2.14. Spring module dependency changes	6
2.15. Total order executor is now removed	7
2.16. HikariCP is now the default implementation for JDBC PooledConnectionFactory	7
2.17. RocksDB in place of LevelDB	7
2.18. JDBC Mixed and Binary stores removed	7
2.19. @Store Annotation Introduced	7
2.20. Server authentication changes	8
2.21. Package org.infinispan.util.concurrent.jdk8backported has been removed	8
2.21.1. Moved classes	8
2.21.2. Removed classes	8
2.22. Store as Binary is deprecated	8
2.23. DataContainer collection methods are deprecated	8
3. Upgrading from 8.1 to 8.2	9
3.1. Entry Retriever is deprecated	9
3.2. Map / Reduce is deprecated	9

4. Upgrading from 8.x to 8.1	10
4.1. Packaging changes	10
4.1.1. CDI module split	10
4.1.2. Spring module split	10
4.2. Spring 3 support is deprecated	10
5. Upgrading from 7.x to 8.0	11
5.1. Configuration changes	11
5.1.1. Removal of Async Marshalling	11
5.1.2. Reenabling of isolation level configurations in server	11
5.1.3. Subsystem renaming in server	11
5.1.4. Server domain mode	11
6. Upgrading from 6.0 to 7.0	12
6.1. API Changes	12
6.1.1. Cache Loader	12
6.1.2. Cache Writer	12
6.1.3. Filters	12
6.2. Declarative configuration	12
7. Upgrading from 5.3 to 6.0	13
7.1. Declarative configuration	13
7.2. Deprecated API removal	13
8. Upgrading from 5.2 to 5.3	14
8.1. Declarative configuration	14
9. Upgrading from 5.1 to 5.2	15
9.1. Declarative configuration	15
9.2. Transaction	15
9.3. Cache Loader and Store configuration	15
9.4. Virtual Nodes and Segments	15
10. Upgrading from 5.0 to 5.1	16
10.1. API	16
10.2. Eviction and Expiration	16
10.3. Transactions	17
10.4. State transfer	17
10.5. Configuration	17
10.6. Flags and ClassLoaders	19
10.7. JGroups Bind Address	19

This guide walks you through the process of upgrading {brandname}.

Chapter 1. Upgrading from 9.0 to 9.1

1.1. Kubernetes Ping changes

The latest version of Kubernetes Ping uses unified environmental variables for both Kubernetes and OpenShift. Some of them were shortened for example `OPENSIFT_KUBE_PING_NAMESPACE` was changed to `KUBERNETES_NAMESPACE`. Please refer to [Kubernetes Ping documentation](#).

1.2. Stat Changes

Average values for read, write and removals are now returned in Nanoseconds, opposed to Milliseconds.

1.3. (FineGrained)AtomicMap reimplemented

`{brandname}` now contains a new implementation of both `AtomicMap` and `FineGrainedAtomicMap`, but the semantics has been preserved. The new implementation does not use `DeltaAware` interface but the [Functional API](#) instead.

There are no changes needed for `AtomicMap`, but it now supports non-transactional use case as well.

`FineGrainedAtomicMap` now uses the [Grouping API](#) and therefore you need to [enable groups in configuration](#). Also it holds entries as regular cache entries, plus one cache entry for cached key set (the map itself). Therefore the cache size or iteration/streaming results may differ. Note that fine grained atomic maps are still supported on transactional caches only.

1.4. RemoteCache keySet/entrySet/values

`RemoteCache` now implements all of the collection backed methods from `Map` interface. Previously `keySet` was implemented, however it was a deep copy. This has now changed and it is a backing set. That is that the set retrieves the updated values on each invocation or updates to the backing remote cache for writes. The `entrySet` and `values` methods are also now supported as backing variants as well.

If you wish to have a copy like was provided before it is recommended to copy the contents into a in memory local set such as

```
Set<K> keysCopy = remoteCache.keySet().stream().collect(Collectors.toSet());
```

1.5. DeltaAware deprecated

Interfaces `DeltaAware`, `Delta` and `CopyableDeltaAware` have been deprecated. Method `AdvancedCache.applyDelta()` has been deprecated and the implementation does not allow custom set of locked keys. `ApplyDeltaCommand` and its uses in interceptor stack are deprecated.

Any partial updates to an entry should be replaced using the [Functional API](#).

1.6. {brandname} Query Configuration

The configuration property `directory_provider` now accepts a new value `local-heap`. This value replaces the now deprecated `ram`, and as its predecessor will cause the index to be stored in a `org.apache.lucene.store.RAMDirectory`.

The configuration value `ram` is still accepted and will have the same effect, but failing to replace `ram` with `local-heap` will cause a warning to be logged. We suggest to perform this replacement, as the `ram` value will no longer be recognised by {brandname} in a future version.

This change was made as the team believes the `local-heap` name better expresses the storage model, especially as this storage method will not allow real-time replication of the index across multiple nodes. This index storage option is mostly useful for single node integration testing of the query functionality.

1.7. Store Batch Size Changes

`TableManipulation::batchSize` and `JpaStoreConfiguration::batchSize` have been deprecated and replaced by the higher level `AbstractStoreConfiguration::maxBatchSize`.

1.8. Partition Handling changes

In {brandname} 9.1 partition handling has been improved to allow for automatic [conflict resolution on partition merges](#). Consequently, `PartitionHandlingConfiguration::enabled` has been deprecated in favour of `PartitionHandlingConfiguration::whenSplit`. Configuring `whenSplit` to the `DENY_READ_WRITES` strategy is equivalent to setting `enabled` to `true`, whilst specifying `ALLOW_READ_WRITES` is equivalent to disabling partition handling (default).

Furthermore, during a partition merge with `ALLOW_READ_WRITES`, the default `EntryMergePolicy` is `MergePolicies.PREFERRED_ALWAYS` which provides a deterministic way of tie-breaking `CacheEntry` conflicts. If you require the old behaviour, simply set the merge-policy to `null`.

Chapter 2. Upgrading from 8.x to 9.0

2.1. Default transaction mode changed

The default configuration for transactional caches changed from `READ_COMMITTED` and `OPTIMISTIC` locking to `REPEATABLE_READ` and `OPTIMISTIC` locking with `write-skew` enabled.

Also, using the `REPEATABLE_READ` isolation level and `OPTIMISTIC` locking without `write-skew` enabled is no longer allowed. To help with the upgrade, `write-skew` will be automatically enabled in this case.

The following configuration has been deprecated:

- `write-skew`: as said, it is automatically enabled.
- `<versioning>` and its attributes. It is automatically enabled and configured when needed.

2.2. Removed eagerLocking and eagerLockingSingleNode configuration settings

Both were deprecated since version 5.1. `eagerLocking(true)` can be replaced with `lockingMode(LockingMode.PESSIMISTIC)`, and `eagerLockingSingleNode()` does not need a replacement because it was a no-op.

2.3. Removed async transaction support

Asynchronous mode is no longer supported in transactional caches and it will automatically use the synchronous cache mode. In addition, the second phase of a transaction commit is done synchronously. The following methods (and related) are deprecated:

- `TransactionConfigurationBuilder.syncCommitPhase(boolean)`
- `TransactionConfigurationBuilder.syncRollbackPhase(boolean)`

2.4. Deprecated all the dummy related transaction classes.

The following classes have been deprecated and they will be removed in the future:

- `DummyBaseTransactionManager`: replaced by `EmbeddedBasedTransactionManager`;
- `DummyNoXaXid` and `DummyXid`: replaced by `EmbeddedXid`;
- `DummyTransaction`: replaced by `EmbeddedTransaction`;
- `DummyTransactionManager`: replaced by `EmbeddedTransactionManager`;
- `DummyTransactionManagerLookup` and `RecoveryDummyTransactionManagerLookup`: replaced by `EmbeddedTransactionManagerLookup`;
- `DummyUserTransaction`: replaced by `EmbeddedUserTransaction`;

2.5. Clustering configuration changes

The `mode` attribute in the XML declaration of clustered caches is no longer mandatory. It defaults to SYNC.

2.6. Default Cache changes

Up to {brandname} 8.x, the default cache always implicitly existed, even if not declared in the XML configuration. Additionally, the default cache configuration affected all other cache configurations, acting as some kind of base template. Since 9.0, the default cache only exists if it has been explicitly configured. Additionally, even if it has been specified, it will never act as base template for other caches.

2.7. Marshalling Enhancements and Store Compatibility

Internally {brandname} 9.x has introduced many improvements to its marshalling codebase in order to improve performance and allow for greater flexibility. Consequently, data marshalled and persisted by {brandname} 8.x is no longer compatible with {brandname} 9.0. To aid you in migrating your existing stores to 9.0, we have provided a [Store Migrator](#), however at present this only allows the migration of JDBC stores.

2.8. New Cloud module for library mode

In {brandname} 8.x, cloud related configuration were added to `infinispan-core` module. Since 9.0 they were moved to `infinispan-cloud` module.

2.9. Entry Retriever is now removed

The entry retriever feature has been removed. Please update to use the new Streams feature detailed in the User Guide. The `org.infinispan.filter.CacheFilters` class can be used to convert `KeyValueFilter` and `Converter` instances into proper Stream operations that are able to be marshalled.

2.10. Map / Reduce is now removed

Map reduce has been removed in favor of the new Streams feature which should provide more features and performance. There are no bridge classes to convert to the new streams and all references must be rewritten.

2.11. Spring 3 support is now removed

Spring 3 is no longer supported.

2.12. Function classes have moved packages

The class `SerializableSupplier` has moved from the `org.infinispan.stream` package to the `org.infinispan.util.function` package.

The class `CloseableSupplier` has moved from the `org.infinispan.util` package to the `org.infinispan.util.function` package.

The classes `TriConsumer`, `CloseableSupplier`, `SerializableRunnable`, `SerializableFunction` & `SerializableCallable` have all been moved from the `org.infinispan.util` package to the `org.infinispan.util.function` package.

2.13. SegmentCompletionListener interface has moved

The interface `SegmentCompletionListener` has moved from the interface `org.infinispan.CacheStream` to the new `org.infinispan.BaseCacheStream`.

2.14. Spring module dependency changes

All {brandname}, Spring and Logger dependencies are now in the `provided` scope. One can decide whether to use small jars or uber jars but they need to be added to the classpath of the application. It also gives one freedom in choosing Spring (or Spring Boot) version.

Here is an example:

```
<dependencies>
  <dependency>
    <groupId>org.infinispan</groupId>
    <artifactId>infinispan-embedded</artifactId>
  </dependency>
  <dependency>
    <groupId>org.infinispan</groupId>
    <artifactId>infinispan-spring4-embedded</artifactId>
  </dependency>
  <dependency>
    <groupId>org.springframework</groupId>
    <artifactId>spring-context</artifactId>
  </dependency>
  <dependency>
    <groupId>org.springframework.session</groupId>
    <artifactId>spring-session</artifactId>
  </dependency>
</dependencies>
```

Additionally there is no Logger implementation specified (since this may vary depending on use case).

2.15. Total order executor is now removed

The total order protocol now uses the `remote-command-executor`. The attribute `total-order-executor` in `<container>` tag is removed.

2.16. HikariCP is now the default implementation for JDBC PooledConnectionFactory

`HikariCP` offers superior performance to `c3p0` and is now the default implementation. Additional properties for HikariCP can be provided by placing a `hikari.properties` file on the classpath or by specifying the path to the file via `PooledConnectionFactoryConfiguration.propertyFile` or `properties-file` in the connection pool's xml config. N.B. a properties file specified explicitly in the configuration is loaded instead of the `hikari.properties` file on the class path and Connection pool characteristics which are explicitly set in `PooledConnectionFactoryConfiguration` always override the values loaded from a properties file.

Support for `c3p0` has been deprecated and will be removed in a future release. Users can force `c3p0` to be utilised as before by providing the system property `-Dinfinispan.jdbc.c3p0.force=true`.

2.17. RocksDB in place of LevelDB

The LevelDB cache store was replaced with a `RocksDB`. RocksDB is a fork of LevelDB which provides superior performance in high concurrency scenarios. The new cache store can parse old LevelDB configurations but will always use the RocksDB implementation.

2.18. JDBC Mixed and Binary stores removed

The JDBC Mixed and Binary stores have been removed due to the poor performance associated with storing entries in buckets. Storing entries in buckets is non-optimal as each read/write to the store requires an existing bucket for a given hash to be retrieved, deserialised, updated, serialised and then re-inserted back into the db. If you were previously using one of the removed stores, we have provided a migrator tool to assist in migrating data from an existing binary table to a JDBC string based store. See [JDBC Migrator Guide](#) for more details.

2.19. @Store Annotation Introduced

A new annotation, `@Store`, has been added for persistence stores. This allows a store's properties to be explicitly defined and validated against the provided store configuration. Existing stores should be updated to use this annotation and the store's configuration class should also declare the `@ConfigurationFor` annotation. If neither of these annotations are present on the store or configuration class, then a your store will continue to function as before, albeit with a warning that additional store validation cannot be completed.

2.20. Server authentication changes

The no-anonymous policy is now automatically enabled for Hot Rod authentication unless explicitly specified.

2.21. Package

org.infinispan.util.concurrent.jdk8backported has been removed

2.21.1. Moved classes

Classes regarding EntrySizeCalculator have now been moved down to the org.infinispan.util package.

2.21.2. Removed classes

The *ConcurrentHashMapV8 classes and their supporting classes have all been removed. The CollectionFactory#makeBoundedConcurrentMap method should be used if you desire to have a bounded ConcurrentMap.

2.22. Store as Binary is deprecated

Store as Binary configuration is now deprecated and will be removed in a future release. This is replaced by the new memory configuration.

2.23. DataContainer collection methods are deprecated

The keySet, entrySet and values methods on DataContainer have been deprecated. These behavior of these methods are very inconsistent and will be removed later. It is recommended to update references to use iterator or iteratorIncludingExpired methods instead.

Chapter 3. Upgrading from 8.1 to 8.2

3.1. Entry Retriever is deprecated

Entry Retriever is now deprecated and will be removed in {brandname} 9. This is replaced by the new Streams feature.

3.2. Map / Reduce is deprecated

Map reduce is now deprecated and will be removed in {brandname} 9. This is replaced by the new Streams feature.

Chapter 4. Upgrading from 8.x to 8.1

4.1. Packaging changes

4.1.1. CDI module split

CDI module (GroupId:ArtifactId `org.infinispan:infinispan-cdi`) has been split into `org.infinispan:infinispan-cdi-embedded` and `org.infinispan:infinispan-cdi-remote`. Please make sure that you use proper artifact.

4.1.2. Spring module split

Spring module (GroupId:ArtifactId `org.infinispan:infinispan-spring4`) has been split into `org.infinispan:infinispan-spring4-embedded` and `org.infinispan:infinispan-spring4-remote`. Please make sure that you use proper artifact.

4.2. Spring 3 support is deprecated

Spring 3 support (GroupId:ArtifactId `org.infinispan:infinispan-spring`) is deprecated. Please consider migrating into Spring 4 support.

Chapter 5. Upgrading from 7.x to 8.0

5.1. Configuration changes

5.1.1. Removal of Async Marshalling

Async marshalling has been entirely dropped since it was never reliable enough. The "async-marshalling" attribute has been removed from the 8.0 XML schema and will be ignored when parsing 7.x configuration files. The programmatic configuration methods related to `asyncMarshalling/syncMarshalling` are now deprecated and have no effect aside from producing a WARN message in the logs.

5.1.2. Reenabling of isolation level configurations in server

Because of the inability to configure write skew in the server, the isolation level attribute was ignored and defaulted to `READ_COMMITTED`. Now, when enabling `REPEATABLE_READ` together with optimistic locking, write skew is enabled by default in local and synchronous configurations.

5.1.3. Subsystem renaming in server

In order to avoid conflict and confusion with the similar subsystems in WildFly, we have renamed the following subsystems in server: * `infinispan` → `datagrid-infinispan` * `jgroups` → `datagrid-jgroups` * `endpoint` → `datagrid-infinispan-endpoint`

5.1.4. Server domain mode

We no longer support the use of standalone mode for running clusters of servers. Domain mode (`bin/domain.sh`) should be used instead.

Chapter 6. Upgrading from 6.0 to 7.0

6.1. API Changes

6.1.1. Cache Loader

To be more inline with JCache and `java.util.collections` interfaces we have changed the first argument type for the `CacheLoader.load` & `CacheLoader.contains` methods to be `Object` from type `K`.

6.1.2. Cache Writer

To be more inline with JCache and `java.util.collections` interfaces we have changed the first argument type for the `CacheWriter.delete` method to be `Object` from type `K`.

6.1.3. Filters

Over time {brandname} added 2 interfaces with identical names and almost identical methods. The `org.infinispan.notifications.KeyFilter` and `org.infinispan.persistence.spi.AdvancedCacheLoader$KeyFilter` interfaces.

Both of these interfaces are used for the sole purpose of filtering an entry by it's given key. {brandname} 7.0 has also introduced the `KeyValueFilter` which is similar to both but also can filter on the entries value and/or metadata.

As such all of these classes have been moved into a new package `org.infinispan.filter` and all of their related helper classes.

The new `org.infinispan.filter.KeyFilter` interface has replaced both of the previous interfaces and all previous references use the new interface.

6.2. Declarative configuration

The XML schema for the embedded configuration has changed to more closely follow the server configuration. Use the `config-converter.sh` or `config-converter.bat` scripts to convert an {brandname} 6.0 to the current format.

Chapter 7. Upgrading from 5.3 to 6.0

7.1. Declarative configuration

In order to use all of the latest features, make sure you change the namespace declaration at the top of your XML configuration files as follows:

```
<infinispan xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation=
"urn:infinispan:config:6.0 http://www.infinispan.org/schemas/infinispan-config-
6.0.xsd" xmlns="urn:infinispan:config:6.0">
    ...
</infinispan>
```

7.2. Deprecated API removal

- Class `org.infinispan.persistence.remote.wrapperEntryWrapper`.
- Method `ObjectOutput startObjectOutput(OutputStream os, boolean isReentrant)` from class `org.infinispan.commons.marshall.StreamingMarshaller`.
- Method `CacheEntry getCacheEntry(Object key, EnumSet<Flag> explicitFlags, ClassLoader explicitClassLoader)` from class `org.infinispan.AdvancedCache`. Please use instead: `AdvanceCache.withFlags(Flag... flags).with(ClassLoader classLoader).getCacheEntry(K key)`.
- Method `AtomicMap<K, V> getAtomicMap(Cache<MK, ?> cache, MK key, FlagContainer flagContainer)` from class `org.infinispan.atomic.AtomicMapLookup`. Please use instead `AtomicMapLookup.getAtomicMap(cache.getAdvancedCache().withFlags(Flag... flags), MK key)`.
- Package `org.infinispan.config` (and all methods involving the old configuration classes). All methods removed has an overloaded method which receives the new configuration classes as parameters. Please refer to [\[configuration\]](#) for more information about the new configuration classes.



This only affects the programmatic configuration.

- Class `org.infinispan.context.FlagContainer`.
- Method `boolean isLocal(Object key)` from class `org.infinispan.distribution.DistributionManager`. Please use instead `DistributionManager.getLocality(Object key)`.
- JMX operation `void setStatisticsEnabled(boolean enabled)` from class `org.infinispan.interceptors.TxInterceptor` Please use instead the `statisticsEnabled` attribute.
- Method `boolean delete(boolean synchronous)` from class `org.infinispan.io.GridFile`. Please use instead `GridFile.delete()`.
- JMX attribute `long getLocallyInterruptedTransactions()` from class `org.infinispan.util.concurrent.locks.DeadlockDetectingLockManager`.

Chapter 8. Upgrading from 5.2 to 5.3

8.1. Declarative configuration

In order to use all of the latest features, make sure you change the namespace declaration at the top of your XML configuration files as follows:

```
<infinispan xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation=
"urn:infinispan:config:5.2 http://www.infinispan.org/schemas/infinispan-config-
5.2.xsd" xmlns="urn:infinispan:config:5.3">
  ...
</infinispan>
```

Chapter 9. Upgrading from 5.1 to 5.2

9.1. Declarative configuration

In order to use all of the latest features, make sure you change the namespace declaration at the top of your XML configuration files as follows:

```
<infinispan xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation=
"urn:infinispan:config:5.2 http://www.infinispan.org/schemas/infinispan-config-
5.2.xsd" xmlns="urn:infinispan:config:5.2">
    ...
</infinispan>
```

9.2. Transaction

The default transaction enlistment model has changed ([ISPN-1284](#)) from `XAResource` to `Synchronization`. Also now, if the `XAResource` enlistment is used, then `recovery` is enabled by default.

In practical terms, if you were using the default values, this should not cause any backward compatibility issues but an increase in performance of about 5-7%. However in order to use the old configuration defaults, you need to configure the following:

```
<transaction useSynchronization="false">
    <recovery enabled="false"/>
</transaction>
```

or the programmatic configuration equivalent:

```
ConfigurationBuilder builder = new ConfigurationBuilder();
builder.transaction().useSynchronization(false).recovery().enabled(false)
```

9.3. Cache Loader and Store configuration

Cache Loader and Store configuration has changed greatly in {brandname} 5.2. Please refer to the [Cache Loaders and Stores](#) documentation.

9.4. Virtual Nodes and Segments

The concept of Virtual Nodes doesn't exist anymore in {brandname} 5.2 and has been replaced by Segments. Please refer to the [Clustering modes](#) documentation for details.

Chapter 10. Upgrading from 5.0 to 5.1

10.1. API

The cache and cache manager hierarchies have changed slightly in 5.1 with the introduction of `BasicCache` and `BasicCacheContainer`, which are parent classes of existing `Cache` and `CacheContainer` classes respectively. What's important is that Hot Rod clients must now code against `BasicCache` and `BasicCacheContainer` rather than `Cache` and `CacheContainer`. So previous code that was written like this will no longer compile.

WontCompile.java

```
import org.infinispan.Cache;
import org.infinispan.manager.CacheContainer;
import org.infinispan.client.hotrod.RemoteCacheManager;
...
CacheContainer cacheContainer = new RemoteCacheManager();
Cache cache = cacheContainer.getCache();
```

Instead, if Hot Rod clients want to continue using interfaces higher up the hierarchy from the remote cache/container classes, they'll have to write:

Correct.java

```
import org.infinispan.BasicCache;
import org.infinispan.manager.BasicCacheContainer;
import org.infinispan.client.hotrod.RemoteCacheManager;
...
BasicCacheContainer cacheContainer = new RemoteCacheManager();
BasicCache cache = cacheContainer.getCache();
```

However, previous code that interacted against the `RemoteCache` and `RemoteCacheManager` will work as it used to:

AlsoCorrect.java

```
import org.infinispan.client.hotrod.RemoteCache;
import org.infinispan.client.hotrod.RemoteCacheManager;
...
RemoteCacheManager cacheContainer = new RemoteCacheManager();
RemoteCache cache = cacheContainer.getCache();
```

10.2. Eviction and Expiration

- The eviction XML element no longer defines the `wakeupInterval` attribute. This is now configured via the `expiration` element:

```
<expiration wakeUpInterval="60000"... />
```

Eviction's `maxEntries` is used as guide for the entire cache, but eviction happens on a per cache segment, so when the segment is full, the segment is evicted. That's why `maxEntries` is a theoretical limit but in practical terms, it'll be a bit less than that. This is done for performance reasons.

10.3. Transactions

- A cache marked as `TRANSACTIONAL` cannot be accessed outside of a transaction, and a `NON_TRANSACTIONAL` cache cannot be accessed within a transaction. In 5.0, a transactional cache would support non-transactional calls as well. This change was done to be in-line with expectations set out in [JSR-107](#) as well as to provide more consistent behavior.
- In 5.0, commit and rollback phases were asynchronous by default. Starting with 5.1, these are now synchronous by default, to provide the guarantees required by a single lock-owner model.

10.4. State transfer

One of the big changes we made in 5.1 was to use the same push-based state transfer we introduced in 5.0 both for rehashing in distributed mode and for state retrieval in replicated mode. We even borrow the consistent hash concept in replicated mode to transfer state from all previous cache members at once in order to speed up transfer.

As a consequence we've unified the state transfer configuration as well, there is now a `stateTransfer` element containing a simplified state transfer configuration. The corresponding attributes in the `stateRetrieval` and `hash` elements have been deprecated, as have been some attributes that are no longer used.

10.5. Configuration

If you use XML to configure {brandname}, you shouldn't notice any change, except a much faster startup, courtesy of the [StAX](#) based parser. However, if you use programmatic configuration, read on for the important differences.

Configuration is now packaged in `org.infinispan.configuration`, and you must use a fluent, builder style:

```
Configuration c1 = new ConfigurationBuilder()
    // Adjust any configuration defaults you want
    .clustering()
        .l1()
            .disable()
        .mode(DIST_SYNC)
        .hash()
            .numOwners(5)
    .build();
```

- The old javabean style configuration is now deprecated and will be removed in a later version.
- Configuration properties which can be safely changed at runtime are mutable, and all others are immutable.
- To copy a configuration, use the `read()` method on the builder, for example:

```
Configuration c2 = new ConfigurationBuilder()
    // Read in C1 to provide defaults
    .read(c1)
    .clustering()
        .l1()
        .enable()
    // This cache is DIST_SYNC, will have 5 owners, with L1 cache enabled
    .build();
```

This completely replaces the old system of defining a set of overrides on bean properties. Note that this means the behaviour of {brandname} configuration is somewhat different when used programmatically. Whilst before, you could define a default configuration, and any overrides would be applied on top of *your* defaults when defined, now you must explicitly read in your defaults to the builder. This allows for much greater flexibility in your code (you can have as many "default" configurations as you want), and makes your code more explicit and type safe (finding references works).

The schema is unchanged from before. {brandname} 4.0 configurations are currently not being parsed. To upgrade, just change the schema definition from:

```
<infinispan
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:infinispan:config:4.1
http://www.infinispan.org/schemas/infinispan-config-4.1.xsd"
  xmlns="urn:infinispan:config:4.1">
```

to

```
<infinispan
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:infinispan:config:5.1
http://www.infinispan.org/schemas/infinispan-config-5.1.xsd"
  xmlns="urn:infinispan:config:5.1">
```

The schema documentation has changed format, as it is now produced using the standard tool `xsd doc`. This should be a significant improvement, as better navigation is offered. Some elements and attributes are missing docs right now, we are working on adding this. As an added benefit, your IDE should now show documentation when an xsd referenced (as above)

We are in the process of adding in support for this configuration style for modules (such as cache stores). In the meantime, please use the old configuration or XML if you require support for cache

store module configuration.

10.6. Flags and ClassLoaders

The `Flags` and `ClassLoader` API has changed. In the past, the following would work:

```
cache.withFlags(f1, f2); cache.withClassLoader(cl); cache.put(k, v);
```

In 5.1.0, these `withX()` methods return a new instance and not the cache itself, so thread locals are avoided and the code above will not work. If used in a fluent manner however, things still work:

```
cache.withFlags(f1, f2).withClassLoader(cl).put(k, v);
```

The above pattern has always been the intention of this API anyway.

10.7. JGroups Bind Address

Since upgrading to JGroups 3.x, `-Dbind.address` is ignored. This should be replaced with `-Djgroups.bind_addr`.