

Deploying and Configuring Infinispan

10.1 Servers

Table of Contents

1. Getting Started with Infinispan Server	2
1.1. Infinispan Server Requirements	2
1.2. Downloading Server Distributions	2
1.3. Installing Infinispan Server	2
1.4. Running Infinispan Servers	3
1.4.1. Starting Infinispan Servers	3
1.4.2. Verifying Infinispan Cluster Discovery	3
1.4.3. Performing Operations with the Infinispan CLI	4
2. Setting Up Infinispan Endpoints	8
2.1. Infinispan Endpoints	8
2.1.1. Hot Rod	8
2.1.2. REST	8
2.1.3. Memcached	9
2.1.4. Protocol Comparison	9
2.2. Endpoint Connectors	9
2.2.1. Hot Rod Connectors	10
2.2.2. REST Connectors	10
2.2.3. Memcached Connectors	11
3. Monitoring Infinispan Servers	12
3.1. Working with Infinispan Server Logs	12
3.1.1. Infinispan Log Files	12
3.1.2. Configuring Infinispan Log Properties	12
3.1.3. Logging Framework	15
3.1.4. Access Logs	15
3.2. Retrieving Server Health Statistics	16
3.2.1. Accessing the Health API via JMX	16
3.2.2. Accessing the Health API via REST	17
3.3. Monitoring Infinispan Servers Over JMX	18
3.3.1. JMX Monitoring	18
4. Securing Infinispan Servers	23
4.1. Security Concepts	23
4.1.1. Authorization	23
4.1.2. Server Realms	23
4.2. Hot Rod Authentication	24
4.2.1. SASL Quality of Protection	25
4.2.2. SASL Policies	25
4.2.3. Using GSSAPI/Kerberos	26
4.3. Hot Rod and REST encryption (TLS/SSL)	27

5. Single Port	31
5.1. Single-Port router	31
5.1.1. Testing the Single-Port router	31
5.2. Hot Rod Protocol Detection	32
5.2.1. TLS/ALPN protocol selection	32
6. Remotely Executing Server-Side Tasks	34
6.1. Creating Server Tasks	34
6.1.1. Server Tasks	34
6.1.2. Deploying Server Tasks to Infinispan Servers	35
6.2. Creating Server Scripts	36
6.2.1. Server Scripts	36
6.2.2. Adding Scripts to Infinispan Servers	38
6.2.3. Programmatically Creating Scripts	39
6.3. Running Server-Side Tasks and Scripts	39
6.3.1. Running Tasks and Scripts	39
6.3.2. Programmatically Running Scripts	39
6.3.3. Programmatically Running Tasks	40
7. Adding Custom Marshaller Bridges	41
7.1. Protostuff	41
7.2. Kryo	41
7.3. Storing deserialized objects	42
7.4. Registering Custom Schemas/Serializers	42
8. Performing Rolling Upgrades	43
8.1. Rolling Upgrades	43
8.2. Setting Up Target Clusters	43
8.3. Synchronizing Data from Source Clusters	44

Infinispan server is a managed, distributed, and clusterable data grid that provides elastic scaling and high performance access to caches from multiple endpoints, such as Hot Rod and REST.

Chapter 1. Getting Started with Infinispan Server

Quickly set up Infinispan server and learn the basics.

1.1. Infinispan Server Requirements

Check host system requirements for the Infinispan server.

Infinispan server requires a Java Virtual Machine and supports:

- Java 8
- Java 11

1.2. Downloading Server Distributions

The Infinispan server distribution is an archive of Java libraries (**JAR** files), configuration files, and a **data** directory.

Procedure

Download the Infinispan 10.1 server from [Infinispan downloads](#).

Verification

Use the checksum to verify the integrity of your download.

1. Run the **sha1sum** command with the server download archive as the argument, for example:

```
$ sha1sum infinispan-server- $\{\text{version}\}$ .zip
```

2. Compare with the **SHA-1** checksum value on the Infinispan downloads page.

Reference

[Infinispan Server README](#) describes the contents of the server distribution.

1.3. Installing Infinispan Server

Extract the Infinispan server archive to any directory on your host.

Procedure

Use any extraction tool with the server archive, for example:

```
$ unzip infinispan-server- $\{\text{version}\}$ .zip
```

The resulting directory is your **\$ISPN_HOME**.

1.4. Running Infinispan Servers

Spin up Infinispan server instances that automatically form clusters. Learn how to create cache definitions to store your data.

1.4.1. Starting Infinispan Servers

Launch Infinispan server with the startup script.

Procedure

1. Open a terminal in `$ISPN_HOME`.
2. Run the `server` script.

Linux

```
$ bin/server.sh
```

Microsoft Windows

```
bin\server.bat
```

The server gives you these messages when it starts:

```
INFO [org.infinispan.SERVER] (main) ISPN080004: Protocol SINGLE_PORT listening
on 127.0.0.1:11222
INFO [org.infinispan.SERVER] (main) ISPN080001: Infinispan Server ${version}
started in 7453ms
```

Hello Infinispan!

- Open `127.0.0.1:11222` in any browser to see the Infinispan server welcome message.

Reference

[Infinispan Server README](#) describes command line arguments for the `server` script.

1.4.2. Verifying Infinispan Cluster Discovery

Infinispan servers running on the same network discover each other with the `MPING` protocol.

This procedure shows you how to use Infinispan server command arguments to start two instances on the same host and verify that the cluster view forms.

Prerequisites

Start a Infinispan server.

Procedure

1. Install and run a new Infinispan server instance.
 - a. Open a terminal in `$ISPN_HOME`.
 - b. Copy the root directory to `server2`.

```
$ cp -r server server2
```

2. Specify a port offset and the location of the `server2` root directory.

```
$ bin/server.sh -o 100 -s server2
```

Verification

Running servers return the following messages when new servers join clusters:

```
INFO [org.infinispan.CLUSTER] (jgroups-11,<server_hostname>)
ISPN000094: Received new cluster view for channel cluster:
[<server_hostname>|3] (2) [<server_hostname>, <server2_hostname>]
INFO [org.infinispan.CLUSTER] (jgroups-11,<server_hostname>)
ISPN100000: Node <server2_hostname> joined the cluster
```

Servers return the following messages when they join clusters:

```
INFO [org.infinispan.remoting.transport.jgroups.JGroupsTransport] (main)
ISPN000078: Starting JGroups channel cluster
INFO [org.infinispan.CLUSTER] (main)
ISPN000094: Received new cluster view for channel cluster:
[<server_hostname>|3] (2) [<server_hostname>, <server2_hostname>]
```

Reference

[Infinispan Server README](#) describes command line arguments for the `server` script.

1.4.3. Performing Operations with the Infinispan CLI

Connect to servers with the Infinispan command line interface (CLI) to access data and perform administrative functions.

Starting the Infinispan CLI

Start the Infinispan CLI as follows:

1. Open a terminal in `$ISPN_HOME`.
2. Run the CLI.

```
$ bin/cli.sh  
[disconnected]>
```

Connecting to Infinispan Servers

Do one of the following:

- Run the `connect` command to connect to a Infinispan server on the default port of `11222`:

```
[disconnected]> connect  
[hostname1@cluster//containers/default]>
```

- Specify the location of a Infinispan server. For example, connect to a local server that has a port offset of 100:

```
[disconnected]> connect 127.0.0.1:11322  
[hostname2@cluster//containers/default]>
```



Press the tab key to display available commands and options. Use the `-h` option to display help text.

Creating Caches from Templates

Use Infinispan cache templates to add caches with recommended default settings.

Procedure

1. Create a distributed, synchronous cache from a template and name it "mycache".

```
[//containers/default]> create cache --template=org.infinispan.DIST_SYNC mycache
```



Press the tab key after the `--template=` argument to list available cache templates.

2. Retrieve the cache configuration.


```
[//containers/default]> describe caches/mycache
{
  "distributed-cache" : {
    "mode" : "SYNC",
    "remote-timeout" : 17500,
    "state-transfer" : {
      "timeout" : 60000
    },
    "transaction" : {
      "mode" : "NONE"
    },
    "locking" : {
      "concurrency-level" : 1000,
      "acquire-timeout" : 15000,
      "striping" : false
    }
  }
}
```

Adding Cache Entries

Add data to caches with the Infinispan CLI.

Prerequisites

- Create a cache named "mycache".

Procedure

1. Add a key/value pair to **mycache**.

```
[//containers/default]> put --cache=mycache hello world
```



If the CLI is in the context of a cache, do **put k1 v1** for example:

```
[//containers/default]> cd caches/mycache
[//containers/default/caches/mycache]> put hello world
```

2. List keys in the cache.

```
[//containers/default]> ls caches/mycache
hello
```

3. Get the value for the **hello** key.
 - a. Navigate to the cache.

```
[//containers/default]> cd caches/mycache
```

- b. Use the `get` command to retrieve the key value.

```
[//containers/default/caches/mycache]> get hello  
world
```

Shutting Down Infinispan Servers

Use the CLI to gracefully shutdown running servers. This ensures that Infinispan passivates all entries to disk and persists state.

- Use the `shutdown server` command to stop individual servers, for example:

```
[//containers/default]> shutdown server server_hostname
```

- Use the `shutdown cluster` command to stop all servers joined to the cluster, for example:

```
[//containers/default]> shutdown cluster
```

Infinispan servers log the following shutdown messages:

```
INFO [org.infinispan.SERVER] (pool-3-thread-1) ISPN080002: Infinispan Server stopping  
INFO [org.infinispan.CONTAINER] (pool-3-thread-1) ISPN000029: Passivating all entries  
to disk  
INFO [org.infinispan.CONTAINER] (pool-3-thread-1) ISPN000030: Passivated 28 entries  
in 46 milliseconds  
INFO [org.infinispan.CLUSTER] (pool-3-thread-1) ISPN000080: Disconnecting JGroups  
channel cluster  
INFO [org.infinispan.CONTAINER] (pool-3-thread-1) ISPN000390: Persisted state,  
version=<Infinispan version> timestamp=YYYY-MM-DDTHH:MM:SS  
INFO [org.infinispan.SERVER] (pool-3-thread-1) ISPN080003: Infinispan Server stopped  
INFO [org.infinispan.SERVER] (Thread-0) ISPN080002: Infinispan Server stopping  
INFO [org.infinispan.SERVER] (Thread-0) ISPN080003: Infinispan Server stopped
```

When you shutdown Infinispan clusters, the shutdown messages include:

```
INFO [org.infinispan.SERVER] (pool-3-thread-1) ISPN080029: Cluster shutdown  
INFO [org.infinispan.CLUSTER] (pool-3-thread-1) ISPN000080: Disconnecting JGroups  
channel cluster
```

Chapter 2. Setting Up Infinispan Endpoints

Infinispan servers provide listener endpoints that handle requests from remote client applications.

2.1. Infinispan Endpoints

Infinispan endpoints expose the `CacheManager` interface over different connector protocols so you can remotely access data and perform operations to manage and maintain Infinispan clusters.

You can define multiple endpoint connectors on different socket bindings.

2.1.1. Hot Rod

Hot Rod is a binary TCP client-server protocol designed to provide faster data access and improved performance in comparison to text-based protocols.

Infinispan provides Hot Rod client libraries in Java, C++, C#, Node.js and other programming languages.

Topology state transfer

Infinispan uses topology caches to provide clients with cluster views. Topology caches contain entries that map internal JGroups transport addresses to exposed Hot Rod endpoints.

When client send requests, Infinispan servers compare the topology ID in request headers with the topology ID from the cache. Infinispan servers send new topology views if client have older topology IDs.

Cluster topology views allow Hot Rod clients to immediately detect when nodes join and leave, which enables dynamic load balancing and failover.

In distributed cache modes, the consistent hashing algorithm also makes it possible to route Hot Rod client requests directly to primary owners.

Reference

- [Infinispan Hot Rod Server](#)
- [Hot Rod client implementations](#)

2.1.2. REST

Infinispan exposes a RESTful interface that allows HTTP clients to access data, monitor and maintain clusters, and perform administrative operations.

You can use standard HTTP load balancers to provide clients with load balancing and failover capabilities. However, HTTP load balancers maintain static cluster views and require manual updates when cluster topology changes occur.

Reference

- [Infinispan REST Server](#)

- [mod_cluster HTTP load balancer](#)

2.1.3. Memcached

Infinispan provides an implementation of the Memcached text protocol for remote client access.

The Infinispan Memcached server supports clustering with replicated and distributed cache modes.

There are some Memcached client implementations, such as the Cache::Memcached Perl client, that can offer load balancing and failover detection capabilities with static lists of Infinispan server addresses that require manual updates when cluster topology changes occur.

Reference

- [Infinispan Memcached Server](#)
- [Memcached clients](#)
- [Memcached text protocol](#)

2.1.4. Protocol Comparison

	Hot Rod	HTTP / REST	Memcached
Topology-aware	Y	N	N
Hash-aware	Y	N	N
Encryption	Y	Y	N
Authentication	Y	Y	N
Conditional ops	Y	Y	Y
Bulk ops	Y	N	N
Transactions	N	N	N
Listeners	Y	N	N
Query	Y	Y	N
Execution	Y	N	N
Cross-site failover	Y	N	N

2.2. Endpoint Connectors

You configure Infinispan server endpoints with connector declarations that specify socket bindings, authentication mechanisms, and encryption configuration.

The default endpoint connector configuration is as follows:

```
<endpoints socket-binding="default">
  <hotrod-connector name="hotrod"/>
  <rest-connector name="rest"/>
  <memcached-connector socket-binding="memcached"/>
</endpoints>
```

- `endpoints` contains endpoint connector declarations and defines global configuration for endpoints such as default socket bindings, security realms, and whether clients must present valid TLS certificates.
- `<hotrod-connector name="hotrod"/>` declares a Hot Rod connector.
- `<rest-connector name="rest"/>` declares a Hot Rod connector.
- `<memcached-connector socket-binding="memcached"/>` declares a Memcached connector that uses the memcached socket binding.

Reference

[urn:infinispan:server](#) schema provides all available endpoint configuration.

2.2.1. Hot Rod Connectors

Hot Rod connector declarations enable Hot Rod servers.

```
<hotrod-connector name="hotrod">
  <topology-state-transfer />
  <authentication>
    ...
  </authentication>
  <encryption>
    ...
  </encryption>
</hotrod-connector>
```

- `name="hotrod"` logically names the Memcached connector.
- `topology-state-transfer` tunes the state transfer operations that provide Hot Rod clients with cluster topology.
- `authentication` configures SASL authentication mechanisms.
- `encryption` configures TLS settings for client connections.

Reference

[urn:infinispan:server](#) schema provides all available Hot Rod connector configuration.

2.2.2. REST Connectors

REST connector declarations enable REST servers.

```
<rest-connector name="rest">
  <authentication>
    ...
  </authentication>
  <cors-rules>
    ...
  </cors-rules>
  <encryption>
    ...
  </encryption>
</rest-connector>
```

- `name="rest"` logically names the REST connector.
- `authentication` configures authentication mechanisms.
- `cors-rules` specifies CORS (Cross Origin Resource Sharing) rules for cross-domain requests.
- `encryption` configures TLS settings for client connections.

Reference

[urn:infinispan:server](#) schema provides all available REST connector configuration.

2.2.3. Memcached Connectors

Memcached connector declarations enable Memcached servers.



Infinispan servers do not enable Memcached connectors by default.

```
<memcached-connector name="memcached" socket-binding="memcached" cache="mycache" />
```

- `name="memcached"` logically names the Memcached connector.
- `socket-binding="memcached"` declares a unique socket binding for the Memcached connector.
- `cache="mycache"` names the cache that the Memcached connector exposes. The default is `memcachedCache`.

Memcached connectors expose a single cache only. To expose multiple caches through the Memcached endpoint, you must declare additional connectors. Each Memcached connector must also have a unique socket binding.

Reference

[urn:infinispan:server](#) schema provides all available Memcached connector configuration.

Chapter 3. Monitoring Infinispan Servers

3.1. Working with Infinispan Server Logs

Infinispan uses JBoss LogManager to provide configurable logging mechanisms that capture details about the environment and record cache operations for troubleshooting purposes and root cause analysis.

3.1.1. Infinispan Log Files

Infinispan writes log messages to the following directory:

```
$ISPN_HOME/${infinispan.server.root}/log
```

server.log

Messages in human readable format, including boot logs that relate to the server startup. Infinispan creates this file by default when you launch servers.

server.log.json

Messages in JSON format that let you parse and analyze Infinispan logs. Infinispan creates this file when you enable the JSON-FILE log handler.

3.1.2. Configuring Infinispan Log Properties

You configure Infinispan logs with `logging.properties`, which is a standard Java properties file with the `property=value` format.

Procedure

1. Open `$ISPN_HOME/${infinispan.server.root}/conf/logging.properties` with any text editor.
2. Configure logging properties as appropriate.
3. Save and close `logging.properties`.

Log Levels

Log levels indicate the nature and severity of messages.

Log level	Description
TRACE	Provides detailed information about running state of applications. This is the most verbose log level.
DEBUG	Indicates the progress of individual requests or activities.
INFO	Indicates overall progress of applications, including lifecycle events.
WARN	Indicates circumstances that can lead to error or degrade performance.

Log level	Description
ERROR	Indicates error conditions that might prevent operations or activities from being successful but do not prevent applications from running.
FATAL	Indicates events that could cause critical service failure and application shutdown.

Infinispan Log Categories

Infinispan provides categories for INFO, WARN, ERROR, FATAL level messages that organize logs by functional area.

org.infinispan.CLUSTER

Messages specific to Infinispan clustering that include state transfer operations, rebalancing events, partitioning, and so on.

org.infinispan.CONFIG

Messages specific to Infinispan configuration.

org.infinispan.CONTAINER

Messages specific to the data container that include expiration and eviction operations, cache listener notifications, transactions, and so on.

org.infinispan.PERSISTENCE

Messages specific to cache loaders and stores.

org.infinispan.SECURITY

Messages specific to Infinispan security.

org.infinispan.SERVER

Messages specific to Infinispan servers.

org.infinispan.XSITE

Messages specific to cross-site replication operations.

Root Logger

The root logger is `org.infinispan` and is always configured. This logger captures all messages for Infinispan log categories.

Log Handlers

Log handlers define how Infinispan records log messages.

CONSOLE

Write log messages to the host standard out (`stdout`) or standard error (`stderr`) stream. Uses the `org.jboss.logmanager.handlers.ConsoleHandler` class by default.

FILE

Write log messages to a file.

Uses the `org.jboss.logmanager.handlers.PeriodicRotatingFileHandler` class by default.

JSON-FILE

Write log messages to a file in JSON format.

Uses the `org.jboss.logmanager.handlers.PeriodicRotatingFileHandler` class by default.

Log Formatters

Log formatters:

- Configure log handlers and define the appearance of log messages.
- Are strings that use syntax based on the `java.util.logging.Formatter` class.

An example is the default pattern for messages with the FILE log handler:

```
%d{yyyy-MM-dd HH:mm:ss,SSS} %-5p [%c] (%t) %s%e%n
```

- `%d` adds the current time and date.
- `%-5p` specifies the log level.
- `%c` specifies the logging category.
- `%t` adds the name of the current thread.
- `%s` specifies the simple log message.
- `%e` adds the exception stack trace.
- `%n` adds a new line.

Reference

[java.util.logging.Formatter](#)

Enabling and Configuring the JSON Log Handler

Infinispan provides a JSON log handler to write messages in JSON format.

Prerequisites

Ensure that Infinispan is not running. You cannot dynamically enable log handlers.

Procedure

1. Open `$ISPN_HOME/${infinispan.server.root}/conf/logging.properties` with any text editor.
2. Add `JSON-FILE` as a log handler, for example:

```
logger.handlers=CONSOLE,FILE,JSON-FILE
```

3. Optionally configure the JSON log handler and formatter.
 - a. Use the `handler.JSON-FILE` property to configure the JSON log handler.
 - b. Use the `formatter.JSON-FORMATTER` property to configure the JSON log formatter.

4. Save and close `logging.properties`.

When you start Infinispan, it writes each log message as a JSON map in the following file:

```
$ISPN_HOME/${infinispan.server.root}/log/server.log.json
```

3.1.3. Logging Framework

Infinispan uses the JBoss Logging Framework and delegates to logging providers in the following order:

1. JBoss Logging, if you are running Infinispan servers.
2. Apache Log4j, if `org.apache.log4j.LogManager` and `org.apache.log4j.Hierarchy` are on the classpath.
3. LogBack, if `ch.qos.logback.classic.Logger` is on the classpath.
4. JDK logging (`java.util.logging`), if no other logging provider is available.

3.1.4. Access Logs

Hot Rod and REST endpoints can record all inbound client requests as log entries with the following categories:

- `org.infinispan.HOTROD_ACCESS_LOG` logging category for the Hot Rod endpoint.
- `org.infinispan.REST_ACCESS_LOG` logging category for the REST endpoint.

Enabling Access Logs

Access logs for Hot Rod and REST endpoints are disabled by default. To enable either logging category, set the level to `TRACE` in the Infinispan server configuration, as in the following example:

```
<logger category="org.infinispan.HOTROD_ACCESS_LOG" use-parent-handlers="false">
  <level name="TRACE"/>
  <handlers>
    <handler name="HR-ACCESS-FILE"/>
  </handlers>
</logger>
```

Access Log Properties

The default format for access logs is as follows:

```
`%X{address} %X{user} [%d{dd/MMM/yyyy:HH:mm:ss z}] &quot;%X{method} %m
%X{protocol}&quot;; %X{status} %X{requestSize} %X{responseSize} %X{duration}%n`
```

The preceding format creates log entries such as the following:

```
127.0.0.1 - [DD/MM/YYYY:HH:MM:SS] "PUT /rest/v2/caches/default/key HTTP/1.1" 404 5 77 10
```

Logging properties use the `%X{name}` notation and let you modify the format of access logs. The following are the default logging properties:

Property	Description
<code>address</code>	Either the <code>X-Forwarded-For</code> header or the client IP address.
<code>user</code>	Principal name, if using authentication.
<code>method</code>	Method used. <code>PUT</code> , <code>GET</code> , and so on.
<code>protocol</code>	Protocol used. <code>HTTP/1.1</code> , <code>HTTP/2</code> , <code>HOTROD/2.9</code> , and so on.
<code>status</code>	An HTTP status code for the REST endpoint. <code>OK</code> or an exception for the Hot Rod endpoint.
<code>requestSize</code>	Size, in bytes, of the request.
<code>responseSize</code>	Size, in bytes, of the response.
<code>duration</code>	Number of milliseconds that the server took to handle the request.



Use the header name prefixed with `h:` to log headers that were included in requests; for example, `%X{h:User-Agent}`.

3.2. Retrieving Server Health Statistics

Monitor the health of your Infinispan clusters in the following ways:

- Programmatically with `embeddedCacheManager.getHealth()` method calls.
- JMX MBeans
- Infinispan REST Server

3.2.1. Accessing the Health API via JMX

Retrieve Infinispan cluster health statistics via JMX.

Procedure

1. Connect to Infinispan server using any JMX capable tool such as JConsole and navigate to the following object:

```
org.infinispan:type=CacheManager,name="default",component=CacheContainerHealth
```

2. Select available MBeans to retrieve cluster health statistics.

3.2.2. Accessing the Health API via REST

Get Infinispan cluster health via the REST API.

Procedure

- Invoke a **GET** request to retrieve cluster health.

```
GET /rest/v2/cache-managers/{cacheManagerName}/health
```

Infinispan responds with a **JSON** document such as the following:

```
{
  "cluster_health":{
    "cluster_name":"ISPN",
    "health_status":"HEALTHY",
    "number_of_nodes":2,
    "node_names":[
      "NodeA-36229",
      "NodeB-28703"
    ]
  },
  "cache_health":[
    {
      "status":"HEALTHY",
      "cache_name":"__protobuf_metadata"
    },
    {
      "status":"HEALTHY",
      "cache_name":"cache2"
    },
    {
      "status":"HEALTHY",
      "cache_name":"mycache"
    },
    {
      "status":"HEALTHY",
      "cache_name":"cache1"
    }
  ]
}
```



Get cache manager status as follows:

```
GET /rest/v2/cache-managers/{cacheManagerName}/health/status
```

See the *REST v2 (version 2) API* documentation for more information.

3.3. Monitoring Infinispan Servers Over JMX

You can monitor an Infinispan Server over JMX in two ways:

- Use JConsole or VisualVM running locally as the same user. This will use a local `jvmstat` connection and requires no additional setup
- Use JMX remoting (aka JSR-160) to connect from any host. This requires connecting through the management port (usually 9990) using a special protocol which respects the server security configuration

To setup a client for JMX remoting you need to add the `$ISPN_HOME/bin/client/jboss-client.jar` to your client's classpath and use one of the following service URLs:

- `service:jmx:remote-http-jmx://host:port` for plain connections through the management interface
- `service:jmx:remote-https-jmx://host:port` for TLS connections through the management interface (although this requires having the appropriate keys available)
- `service:jmx:remoting-jmx://localhost:port` for connections through the remoting interface (necessary for connecting to individual servers in a domain)

The JMX subsystem registers a service with the Remoting endpoint so that remote access to JMX can be obtained over the exposed Remoting connector. This is switched on by default in standalone mode and accessible over port 9990 but in domain mode it is switched off so it needs to be enabled. In domain mode the port will be the port of the Remoting connector for the Server instance to be monitored.

```
<subsystem xmlns="urn:jboss:domain:jmx:1.3">
  <expose-resolved-model/>
  <expose-expression-model/>
  <remoting-connector use-management-endpoint="false"/>
</subsystem>
```

3.3.1. JMX Monitoring

Management of Infinispan instances is all about exposing as much relevant statistical information that allows administrators to get a view of the state of each Infinispan instance. Taking in account that a single installation could be made up of several tens or hundreds Infinispan instances, providing clear and concise information in an efficient manner is imperative. The following sections dive into the range of management tooling that Infinispan provides.



Any management tool that supports JMX already has basic support for Infinispan. However, custom plugins could be written to adapt the JMX information for easier consumption.

JMX

Over the years, [JMX](#) has become the de facto standard for management and administration of middleware and as a result, the Infinispan team has decided to standardize on this technology for the exposure of management and statistical information.

Understanding The Exposed MBeans

By connecting to the VM(s) where Infinispan is running with a standard JMX GUI such as [JConsole](#) or [VisualVM](#) you should find the following MBeans:

- For CacheManager level JMX statistics, without further configuration, you should see an MBean called `org.infinispan:type=CacheManager,name="DefaultCacheManager"` with [properties specified by the CacheManager MBean](#) .
- Using the `cacheManagerName` attribute in `globalJmxStatistics` XML element, or using the corresponding `GlobalJmxStatisticsConfigurationBuilder.cacheManagerName(String cacheManagerName)` call, you can name the cache manager in such way that the name is used as part of the JMX object name. So, if the name had been "Hibernate2LC", the JMX name for the cache manager would have been: `org.infinispan:type=CacheManager,name="Hibernate2LC"` . This offers a nice and clean way to manage environments where multiple cache managers are deployed, which follows [JMX best practices](#) .
- For Cache level JMX statistics, you should see several different MBeans depending on which configuration options have been enabled. For example, if you have configured a write behind cache store, you should see an MBean exposing properties belonging to the cache store component. All Cache level MBeans follow the same format though which is the following: `org.infinispan:type=Cache,name="{name-of-cache}({cache-mode})",manager="{name-of-cache-manager}",component={component-name}` where:
 - `{name-of-cache}` has been substituted by the actual cache name. If this cache represents the default cache, its name will be `__defaultCache`.
 - `{cache-mode}` has been substituted by the cache mode of the cache. The cache mode is represented by the lower case version of the possible enumeration values shown [here](#).
 - `{name-of-cache-manager}` has been substituted by the name of the cache manager to which this cache belongs. The name is derived from the `cacheManagerName` attribute value in `globalJmxStatistics` element.
 - `{component-name}` has been substituted by one of the JMX component names in the [JMX reference documentation](#) .

For example, the cache store JMX component MBean for a default cache configured with synchronous distribution would have the following name: `org.infinispan:type=Cache,name="__defaultcache(dist_sync)",manager="DefaultCacheManager",component=CacheStore`

Please note that cache and cache manager names are quoted to protect against illegal characters being used in these user-defined names.

Enabling JMX Statistics

The MBeans mentioned in the previous section are always created and registered in the

MBeanServer allowing you to manage your caches but some of their attributes do not expose meaningful values unless you take the extra step of enabling collection of statistics. Gathering and reporting statistics via JMX can be enabled at 2 different levels:

CacheManager level

The CacheManager is the entity that governs all the cache instances that have been created from it. Enabling CacheManager statistics collections differs depending on the configuration style:

- If configuring the CacheManager via XML, make sure you add the following XML under the `<cache-container />` element:

```
<cache-container statistics="true"/>
```

- If configuring the CacheManager programmatically, simply add the following code:

```
GlobalConfigurationBuilder globalConfigurationBuilder = ...
globalConfigurationBuilder.globalJmxStatistics().enable();
```

Cache level

At this level, you will receive management information generated by individual cache instances. Enabling Cache statistics collections differs depending on the configuration style:

- If configuring the Cache via XML, make sure you add the following XML under the one of the top level cache elements, such as `<local-cache />`:

```
<local-cache statistics="true"/>
```

- If configuring the Cache programmatically, simply add the following code:

```
ConfigurationBuilder configurationBuilder = ...
configurationBuilder.jmxStatistics().enable();
```

Monitoring cluster health

It is also possible to monitor Infinispan cluster health using JMX. On CacheManager there's an additional object called `CacheContainerHealth`. It contains the following attributes:

- `cacheHealth` - a list of caches and corresponding statuses (HEALTHY, DEGRADED or HEALTHY_REBALANCING)
- `clusterHealth` - overall cluster health
- `clusterName` - cluster name
- `freeMemoryKb` - Free memory obtained from JVM runtime measured in KB
- `numberOfCpus` - The number of CPUs obtained from JVM runtime

- numberOfNodes - The number of nodes in the cluster
- totalMemoryKb - Total memory obtained from JVM runtime measured in KB

Multiple JMX Domains

There can be situations where several CacheManager instances are created in a single VM, or Cache names belonging to different CacheManagers under the same VM clash.

Using different JMX domains for multi cache manager environments should be last resort. Instead, it's possible to name a cache manager in such way that it can easily be identified and used by monitoring tools. For example:

- Via XML:

```
<cache-container statistics="true" name="Hibernate2LC"/>
```

- Programmatically:

```
GlobalConfigurationBuilder globalConfigurationBuilder = ...
globalConfigurationBuilder.globalJmxStatistics()
    .enable()
    .cacheManagerName("Hibernate2LC");
```

Using either of these options should result on the CacheManager MBean name being: `org.infinispan:type=CacheManager,name="Hibernate2LC"`

For the time being, you can still set your own jmxDomain if you need to and we also allow duplicate domains, or rather duplicate JMX names, but these should be limited to very special cases where different cache managers within the same JVM are named equally.

Registering MBeans In Non-Default MBean Servers

Let's discuss where Infinispan registers all these MBeans. By default, Infinispan registers them in the [standard JVM MBeanServer platform](#). However, users might want to register these MBeans in a different MBeanServer instance. For example, an application server might work with a different MBeanServer instance to the default platform one. In such cases, users should implement the [MBeanServerLookup interface](#) provided by Infinispan so that the `getMBeanServer()` method returns the MBeanServer under which Infinispan should register the management MBeans. Once you have your implementation ready, simply configure Infinispan with the fully qualified name of this class. For example:

- Via XML:

```
<cache-container statistics="true">
  <jmx mbean-server-lookup="com.acme.MyMBeanServerLookup" />
</cache-container>
```


- Programmatically:

```
GlobalConfigurationBuilder globalConfigurationBuilder = ...
globalConfigurationBuilder.globalJmxStatistics()
    .enable()
    .mBeanServerLookup(new com.acme.MyMBeanServerLookup());
```

Available MBeans

For a complete list of available MBeans, refer to the [JMX reference documentation](#)

Chapter 4. Securing Infinispan Servers

4.1. Security Concepts

4.1.1. Authorization

Just like embedded mode, the server supports cache authorization using the same configuration, e.g.:

```
<infinispan>
  <cache-container default-cache="secured" name="secured">
    <security>
      <authorization>
        <identity-role-mapper />
        <role name="admin" permissions="ALL" />
        <role name="reader" permissions="READ" />
        <role name="writer" permissions="WRITE" />
        <role name="supervisor" permissions="READ WRITE EXEC"/>
      </authorization>
    </security>
    <local-cache name="secured">
      <security>
        <authorization roles="admin reader writer supervisor" />
      </security>
    </local-cache>
  </cache-container>
</infinispan>
```

4.1.2. Server Realms

Infinispan Server security is built around the features provided by the underlying server realm and security domains. Security Realms are used by the server to provide authentication and authorization information for both the management and application interfaces.

```
<server xmlns="urn:jboss:domain:2.1">
  ...
  <management>
    ...
    <security-realm name="ApplicationRealm">
      <authentication>
        <properties path="application-users.properties" relative-to=
"jboss.server.config.dir"/>
      </authentication>
      <authorization>
        <properties path="application-roles.properties" relative-to=
"jboss.server.config.dir"/>
      </authorization>
    </security-realm>
    ...
  </management>
  ...
</server>
```

Infinispan Server comes with an `add-user.sh` script (`add-user.bat` for Windows) to ease the process of adding new user/role mappings to the above property files. An example invocation for adding a user to the `ApplicationRealm` with an initial set of roles:

```
./bin/add-user.sh -a -u myuser -p "qwer1234!" -ro supervisor,reader,writer
```

It is also possible to authenticate/authorize against alternative sources, such as LDAP, JAAS, etc.

Bear in mind that the choice of authentication mechanism you select for the protocols limits the type of authentication sources, since the credentials must be in a format supported by the algorithm itself (e.g. pre-digested passwords for the digest algorithm)

4.2. Hot Rod Authentication

The Hot Rod protocol supports authentication by leveraging the SASL mechanisms. The supported SASL mechanisms (usually shortened as mechs) are:

- **PLAIN** - This is the most insecure mech, since credentials are sent over the wire in plain-text format, however it is the simplest to get to work. In combination with encryption (i.e. TLS) it can be used safely
- **DIGEST-MD5** - This mech hashes the credentials before sending them over the wire, so it is more secure than PLAIN
- **GSSAPI** - This mech uses Kerberos tickets, and therefore requires the presence of a properly configured Kerberos Domain Controller (such as Microsoft Active Directory)
- **EXTERNAL** - This mech obtains credentials from the underlying transport (i.e. from a X.509 client certificate) and therefore requires encryption using client-certificates to be enabled.

The following configuration enables authentication against ApplicationRealm, using the DIGEST-MD5 SASL mechanism and only enables the **auth** QoP (see [SASL Quality of Protection](#)):

Hot Rod connector configuration

```
<hotrod-connector socket-binding="hotrod" cache-container="default">
  <authentication security-realm="ApplicationRealm">
    <sasl server-name="myhotrodserver" mechanisms="DIGEST-MD5" qop="auth" />
  </authentication>
</hotrod-connector>
```

Notice the server-name attribute: it is the name that the server declares to incoming clients and therefore the client configuration must match. It is particularly important when using GSSAPI as it is equivalent to the Kerberos service name. You can specify multiple mechanisms and they will be attempted in order.

4.2.1. SASL Quality of Protection

While the main purpose of SASL is to provide authentication, some mechanisms also support integrity and privacy protection, also known as Quality of Protection (or qop). During authentication negotiation, ciphers are exchanged between client and server, and they can be used to add checksums and encryption to all subsequent traffic. You can tune the required level of qop as follows:

QOP	Description
auth	Authentication only
auth-int	Authentication with integrity protection
auth-conf	Authentication with integrity and privacy protection

4.2.2. SASL Policies

You can further refine the way a mechanism is chosen by tuning the SASL policies. This will effectively include / exclude mechanisms based on whether they match the desired policies.

Policy	Description
forward-secrecy	Specifies that the selected SASL mechanism must support forward secrecy between sessions. This means that breaking into one session will not automatically provide information for breaking into future sessions.
pass-credentials	Specifies that the selected SASL mechanism must require client credentials.
no-plain-text	Specifies that the selected SASL mechanism must not be susceptible to simple plain passive attacks.
no-active	Specifies that the selected SASL mechanism must not be susceptible to active (non-dictionary) attacks. The mechanism might require mutual authentication as a way to prevent active attacks.

Policy	Description
no-dictionary	Specifies that the selected SASL mechanism must not be susceptible to passive dictionary attacks.
no-anonymous	Specifies that the selected SASL mechanism must not accept anonymous logins.

Each policy's value is either "true" or "false". If a policy is absent, then the chosen mechanism need not have that characteristic (equivalent to setting the policy to "false"). One notable exception is the **no-anonymous** policy which, if absent, defaults to true, thus preventing anonymous connections.



It is possible to have mixed anonymous and authenticated connections to the endpoint, delegating actual access logic to cache authorization configuration. To do so, set the **no-anonymous** policy to false and turn on cache authorization.

The following configuration selects all available mechanisms, but effectively only enables GSSAPI, since it is the only one that respects all chosen policies:

Hot Rod connector policies

```
<hotrod-connector socket-binding="hotrod" cache-container="default">
  <authentication security-realm="ApplicationRealm">
    <sasl server-name="myhotrodserver" mechanisms="PLAIN DIGEST-MD5 GSSAPI EXTERNAL"
qop="auth">
      <policy>
        <no-active value="true" />
        <no-anonymous value="true" />
        <no-plain-text value="true" />
      </policy>
    </sasl>
  </authentication>
</hotrod-connector>
```

4.2.3. Using GSSAPI/Kerberos

If you want to use GSSAPI/Kerberos, setup and configuration differs. First we need to define a Kerberos login module using the security domain subsystem:

Security domain configuration

```
<system-properties>
  <property name="java.security.krb5.conf" value="/tmp/infinispan/krb5.conf"/>
  <property name="java.security.krb5.debug" value="true"/>
  <property name="jboss.security.disable.secdomain.option" value="true"/>
</system-properties>

<security-domain name="infinispan-server" cache-type="default">
  <authentication>
    <login-module code="Kerberos" flag="required">
      <module-option name="debug" value="true"/>
      <module-option name="storeKey" value="true"/>
      <module-option name="refreshKrb5Config" value="true"/>
      <module-option name="useKeyTab" value="true"/>
      <module-option name="doNotPrompt" value="true"/>
      <module-option name="keyTab" value="/tmp/infinispan/infinispan.keytab"/>
      <module-option name="principal" value="HOTROD/localhost@INFINISPAN.ORG"/>
    </login-module>
  </authentication>
</security-domain>
```

Next we need to modify the Hot Rod connector

Hot Rod connector configuration

```
<hotrod-connector socket-binding="hotrod" cache-container="default">
  <authentication security-realm="ApplicationRealm">
    <sasl server-name="infinispan-server" server-context-name="infinispan-server"
mechanisms="GSSAPI" qop="auth" />
  </authentication>
</hotrod-connector>
```

4.3. Hot Rod and REST encryption (TLS/SSL)

Both Hot Rod and REST protocols support encryption using SSL/TLS with optional TLS/SNI support ([Server Name Indication](#)). To set this up you need to create a keystore using the keytool application which is part of the JDK to store your server certificate. Then add a <server-identities> element to your security realm:

```
<security-realm name="ApplicationRealm">
  <server-identities>
    <ssl>
      <keystore path="keystore_server.jks" relative-to="jboss.server.config.dir"
keystore-password="secret" />
    </ssl>
  </server-identities>
</security-realm>
```



When using SNI support there might be multiple Security Realms configured.

It is also possible to generate development certificates on server startup. In order to do this, just specify `generate-self-signed-certificate-host` in the keystore element as shown below:

Generating Keystore automatically

```
<security-realm name="ApplicationRealm">
  <server-identities>
    <ssl>
      <keystore path="keystore_server.jks" relative-to="jboss.server.config.dir"
keystore-password="secret" generate-self-signed-certificate-host="localhost"/>
    </ssl>
  </server-identities>
</security-realm>
```



There are three basic principles that you should remember when using automatically generated keystores:

- They shouldn't be used on a production environment
- They are generated when necessary (e.g. while obtaining the first connection from the client)
- They contain also certificates so they might be used in a Hot Rod client directly

Next modify the `<hotrod-connector>` and/or `<rest-connector>` elements in the endpoint subsystem to require encryption. Optionally add SNI configuration:

```
<hotrod-connector socket-binding="hotrod" cache-container="local">
  <encryption security-realm="ApplicationRealm" require-ssl-client-auth="false">
    <sni host-name="domain1" security-realm="Domain1ApplicationRealm" />
    <sni host-name="domain2" security-realm="Domain2ApplicationRealm" />
  </encryption>
</hotrod-connector>
<rest-connector socket-binding="rest" cache-container="local">
  <encryption security-realm="ApplicationRealm" require-ssl-client-auth="false">
    <sni host-name="domain1" security-realm="Domain1ApplicationRealm" />
    <sni host-name="domain2" security-realm="Domain2ApplicationRealm" />
  </encryption>
</rest-connector>
```



To configure the client In order to connect to the server using the Hot Rod protocol, the client needs a trust store containing the public key of the server(s) you are going to connect to, unless the key was signed by a Certification Authority (CA) trusted by the JRE.

```
ConfigurationBuilder clientBuilder = new ConfigurationBuilder();
clientBuilder
  .addServer()
    .host("127.0.0.1")
    .port(11222)
  .security()
    .ssl()
      .enabled(true)
      .sniHostName("domain1")
      .trustStoreFileName("truststore_client.jks")
      .trustStorePassword("secret".toCharArray());
remoteCacheManager = new RemoteCacheManager(clientBuilder.build());
```

Additionally, you might also want to enable client certificate authentication (and optionally also allow the use of the EXTERNAL SASL mech to authenticate and authorize clients). To enable this you will need the security realm on the server to be able to trust incoming client certificates by adding a trust store:


```
<security-realm name="ApplicationRealm">
  <authentication>
    <truststore path="truststore_server.jks" relative-to="jboss.server.config.dir"
keystore-password="secret"/>
  </authentication>
  <server-identities>
    <ssl>
      <keystore path="keystore_server.jks" relative-to="jboss.server.config.dir"
keystore-password="secret" />
    </ssl>
  </server-identities>
</security-realm>
```

And then tell the connector to require a client certificate:

```
<hotrod-connector socket-binding="hotrod" cache-container="local">
  <encryption security-realm="ApplicationRealm" require-ssl-client-auth="true" />
</hotrod-connector>
```

The client, at this point, will also need to specify a keyStore which contains its certificate on top of the trustStore which trusts the server certificate.

Chapter 5. Single Port

Single-Port is a special type of router connector which allows exposing multiple protocols over the same TCP port. This approach is very convenient because it reduces the number of ports required by a server, with advantages in security, configuration and management. Protocol switching is handled in three ways:

- **HTTP/1.1 Upgrade header:** initiate an HTTP/1.1 connection and send an `Upgrade: protocol` header where protocol is the name assigned to the desired endpoint.
- **TLS/ALPN:** protocol selection is performed based on the SNI specified by the client.
- **Hot Rod header detection:** if a Hot Rod endpoint is present in the router configuration, then any attempt to send a Hot Rod header will be detected and the protocol will be switched automatically.



The initial implementation supports only HTTP/1.1, HTTP/2 and Hot Rod protocols. The Memcached protocol is not supported.

5.1. Single-Port router

Internally, Single-Port is based on the same router component used to enable multi-tenancy, and therefore it shares the same configuration.

```
<!-- TLS/ALPN negotiation -->
<router-connector name="router-ssl" single-port-socket-binding="rest-ssl">
  <single-port security-realm="SSLRealm1">
    <hotrod name="hotrod" />
    <rest name="rest" />
  </single-port>
</router-connector>
<!-- HTTP 1.1/Upgrade procedure -->
<router-connector name="router" single-port-socket-binding="rest">
  <single-port>
    <hotrod name="hotrod" />
    <rest name="rest" />
  </single-port>
</router-connector>
```

With the configuration above, the Single-Port Router will operate on `rest` and `rest-ssl` socket bindings. The router named `router` should typically operate on port `8080` and will use HTTP/1.1 Upgrade (also known as *cleartext upgrade*) procedure. The other router instance (called `router-ssl`) should typically operate on port `8443` and will use TLS/ALPN.

5.1.1. Testing the Single-Port router

A tool such as `curl` can be used to access cache using both *cleartext upgrade* or TLS/ALPN. Here's an example:

```
$ curl -v -k --http2-prior-knowledge
https://127.0.0.1:8443/rest/v2/caches/default/test
```

The `--http2-prior-knowledge` can be exchanged with `--http2` switch allowing to control how the switch procedure is being done (via Plain-Text Upgrade or TLS/ALPN).

5.2. Hot Rod Protocol Detection

The single-port router has built-in automatic detection of Hot Rod messages which trigger a transparent "upgrade" to the Hot Rod protocol. This means that no changes are required on the client side to connect to a single-port endpoint. It also means that older clients will also be able to function seamlessly.

5.2.1. TLS/ALPN protocol selection

Another supported way to select the protocol is to use TLS/ALPN which uses the [Application-Layer Protocol Negotiation](#) spec. This feature requires that you have configured your endpoint to enable TLS.

Enabling ALPN

If you are using JDK 9 or greater, ALPN is supported by default. However, if you are using JDK 8, you will need to use [Netty's BoringSSL](#) library, which leverages native libraries to enable ALPN.

1. Add Netty dependencies.

```
<dependencyManagement>
  <dependency>
    <groupId>io.netty</groupId>
    <artifactId>netty-bom</artifactId>
    <!-- Pulled from Infinispan BOM -->
    <version>${version.netty}</version>
    <type>pom</type>
    <scope>import</scope>
  </dependency>
</dependencyManagement>

<dependency>
  <groupId>io.netty</groupId>
  <artifactId>netty-tcnative-boringssl-static</artifactId>
  <!-- The version is defined in Netty BOM -->
</dependency>
```

2. Configure your trust store accordingly:

```
ConfigurationBuilder builder = new ConfigurationBuilder()
    .addServers("127.0.0.1:8443");

builder.security().ssl().enable()
    .trustStoreFileName("truststore.pkcs12")
    .trustStorePassword(DEFAULT_TRUSTSTORE_PASSWORD.toCharArray());

RemoteCacheManager remoteCacheManager = new RemoteCacheManager(builder.build());
RemoteCache<String, String> cache = remoteCacheManager.getCache("default");
```

Chapter 6. Remotely Executing Server-Side Tasks

Define and add tasks to Infinispan servers that you can invoke from the Infinispan command line interface, REST API, or from Hot Rod clients.

You can implement tasks as custom Java classes or define scripts in languages such as JavaScript.

6.1. Creating Server Tasks

Create custom task implementations and add them to Infinispan servers.

6.1.1. Server Tasks

Infinispan server tasks are classes that extend the `org.infinispan.tasks.ServerTask` interface and generally include the following method calls:

`setTaskContext()`

Allows access to execution context information including task parameters, cache references on which tasks are executed, and so on. In most cases, implementations store this information locally and use it when tasks are actually executed.

`getName()`

Returns unique names for tasks. Clients invoke tasks with these names.

`getExecutionMode()`

Returns the execution mode for tasks.

- `TaskExecutionMode.ONE_NODE` only the node that handles the request executes the script. Although scripts can still invoke clustered operations.
- `TaskExecutionMode.ALL_NODES` Infinispan uses clustered executors to run scripts across nodes. For example, server tasks that invoke stream processing need to be executed on a single node because stream processing is distributed to all nodes.

`call()`

Computes a result. This method is defined in the `java.util.concurrent.Callable` interface and is invoked with server tasks.



Server task implementations must adhere to service loader pattern requirements. For example, implementations must have a zero-argument constructors.

The following `HelloTask` class implementation provides an example task that has one parameter:

```

package example;

import org.infinispan.tasks.ServerTask;
import org.infinispan.tasks.TaskContext;

public class HelloTask implements ServerTask<String> {

    private TaskContext ctx;

    @Override
    public void setTaskContext(TaskContext ctx) {
        this.ctx = ctx;
    }

    @Override
    public String call() throws Exception {
        String name = (String) ctx.getParameters().get().get("name");
        return "Hello " + name;
    }

    @Override
    public String getName() {
        return "hello-task";
    }

}

```

Reference

- [org.infinispan.tasks.ServerTask](#)
- [java.util.concurrent.Callable.call\(\)](#)
- [java.util.ServiceLoader](#)

6.1.2. Deploying Server Tasks to Infinispan Servers

Add your custom server task classes to Infinispan servers.

Prerequisites

Stop any running Infinispan servers. Infinispan does not support runtime deployment of custom classes.

Procedure

1. Package your server task implementation in a JAR file.
2. Add a `META-INF/services/org.infinispan.tasks.ServerTask` file that contains the fully qualified names of server tasks, for example:

```
example.HelloTask
```

3. Copy the JAR file to the `$ISPN_HOME/server/lib` directory of your Infinispan server.
4. Add your classes to the deserialization whitelist in your Infinispan configuration. Alternatively set the whitelist using system properties.

Reference

- [Adding Java Classes to Deserialization White Lists](#)
- [Infinispan 10.1 Configuration Schema](#)

6.2. Creating Server Scripts

Create custom scripts and add them to Infinispan servers.

6.2.1. Server Scripts

Infinispan server scripting is based on the `javax.script` API and is compatible with any JVM-based ScriptEngine implementation. Nashorn is the default JDK ScriptEngine and provides JavaScript capabilities.

Hello World Script Example

The following script provides a simple example that runs on a single Infinispan server, has one parameter, and uses JavaScript:

```
// mode=local,language=javascript,parameters=[greetee]
"Hello " + greetee
```

When you run the preceding script, you pass a value for the `greetee` parameter and Infinispan returns `"Hello ${value}"`.

Script Metadata

Metadata provides additional information about scripts that Infinispan servers use when running scripts.

Script metadata are `property=value` pairs that you add to comments in the first lines of scripts, such as the following example:

```
// name=test, language=javascript
// mode=local, parameters=[a,b,c]
```

- Use comment styles that match the scripting language (`//`, `;;`, `#`).
- Separate `property=value` pairs with commas.
- Separate values with single (') or double (") quote characters.

Table 1. Metadata Properties

Property	Description
<code>mode</code>	<p>Defines the execution mode and has the following values:</p> <p><code>local</code> only the node that handles the request executes the script. Although scripts can still invoke clustered operations.</p> <p><code>distributed</code> Infinispan uses clustered executors to run scripts across nodes.</p>
<code>language</code>	Specifies the ScriptEngine that executes the script.
<code>extension</code>	Specifies filename extensions as an alternative method to set the ScriptEngine.
<code>role</code>	Specifies roles that users must have to execute scripts.
<code>parameters</code>	Specifies an array of valid parameter names for this script. Invocations which specify parameters not included in this list cause exceptions.
<code>datatype</code>	<p>Optionally sets the MediaType (MIME type) for storing data as well as parameter and return values. This property is useful for remote clients that support particular data formats only.</p> <p>Currently you can set only <code>text/plain; charset=utf-8</code> to use the String UTF-8 format for data.</p>

Script Bindings

Infinispan exposes internal objects as bindings for script execution.

Binding	Description
<code>cache</code>	Specifies the cache against which the script is run.
<code>marshaller</code>	Specifies the marshaller to use for serializing data to the cache.
<code>cacheManager</code>	Specifies the <code>cacheManager</code> for the cache.
<code>scriptingManager</code>	Specifies the instance of the script manager that runs the script. You can use this binding to run other scripts from a script.

Script Parameters

Infinispan lets you pass named parameters as bindings for running scripts.

Parameters are `name,value` pairs, where `name` is a string and `value` is any value that the marshaller can interpret.

The following example script has two parameters, `multiplicand` and `multiplier`. The script takes the value of `multiplicand` and multiplies it with the value of `multiplier`.

```
// mode=local,language=javascript
multiplicand * multiplier
```

When you run the preceding script, Infinispan responds with the result of the expression evaluation.

6.2.2. Adding Scripts to Infinispan Servers

Use the command line interface to add scripts to Infinispan servers.

Prerequisites

Infinispan servers store scripts in the `__script_cache` cache. If you enable cache authorization, users must have the `__script_manager` role to access `__script_cache`.

Procedure

1. Define scripts as required.

For example, create a file named `multiplication.js` that runs on a single Infinispan server, has two parameters, and uses JavaScript to multiply a given value:

```
// mode=local,language=javascript
multiplicand * multiplier
```

2. Open a CLI connection to Infinispan and use the `task` command to upload your scripts as in the following example:

```
[//containers/default]> task upload --file=multiplication.js multiplication
```

3. Verify that your scripts are available.

```
[//containers/default]> ls tasks
multiplication
```

6.2.3. Programmatically Creating Scripts

Add scripts with the Hot Rod `RemoteCache` interface as in the following example:

```
RemoteCache<String, String> scriptCache = cacheManager.getCache("__script_cache");
scriptCache.put("multiplication.js",
    "// mode=local,language=javascript\n" +
    "multiplicand * multiplier\n");
```

Reference

org.infinispan.client.hotrod.RemoteCache

6.3. Running Server-Side Tasks and Scripts

Execute tasks and custom scripts on Infinispan servers.

6.3.1. Running Tasks and Scripts

Use the command line interface to run tasks and scripts on Infinispan servers.

Prerequisites

- Open a CLI connection to Infinispan.

Procedure

Use the `task` command to run tasks and scripts on Infinispan servers, as in the following examples:

- Execute a script named `multiplier.js` and specify two parameters:

```
[//containers/default]> task exec multiplier.js -Pmultiplicand=10 -Pmultiplier=20
200.0
```

- Execute a task named `@@cache@names` to retrieve a list of all available caches:

```
//containers/default]> task exec @@cache@names
["__protobuf_metadata","mycache","__script_cache"]
```

6.3.2. Programmatically Running Scripts

Call the `execute()` method to run scripts with the Hot Rod `RemoteCache` interface, as in the following example:

```
RemoteCache<String, Integer> cache = cacheManager.getCache();
// Create parameters for script execution.
Map<String, Object> params = new HashMap<>();
params.put("multiplicand", 10);
params.put("multiplier", 20);
// Run the script with the parameters.
Object result = cache.execute("multiplication.js", params);
```

Reference

org.infinispan.client.hotrod.RemoteCache

6.3.3. Programmatically Running Tasks

Call the `execute()` method to run tasks with the Hot Rod `RemoteCache` interface, as in the following example:

```
// Add configuration for a locally running server.
ConfigurationBuilder builder = new ConfigurationBuilder();
builder.addServer().host("127.0.0.1").port(11222);

// Connect to the server.
RemoteCacheManager cacheManager = new RemoteCacheManager(builder.build());

// Retrieve the remote cache.
RemoteCache<String, String> cache = cacheManager.getCache();

// Create task parameters.
Map<String, String> parameters = new HashMap<>();
parameters.put("name", "developer");

// Run the server task.
String greet = cache.execute("hello-task", parameters);
System.out.println(greet);
```

Reference

org.infinispan.client.hotrod.RemoteCache

Chapter 7. Adding Custom Marshaller Bridges

Infinispan provides two marshalling bridges for marshalling client/server requests using the Kryo and Protostuff libraries. To utilise either of thesemarshallers, you simply place the dependency of the marshaller you require in your client pom. Custom schemas for object marshalling must then be registered with the selected library using the library's api on the client or by implementing a RegistryService for the given marshaller bridge.

7.1. Protostuff

Add the protostuff marshaller dependency to your pom:

```
<dependency>
  <groupId>org.infinispan</groupId>
  <artifactId>infinispan-marshaller-protostuff</artifactId>
  <!-- Replace ${version.infinispan} with the
  version of Infinispan that you're using. -->
  <version>${version.infinispan}</version>
</dependency>
```

To register custom Protostuff schemas in your own code, you must register the custom schema with Protostuff before any marshalling begins. This can be achieved by simply calling:

```
RuntimeSchema.register(ExampleObject.class, new ExampleObjectSchema());
```

Or, you can implement a service provider for the `SchemaRegistryService.java` interface, placing all Schema registrations in the `register()` method. Implementations of this interface are loaded via Java's ServiceLoader api, therefore the full path of the implementing class(es) should be provided in a `META-INF/services/org.infinispan.marshaller.protostuff.SchemaRegistryService` file within your deployment jar.

7.2. Kryo

Add the kryo marshaller dependency to your pom:

```
<dependency>
  <groupId>org.infinispan</groupId>
  <artifactId>infinispan-marshaller-kryo</artifactId>
  <!-- Replace ${version.infinispan} with the
  version of Infinispan that you're using. -->
  <version>${version.infinispan}</version>
</dependency>
```

To register custom Kryo serializer in your own code, you must register the custom serializer with Kryo before any marshalling begins. This can be achieved by implementing a service provider for the `SerializerRegistryService.java` interface, placing all serializer registrations in the `register(Kryo)` method; where serializers should be registered with the supplied `Kryo` object using the Kryo api. e.g. `kryo.register(ExampleObject.class, new ExampleObjectSerializer())`. Implementations of this interface are loaded via Java's ServiceLoader api, therefore the full path of the implementing class(es) should be provided in a `META-INF/services/org/infinispan/marshaller/kryo/SerializerRegistryService` file within your deployment jar.

7.3. Storing deserialized objects

When using the Protostuff/Kryo bridges in caches configured with `application/x-java-object` as `MediaType` (storing POJOs instead of binary content) it is necessary for the class files of all custom objects to be placed on the classpath of the server. To achieve this, you should place a jar containing all of their custom classes on the server's classpath.

When utilising a custom marshaller, it is also necessary for the marshaller and its runtime dependencies to be on the server's classpath. To aid with this step we have created a "bundle" jar for each of the bridge implementations which includes all of the runtime class files required by the bridge and underlying library. Therefore, it is only necessary to include this single jar on the server's classpath.

Bundle jar downloads:

- [Kryo Bundle](#)
- [Protostuff Bundle](#)



Jar files containing custom classes must be placed in the same module/directory as the custom marshaller bundle so that the marshaller can load them. i.e. if you register the marshaller bundle in `modules/system/layers/base/org/infinispan/main/modules.xml`, then you must also register your custom classes here.

7.4. Registering Custom Schemas/Serializers

Custom serializers/schemas for the Kryo/Protostuff marshallers must be registered via their respective service interfaces in order to store deserialized objects. To achieve this, it is necessary for a **JAR** that contains the service provider to be registered in the same directory or module as the marshaller bundle and custom classes.



It is not necessary for the service provider implementation to be provided in the same **JAR** as the user's custom classes. However, the **JAR** that contains the provider must be in the same directory/module as the marshaller and custom class **JAR** files.

Chapter 8. Performing Rolling Upgrades

Upgrade Infinispan without downtime or data loss. You can perform rolling upgrades for Infinispan servers to start using a more recent version of Infinispan.



This section explains how to upgrade Infinispan servers, see the appropriate documentation for your Hot Rod client for upgrade procedures.

8.1. Rolling Upgrades

From a high-level, you do the following to perform rolling upgrades:

1. Set up a target cluster. The target cluster is the Infinispan version to which you want to migrate data. The source cluster is the Infinispan deployment that is currently in use. After the target cluster is running, you configure all clients to point to it instead of the source cluster.
2. Synchronize data from the source cluster to the target cluster.

8.2. Setting Up Target Clusters

1. Start the target cluster with unique network properties or a different JGroups cluster name to keep it separate from the source cluster.
2. Configure a `RemoteCacheStore` on the target cluster for each cache you want to migrate from the source cluster.

`RemoteCacheStore` settings

- `remote-server` must point to the source cluster via the `outbound-socket-binding` property.
- `remoteCacheName` must match the cache name on the source cluster.
- `hotrod-wrapping` must be `true` (enabled).
- `shared` must be `true` (enabled).
- `purge` must be `false` (disabled).
- `passivation` must be `false` (disabled).
- `protocol-version` matches the Hot Rod protocol version of the source cluster.

Example RemoteCacheStore Configuration

```
<distributed-cache>
  <remote-store cache="MyCache" socket-timeout="60000" tcp-no-delay="true"
  protocol-version="2.5" shared="true" hotrod-wrapping="true" purge="false"
  passivation="false">
    <remote-server outbound-socket-binding="remote-store-hotrod-server"/>
  </remote-store>
</distributed-cache>
...
<socket-binding-group name="standard-sockets" default-interface="public"
  port-offset="{jboss.socket.binding.port-offset:0}">
  ...
  <outbound-socket-binding name="remote-store-hotrod-server">
    <remote-destination host="198.51.100.0" port="11222"/>
  </outbound-socket-binding>
  ...
</socket-binding-group>
```

3. Configure the target cluster to handle all client requests instead of the source cluster:
 - a. Configure all clients to point to the target cluster instead of the source cluster.
 - b. Restart each client node.

The target cluster lazily loads data from the source cluster on demand via `RemoteCacheStore`.

8.3. Synchronizing Data from Source Clusters

1. Call the `synchronizeData()` method in the `TargetMigrator` interface. Do one of the following on the target cluster for each cache that you want to migrate:

JMX

Invoke the `synchronizeData` operation and specify the `hotrod` parameter on the `RollingUpgradeManager` MBean.

CLI

```
$ bin/cli.sh --connect controller=127.0.0.1:9990 -c "/subsystem=datagrid-
infinispan/cache-container=clustered/distributed-cache=MyCache:synchronize-
data(migrator-name=hotrod)"
```

Data migrates to all nodes in the target cluster in parallel, with each node receiving a subset of the data.

Use the following parameters to tune the operation:

- `read-batch` configures the number of entries to read from the source cluster at a time. The default value is `10000`.

- `write-threads` configures the number of threads used to write data. The default value is the number of processors available.

For example:

```
synchronize-data(migrator-name=hotrod, read-batch=100000, write-threads=3)
```

2. Disable the `RemoteCacheStore` on the target cluster. Do one of the following:

JMX

Invoke the `disconnectSource` operation and specify the `hotrod` parameter on the `RollingUpgradeManager` MBean.

CLI

```
$ bin/cli.sh --connect controller=127.0.0.1:9990 -c "/subsystem=datagrid-infinispan/cache-container=clustered/distributed-cache=MyCache:disconnect-source(migrator-name=hotrod)"
```

3. Decommission the source cluster.