

Upgrading Infinispan 11.0

Table of Contents

1. Upgrading from 10.1 to 11.0.....	1
1.1. Query.....	1
2. Upgrading from 10.0 to 10.1 and 10.0 to 11.0.....	2
2.1. Maximum Idle Timeouts with Clustered Cache Modes.....	2
3. Upgrading from 10.0 to 10.1.....	3
3.1. REST Store.....	3
3.2. Infinispan Lucene Directory is deprecated.....	3
3.3. Security role mappers and audit loggers.....	3
3.4. Memcached protocol server is deprecated.....	4
3.5. Hot Rod client default mechanism changed to SCRAM-SHA-512.....	4
4. Upgrading from 9.4 to 10.0.....	5
4.1. Marshalling.....	5
4.1.1. Externalizer Deprecations.....	5
4.1.2. Store Migration.....	5
4.1.3. Store Defaults.....	5
4.2. CacheContainerAdmin.....	6
4.3. Hot Rod 3.0.....	6
4.4. Total Order transaction protocol is deprecated.....	6
4.5. Removed the <code>infinispan.server.hotrod.workerThreads</code> system property.....	6
4.6. Removed AtomicMap and FineGrainedAtomicMap.....	6
4.7. Removed Delta and DeltaAware.....	6
4.8. Removed compatibility mode.....	6
4.9. Removed the implicit default cache.....	7
4.10. Removed DistributedExecutor.....	7
4.11. Removed the Tree module.....	7
4.12. The JDBC PooledConnectionFactory now utilises Agroal.....	7
4.13. XML configuration changes.....	7
4.14. RemoteCache Changes.....	7
4.14.1. Marshalling Changes.....	7
4.14.2. The <code>getBulk</code> methods have been removed.....	8
4.15. Persistence changes.....	8
4.16. Client/Server changes.....	8
4.17. <code>SKIP_LISTENER_NOTIFICATION</code> flag.....	8
4.18. <code>performAsync</code> header removed from REST.....	8
4.19. REST status code change.....	8
4.20. Default JGroups stacks in the XML configuration.....	9
4.21. JGroups <code>S3_PING</code> replaced with <code>NATIVE_S3_PING</code>	9
4.22. Cache and Cache Manager Listeners can now be configured to be non blocking.....	9

4.23. Distributed Streams operations no longer support null values	9
4.24. Removed the infinispn-cloud module	9
4.25. Removed experimental flag GUARANTEED_DELIVERY	9
4.26. Cache Health	10
4.27. Multi-tenancy	10
4.28. OffHeap Automatic Resizing	10
4.29. Deprecated methods from DataContainer removed	10
5. Upgrading from 9.3 to 9.4	11
5.1. Client/Server changes	11
5.1.1. Compatibility mode deprecation	11
5.1.2. Memcached storage	11
5.1.3. Scripts Response	11
5.1.4. WebSocket endpoint removal	11
5.1.5. Hot Rod client connection pool properties	11
5.1.6. Server thread pools	11
5.2. Persistence Changes	12
5.2.1. Shared and Passivation	12
5.3. Query changes	12
5.3.1. AffinityIndexManager	12
6. Upgrading from 9.2 to 9.3	13
6.1. AdvancedCacheLoader changes	13
6.2. Partition Handling Configuration	13
6.3. Stat Changes	13
6.4. Event log changes	13
6.5. Max Idle Expiration Changes	13
6.6. WildFly Modules	13
6.7. Deserialization Whitelist	14
7. Upgrading from 9.0 to 9.1	15
7.1. Kubernetes Ping changes	15
7.2. Stat Changes	15
7.3. (FineGrained)AtomicMap reimplemented	15
7.4. RemoteCache keySet/entrySet/values	15
7.5. DeltaAware deprecated	15
7.6. Infinispan Query Configuration	16
7.7. Store Batch Size Changes	16
7.8. Partition Handling changes	16
8. Upgrading from 8.x to 9.0	17
8.1. Default transaction mode changed	17
8.2. Removed eagerLocking and eagerLockingSingleNode configuration settings	17
8.3. Removed async transaction support	17
8.4. Deprecated all the dummy related transaction classes	17

8.5. Clustering configuration changes	18
8.6. Default Cache changes	18
8.7. Marshalling Enhancements and Store Compatibility	18
8.8. New Cloud module for library mode	18
8.9. Entry Retriever is now removed	18
8.10. Map / Reduce is now removed	18
8.11. Spring 4 support is now removed	18
8.12. Function classes have moved packages	19
8.13. SegmentCompletionListener interface has moved	19
8.14. Spring module dependency changes	19
8.15. Total order executor is now removed	20
8.16. HikariCP is now the default implementation for JDBC PooledConnectionFactory	20
8.17. RocksDB in place of LevelDB	20
8.18. JDBC Mixed and Binary stores removed	20
8.19. @Store Annotation Introduced	20
8.20. Server authentication changes	21
8.21. Package org.infinispan.util.concurrent.jdk8backported has been removed	21
8.21.1. Moved classes	21
8.21.2. Removed classes	21
8.22. Store as Binary is deprecated	21
8.23. DataContainer collection methods are deprecated	21
9. Upgrading from 8.1 to 8.2	22
9.1. Entry Retriever is deprecated	22
9.2. Map / Reduce is deprecated	22
10. Upgrading from 8.x to 8.1	23
10.1. Packaging changes	23
10.1.1. CDI module split	23
10.1.2. Spring module split	23
10.2. Spring 3 support is deprecated	23
11. Upgrading from 7.x to 8.0	24
11.1. Configuration changes	24
11.1.1. Removal of Async Marshalling	24
11.1.2. Reenabling of isolation level configurations in server	24
11.1.3. Subsystem renaming in server	24
11.1.4. Server domain mode	24
12. Upgrading from 6.0 to 7.0	25
12.1. API Changes	25
12.1.1. Cache Loader	25
12.1.2. Cache Writer	25
12.1.3. Filters	25
12.2. Declarative configuration	25

13. Upgrading from 5.3 to 6.0.....	26
13.1. Declarative configuration.....	26
13.2. Deprecated API removal.....	26
14. Upgrading from 5.2 to 5.3.....	27
14.1. Declarative configuration.....	27
15. Upgrading from 5.1 to 5.2.....	28
15.1. Declarative configuration.....	28
15.2. Transaction	28
15.3. Cache Loader and Store configuration	28
15.4. Virtual Nodes and Segments	28
16. Upgrading from 5.0 to 5.1.....	29
16.1. API	29
16.2. Eviction and Expiration	29
16.3. Transactions	30
16.4. State transfer.....	30
16.5. Configuration	30
16.6. Flags and ClassLoaders.....	32
16.7. JGroups Bind Address	32
17. Rolling Upgrades and Updates with Kubernetes and OpenShift	33
17.1. Performing Rolling Updates on Kubernetes	33
17.2. Performing Rolling Upgrades on Kubernetes	35

Chapter 1. Upgrading from 10.1 to 11.0

1.1. Query

The AffinityIndexManager is removed. Indexes should be configured with the `near-real-time` index manager instead.

Chapter 2. Upgrading from 10.0 to 10.1 and 10.0 to 11.0

2.1. Maximum Idle Timeouts with Clustered Cache Modes

Maximum idle expiration has been changed to improve data consistency with clustered cache modes when Infinispan nodes fail.



- `Cache.get()` calls do not return until the touch commands complete. This synchronous behavior increases latency of client requests and reduces performance.
- Maximum idle expiration, `max-idle`, does not currently work with entries stored in off-heap memory.
- Likewise, `max-idle` does not work if caches use cache stores as a persistence layer.

See [Maximum Idle Expiration](#) for complete details.

Chapter 3. Upgrading from 10.0 to 10.1

3.1. REST Store

The following configurations were removed from the REST store: `append-cache-name-to-path` and `path`.

To specify the remote server endpoint path, a single configuration `cache-name` should be used.

3.2. Infinispan Lucene Directory is deprecated

The Infinispan Lucene directory is now deprecated and will be removed in a future release. Consequently, the Infinispan Directory provider for Hibernate Search will also be discontinued, with no replacement.

Both `IndexManagers` that rely on the Lucene Directory are also deprecated, the `InfinispanIndexManager` and the `AffinityIndexManager`. Users are encouraged to reconfigure their indexes as non-shared, using the Near Real Time `IndexManager`, with file system storage:

```
<distributed-cache name="default">
  <indexing index="PRIMARY_OWNER">
    <property name="default.indexmanager">near-real-time</property>
    <property name="default.indexBase">
/opt/infinispan/server/data/indexes</property>
  </indexing>
</distributed-cache>
```

Queries need to be adjusted to use the `BROADCAST` runtime option.

3.3. Security role mappers and audit loggers

The security role mapper implementations have been moved from the `org.infinispan.security.impl` package to the `org.infinispan.security.mappers` package:

- `org.infinispan.security.impl.CommonNameRoleMapper` ⇒ `org.infinispan.security.mappers.CommonNameRoleMapper`
- `org.infinispan.security.impl.ClusterRoleMapper` ⇒ `org.infinispan.security.mappers.ClusterRoleMapper`
- `org.infinispan.security.impl.IdentityRoleMapper` ⇒ `org.infinispan.security.mappers.IdentityRoleMapper`

The security audit logger implementations have been moved from the `org.infinispan.security.impl` package to the `org.infinispan.security.audit` package:

- `org.infinispan.security.impl.LoggingAuditLogger` ⇒ `org.infinispan.security.audit.LoggingAuditLogger`
- `org.infinispan.security.impl.NullAuditLogger` ⇒ `org.infinispan.security.audit.NullAuditLogger`

3.4. Memcached protocol server is deprecated

The Memcached protocol server is now deprecated and will be removed in a future release. This is being done because Infinispan only implements the very dated text-only protocol instead of the binary protocol which means no security (authentication / encryption), no support for some new Memcached features and no integration with Infinispan features like single-port. If someone in the community wishes to implement the binary protocol, we would revert the decision.

3.5. Hot Rod client default mechanism changed to SCRAM-SHA-512

The default Hot Rod client authentication mechanism has been changed from `DIGEST-MD5` to `SCRAM-SHA-512`. If you are using `property` user realms, you must make sure you are using `plain-text` storage.

Chapter 4. Upgrading from 9.4 to 10.0

4.1. Marshalling

The internal marshalling capabilities of Infinispan have undergone a significant refactoring in 10.0. The marshalling of internal Infinispan objects and user objects are now truly isolated. This means that it's now possible to configure [Marshaller](#) implementations in embedded mode or on the server, without having to handle the marshalling of Infinispan internal classes. Consequently, it's possible to easily change the marshaller implementation, in a similar manner to how users of the HotRod client are accustomed.

As a consequence of the above changes, the default marshaller used for marshalling user types is no longer based upon JBoss Marshalling. Instead we now utilise the ProtoStream library to store user types in the language agnostic [Protocol Buffers](#) format. It is still possible to utilise the old default, [JBossUserMarshaller](#), however it's necessary to add the `org.infinispan:infinispan-jboss-marshalling` artifact to your application's classpath.

4.1.1. Externalizer Deprecations

The following interfaces/annotations have been deprecated as a consequence of the marshalling refactoring:

- [Externalizer](#),
- [AdvancedExternalizer](#)
- [SerializeWith](#)

For cluster communication any configured [Externalizer](#)'s are still utilised to marshall objects, however they are ignored when persisting data to cache stores unless the [JBossUserMarshaller](#) is explicitly configured via the global [SerializationConfiguration](#).

It's highly recommended to migrate from the old Externalizer and JBoss marshalling approach to the new ProtoStream based marshalling, as the interfaces listed above and the [JBossUserMarshaller](#) implementation will be removed in future versions.

4.1.2. Store Migration

Unfortunately, the extensive marshalling changes mean that the binary format used by Infinispan stores in `9.4.x` is no longer compatible with `10.0.x`. Therefore, it's necessary for any existing stores to be migrated to the new format via the StoreMigrator tool.



Whilst we regret that 9.4.x stores are no longer binary compatible, these extensive changes should ensure binary compatibility across future major versions.

4.1.3. Store Defaults

Stores now default to being segmented if the property is not configured. Some stores do not support being segmented, which will result in a configuration exception being thrown at startup. The

moving forward position is to use segmented stores when possible to increase cache wide performance and reduce memory requirements for various operations including state transfer.

The file based stores (SingleFileStore and SoftIndexFileStore) both support being segmented, but their current implementation requires opening file descriptors based on how many segments there are. This may cause issues in some configurations and users should be aware. Infinispan will print a single WARN message when such a configuration is found.

4.2. CacheContainerAdmin

Caches created through the CacheContainerAdmin API will now be **PERMANENT** by default. Use the **VOLATILE** flag to obtain the previous behaviour.

4.3. Hot Rod 3.0

Older versions of the Hot Rod protocol treated expiration values greater than the number of milliseconds in 30 days as Unix time. Starting with Hot Rod 3.0 this adjustment no longer happens and expiration is taken literally.

4.4. Total Order transaction protocol is deprecated

Total Order transaction protocol is going to be removed in a future release. Use the default protocol (2PC).

4.5. Removed the `infinispan.server.hotrod.workerThreads` system property

The `infinispan.server.hotrod.workerThreads` property was introduced as a hack to work around the fact that the configuration did not expose it. The property has been removed and endpoint worker threads must now be exclusively configured using the `worker-threads` attribute.

4.6. Removed AtomicMap and FineGrainedAtomicMap

AtomicMapLookup, AtomicMap and FineGrainedAtomicMap have been removed. Please see FunctionalMaps or Cache#Merge for similar functionality.

4.7. Removed Delta and DeltaAware

The previously deprecated Delta and DeltaAware interfaces have been removed.

4.8. Removed compatibility mode

The previously deprecated Compatibility Mode has been removed.

4.9. Removed the implicit default cache

The default cache must now be named explicitly via the [GlobalConfigurationBuilder#defaultCacheName\(\)](#) method.

4.10. Removed DistributedExecutor

The previously deprecated DistributedExecutor is now removed. References should be updated to use ClusterExecutor.

4.11. Removed the Tree module

TreeCache has been unsupported for a long time and was only intended as a quick stopgap for JBossCache users. The module has now been removed completely.

4.12. The JDBC PooledConnectionFactory now utilises Agroal

Previously the JDBC PooledConnectionFactory provided c3p0 and HikariCP based connection pools. From 10.0 we only provide a PooledConnectionFactory based upon the [Agroal project](#). This means that it is no longer possible to utilise `c3p0.properties` and `hikari.properties` files to configure the pool, instead an agroal compatible properties file can be provided.

4.13. XML configuration changes

Several configuration elements and attributes that were deprecated since 9.0 have been removed:

- `<eviction>` - replaced with `memory`
- `<versioning>` - automatically enabled
- `<data-container>` - no longer customizable
- `deadlock-detection-spin` - always disabled
- `write-skew` - enabled automatically

The xsite state transfer chunk size (`<backup><state-transfer chunk-size="X"/></backup>`) can no longer be `>= 0`, same as the regular state transfer chunk size. Previously a value `<= 0` would transfer the entire cache in a single batch, which is almost always a bad idea.

4.14. RemoteCache Changes

4.14.1. Marshalling Changes

The default marshaller is no longer GenericJbossMarshaller. We now utilise the [ProtoStream](#) library as the default. If Java Serialization is required by clients, we strongly recommend utilising the link: [JavaSerializationMarshaller](#) instead. However if the GenericJbossMarshaller must be used,

it's necessary to add the `org.infinispan:infinispan-jboss-marshalling` artifact to your client's classpath and for the `GenericJbossMarshaller` to be configured as the marshaller.

4.14.2. The `getBulk` methods have been removed

The `getBulk` method is an expensive method as it requires holding all keys in memory at once and requires a possibly very single result to populate it. The new `retrieveEntries`, `entrySet`, `keySet` and `values` methods handle this in a much more efficient way. Therefore the `getBulk` methods have been removed in favor of them.

4.15. Persistence changes

- File-based cache stores (`SingleFileStore`, `SoftIndexFileStore`, `RocksDBStore`) filesystem layout has been normalized so that they will use the `GlobalStateConfiguration` persistent location as a default location. Additionally, all stores will now use the cache name as part of the data file/directory naming allowing multiple stores to avoid conflicts and ambiguity.
- The CLI loader (`infinispan-persistence-cli`) has been removed.
- The LevelDB store (`infinispan-cachestore-leveldb`) has been removed. Use the RocksDB store instead, as it is fully backwards compatible.
- The deprecated `singleton` store configuration option and the wrapper class `SingletonCacheWriter` have been removed.

Using `shared=true` is enough, as only the primary owner of each key will write to a shared store.

4.16. Client/Server changes

- The Hot Rod client and server only support protocol versions 2.0 and higher. Support for Hot Rod versions 1.0 to 1.3 has been dropped.

4.17. `SKIP_LISTENER_NOTIFICATION` flag

`SKIP_LISTENER_NOTIFICATION` notification flag has been added in the hotrod client. This flag only works when the client and the server version is 9.4.15 or higher. Spring Session integration uses this flag when a session id has changed. If you are using Spring Session with Infinispan 9.4, consider upgrading the client and the server.

4.18. `performAsync` header removed from REST

The `performAsync` header was removed from the REST server. Clients that want to perform async operations with the REST server should manage the request and response on their side to avoid blocking.

4.19. REST status code change

REST operations that don't return resources and are used with `PUT`, `POST` and `DELETE` methods now

return status **204** (No content) instead of **200**.

4.20. Default JGroups stacks in the XML configuration

With the introduction of inline XML JGroups stacks in the configuration, two default stacks are always enabled: **udp** and **tcp**. If you are declaring your own stacks with the same names, an exception reporting the conflict will be thrown. Simply rename your own configurations to avoid the conflict.

4.21. JGroups S3_PING replaced with NATIVE_S3_PING

Because of changes in AWS's access policy regarding signatures, S3_PING will not work in newer regions and will stop working in older regions too. For this reason, you should migrate to using NATIVE_S3_PING instead.

4.22. Cache and Cache Manager Listeners can now be configured to be non blocking

Listeners in the past that were sync, always ran in the thread that caused the event. We now allow a Listener method to be non-blocking in that it will still fire in the original thread, under the assumption that it will return immediately. Please read the Listener Javadoc for information and examples on this.

4.23. Distributed Streams operations no longer support null values

Distributed Streams has parts rewritten to utilize non blocking reactive streams based operations. As such null values are not supported as values from operations as per the reactive streams spec. Please utilize other means to denote a null value.

4.24. Removed the infinispn-cloud module

The infinispn-cloud module has been removed and the **kubernetes**, **ec2**, **google** and **azure** default configurations have been included in **infinispn-core** and can be referenced as default named JGroups configurations.

4.25. Removed experimental flag GUARANTEED_DELIVERY

Almost as soon as GUARANTEED_DELIVERY was added, UNICAST3 and NAKACK2.resend_last_seqno removed the need for it. It was always documented as experimental, so we removed it without deprecation and we also removed the RSVP protocol from the default JGroups stacks.

4.26. Cache Health

The possible statuses of the cache health are now HEALTHY, HEALTHY_REBALANCING and DEGRADED to better reflect the fact that `rebalancing` doesn't mean a cluster is unhealthy.

4.27. Multi-tenancy

When using multi-tenancy in the WildFly based server, it's necessary to specify the `content-path` for each of the REST connectors, to match the `prefix` element under `multi-tenancy\rest\prefix`.

4.28. OffHeap Automatic Resizing

Off Heap memory containers now will dynamically resize based on number of entries in the container. Due to this the address count configuration value is now deprecated for APIs and has been removed from the xml parser.

4.29. Deprecated methods from DataContainer removed

The deprecated methods `keySet`, `values`, `entrySet` and `executeTask` has been removed.

Chapter 5. Upgrading from 9.3 to 9.4

5.1. Client/Server changes

5.1.1. Compatibility mode deprecation

Compatibility mode has been deprecated and will be removed in the next Infinispan version.

To use a cache from multiple endpoints, it is recommended to store data in binary format and to configure the MediaType for keys and values.

If storing data as unmarshalled objects is still desired, the equivalent of compatibility mode is to configure keys and values to store object content:

```
<encoding>
  <key media-type="application/x-java-object"/>
  <value media-type="application/x-java-object"/>
</encoding>
```

5.1.2. Memcached storage

For better interoperability between endpoints, the Memcached server no longer stores keys as `java.lang.String`, but as UTF-8 `byte[]`.

If using memcached, it's recommended to run a rolling upgrade from 9.3 to store data in the new format, or reload the data in the cache.

5.1.3. Scripts Response

Distributed scripts with text-based data type no longer return `null` when the result from each server is null. The response is now a JSON array with each individual result, e.g. `"[null, null]"`

5.1.4. WebSocket endpoint removal

The WebSocket endpoint has been unmaintained for several years. It has been removed.

5.1.5. Hot Rod client connection pool properties

Since the Hot Rod client was overhauled in 9.2, the way the connection pool configuration is handled has changed. Infinispan 9.4 introduces a new naming scheme for the connection pool properties which deprecates the old *commons-pool* names. For a complete reference of the available configuration options for the properties file please refer to [remote client configuration](#) javadoc.

5.1.6. Server thread pools

The threads that handle the child Netty event loops have been renamed from `*-ServerWorker` to `*-`

5.2. Persistence Changes

5.2.1. Shared and Passivation

A store cannot be configured as both shared and having passivation enabled. Doing so can cause data inconsistencies as there is no way to synchronize data between all the various nodes. As such this configuration will now cause a startup exception. Please update your configuration as appropriate.

5.3. Query changes

5.3.1. AffinityIndexManager

The default number of shards is down to 4, it was previously equals to the number of segments in the cache.

Chapter 6. Upgrading from 9.2 to 9.3

6.1. AdvancedCacheLoader changes

The AdvancedCacheLoader SPI has been enhanced to provide an alternative method to process and instead allows reactive streams based `publishKeys` and `publishEntries` methods which provide benefits in performance, threading and ease of use. Note this change will only affect you if you wish take advantage of it in any custom CacheLoaders you may have implemented.

6.2. Partition Handling Configuration

In 9.3 the default `MergePolicy` is now `MergePolicy.NONE`, opposed to `MergePolicy.PREFERRED_ALWAYS`.

6.3. Stat Changes

We have reverted the stat changes introduced in 9.1, so average values for read, write and removals are once again returned as milliseconds.

6.4. Event log changes

Several new event log messages have been added, and one message has been removed (ISPN100013).

6.5. Max Idle Expiration Changes

The max idle entry expiration information is sent between owners in the cluster. However when an entry expires via max idle on a given node, this was not replicated (only removing it locally). Max idle has been enhanced to now expire an entry across the entire cluster, instead of per node. This includes ensuring that max idle expiration is applied across all owners (meaning if another node has accessed the entry within the given time it will prevent that entry from expiring on other nodes that didn't have an access).

Max idle in a transactional clustered cache does not remove expired entries on access (although it will not be returned). These entries are only removed via the expiration reaper.

Iteration in a clustered cache will still show entries that are expired via `maxIdle` to ensure good performance, but could be removed at any point due to expiration reaper.

6.6. WildFly Modules

The Infinispan WildFly modules are now located in the `system/add-ons/{moduleprefix}` dir as per the [WildFly module conventions](#).

6.7. Deserialization Whitelist

Deserialization of content sent by clients to the server are no longer allowed by default. This applies to JSON, XML, and marshalled byte[] that, depending on the cache configuration, will cause the server to convert it to Java Objects either to store it or to perform any operation that cannot be done on a byte[] directly.

The deserialization needs to be enabled using system properties, either by class name or regular expressions:

```
// Comma separated list of fully qualified class names  
-Dinfinispan.deserialization.whitelist.classes=java.time.Instant,com.myclass.Entity  
  
// Regex expression  
-Dinfinispan.deserialization.whitelist.regexp=.*
```

Chapter 7. Upgrading from 9.0 to 9.1

7.1. 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](#).

7.2. Stat Changes

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

7.3. (FineGrained)AtomicMap reimplemented

Infinispan 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.

7.4. 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());
```

7.5. 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.

7.6. Infinispan 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 Infinispan 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.

7.7. Store Batch Size Changes

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

7.8. Partition Handling changes

In Infinispan 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 8. Upgrading from 8.x to 9.0

8.1. 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.

8.2. 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.

8.3. 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)`

8.4. 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`;

8.5. Clustering configuration changes

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

8.6. Default Cache changes

Up to Infinispan 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.

8.7. Marshalling Enhancements and Store Compatibility

Internally Infinispan 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 Infinispan 8.x is no longer compatible with Infinispan 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.

8.8. New Cloud module for library mode

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

8.9. 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.

8.10. 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.

8.11. Spring 4 support is now removed

Spring 4 is no longer supported.

8.12. 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.

8.13. SegmentCompletionListener interface has moved

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

8.14. Spring module dependency changes

All Infinispan, 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-spring5-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).

8.15. 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.

8.16. 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`.

8.17. 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.

8.18. 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.

8.19. @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.

8.20. Server authentication changes

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

8.21. Package

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

8.21.1. Moved classes

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

8.21.2. 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.

8.22. 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.

8.23. 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 9. Upgrading from 8.1 to 8.2

9.1. Entry Retriever is deprecated

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

9.2. Map / Reduce is deprecated

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

Chapter 10. Upgrading from 8.x to 8.1

10.1. Packaging changes

10.1.1. 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.

10.1.2. Spring module split

Spring module (GroupId:ArtifactId `org.infinispan:infinispan-spring5`) has been split into `org.infinispan:infinispan-spring5-embedded` and `org.infinispan:infinispan-spring5-remote`. Please make sure that you use proper artifact.

10.2. Spring 3 support is deprecated

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

Chapter 11. Upgrading from 7.x to 8.0

11.1. Configuration changes

11.1.1. 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.

11.1.2. 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.

11.1.3. 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`

11.1.4. 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 12. Upgrading from 6.0 to 7.0

12.1. API Changes

12.1.1. 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`.

12.1.2. 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`.

12.1.3. Filters

Over time Infinispan 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. Infinispan 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.

12.2. 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 Infinispan 6.0 to the current format.

Chapter 13. Upgrading from 5.3 to 6.0

13.1. 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 https://infinispan.org/schemas/infinispan-config-6.0.xsd"
xmlns="urn:infinispan:config:6.0">
...
</infinispan>
```

13.2. 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.



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 14. Upgrading from 5.2 to 5.3

14.1. 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 https://infinispan.org/schemas/infinispan-config-5.2.xsd"
xmlns="urn:infinispan:config:5.3">
  ...
</infinispan>
```


Chapter 15. Upgrading from 5.1 to 5.2

15.1. 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 https://infinispan.org/schemas/infinispan-config-5.2.xsd"
xmlns="urn:infinispan:config:5.2">
    ...
</infinispan>
```

15.2. 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)
```

15.3. Cache Loader and Store configuration

Cache Loader and Store configuration has changed greatly in Infinispan 5.2.

15.4. Virtual Nodes and Segments

The concept of Virtual Nodes doesn't exist anymore in Infinispan 5.2 and has been replaced by Segments.

Chapter 16. Upgrading from 5.0 to 5.1

16.1. 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();
```

16.2. 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.

16.3. 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.

16.4. 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.

16.5. Configuration

If you use XML to configure Infinispan, 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 Infinispan 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. Infinispan 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
https://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
https://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.

16.6. 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.

16.7. JGroups Bind Address

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

Chapter 17. Rolling Upgrades and Updates with Kubernetes and OpenShift

Pods running in Kubernetes and OpenShift are immutable. The only way you can alter the configuration is to roll out a new deployment.

Upgrades and updates sound similar but are distinct processes for rolling out new deployments.

17.1. Performing Rolling Updates on Kubernetes

Rolling updates replace existing pods with new ones.

- [Rolling Updates](#)
- [When to Use a Rolling Deployment](#)

Example DeploymentConfiguration for Rolling Updates

```
- apiVersion: v1
  kind: DeploymentConfig
  metadata:
    name: infinispn-cluster
  spec:
    replicas: 3
    strategy:
      type: Rolling
      rollingParams:
        updatePeriodSeconds: 10
        intervalSeconds: 20
        timeoutSeconds: 600
        maxUnavailable: 1
        maxSurge: 1
    template:
      spec:
        containers:
          - args:
              - -Djboss.default.jgroups.stack=kubernetes
            image: jboss/infinispn-server:latest
            name: infinispn-server
            ports:
              - containerPort: 8181
                protocol: TCP
              - containerPort: 9990
                protocol: TCP
              - containerPort: 11211
                protocol: TCP
              - containerPort: 11222
                protocol: TCP
              - containerPort: 57600
```

```

    protocol: TCP
  - containerPort: 7600
    protocol: TCP
  - containerPort: 8080
    protocol: TCP
env:
  - name: KUBERNETES_NAMESPACE
    valueFrom: {fieldRef: {apiVersion: v1, fieldPath: metadata.namespace}}
terminationMessagePath: /dev/termination-log
terminationGracePeriodSeconds: 90
livenessProbe:
  exec:
    command:
      - /usr/local/bin/is_running.sh
  initialDelaySeconds: 10
  timeoutSeconds: 80
  periodSeconds: 60
  successThreshold: 1
  failureThreshold: 5
readinessProbe:
  exec:
    command:
      - /usr/local/bin/is_healthy.sh
  initialDelaySeconds: 10
  timeoutSeconds: 40
  periodSeconds: 30
  successThreshold: 2
  failureThreshold: 5

```



Kubernetes uses very similar concept to Deployment Configurations called **Deployment**.

It is also highly recommended to adjust the JGroups stack to discover new nodes (or leaves) more quickly. One should at least adjust the value of **FD_ALL** timeout and adjust it to the longest GC Pause.

Other hints for tuning configuration parameters are:

- OpenShift should replace running nodes one by one. This can be achieved by adjusting **rollingParams** (**maxUnavailable: 1** and **maxSurge: 1**).
- Depending on the cluster size, one needs to adjust **updatePeriodSeconds** and **intervalSeconds**. The bigger cluster size is, the bigger those values should be used.
- When using Initial State Transfer, the **initialDelaySeconds** value for both probes should be set to higher value.
- During Initial State Transfer nodes might not respond to probes. The best results are achieved with higher values of **failureThreshold** and **successThreshold** values.

17.2. Performing Rolling Upgrades on Kubernetes

Rolling upgrades migrate data from one Infinispan cluster to another.

For both Kubernetes and OpenShift, the rolling upgrade procedure is nearly identical to the procedure for Infinispan server rolling upgrades.

Differences in rolling upgrade procedures

- Depending on configuration, it is a good practice to use [OpenShift Routes](#) or [Kubernetes Ingress API](#) to expose services to the clients. During the upgrade the Route (or Ingress) used by the clients can be altered to point to the new cluster.
- Invoking CLI commands can be done by using Kubernetes (`kubectl exec`) or OpenShift clients (`oc exec`). Here is an example: `oc exec <POD_NAME> — '/opt/jboss/infinispan-server/bin/ispn-cli.sh' '-c' '--controller=$(hostname -i):9990' '/subsystem=datagrid-infinispan/cache-container=clustered/distributed-cache=default:disconnect-source(migrator-name=hotrod)'`

Key differences when upgrading using the library mode:

- Client application needs to expose JMX. It usually depends on application and environment type but the easiest way to do it is to add the following switches into the Java bootstrap script `-Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=<PORT>`.
- Connecting to the JMX can be done by forwarding ports. With OpenShift this might be achieved by using `oc port-forward` command whereas in Kubernetes by `kubectl port-forward`.

The last step in the Rolling Upgrade (removing a Remote Cache Store) needs to be performed differently. We need to use [Kubernetes/OpenShift Rolling update](#) command and replace Pods configuration with the one which does not contain Remote Cache Store.

A detailed instruction might be found in [ISPN-6673](#) ticket.