Embedding Infinispan 12.0

Table of Contents

Chapter 1. Configuring the Infinispan Maven Repository

Infinispan Java distributions are available from Maven.

Infinispan artifacts are available from Maven central. See the org.infinispan group for available Infinispan artifacts.

1.1. Configuring Your Infinispan POM

Maven uses configuration files called Project Object Model (POM) files to define projects and manage builds. POM files are in XML format and describe the module and component dependencies, build order, and targets for the resulting project packaging and output.

Procedure

1. Open your project `pom.xml` for editing.
2. Define the `version.infinispan` property with the correct Infinispan version.
3. Include the `infinispan-bom` in a `dependencyManagement` section.

   The Bill Of Materials (BOM) controls dependency versions, which avoids version conflicts and means you do not need to set the version for each Infinispan artifact you add as a dependency to your project.

4. Save and close `pom.xml`.

The following example shows the Infinispan version and BOM:

```
<properties>
  <version.infinispan>12.0.0.Final</version.infinispan>
</properties>

<dependencyManagement>
  <dependencies>
    <dependency>
      <groupId>org.infinispan</groupId>
      <artifactId>infinispan-bom</artifactId>
      <version>${version.infinispan}</version>
      <type>pom</type>
      <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>
```

Next Steps

Add Infinispan artifacts as dependencies to your `pom.xml` as required.

Chapter 2. Installing Infinispan in Library Mode

Add Infinispan as an embedded library in your project.

Procedure

- Add the `infinispan-core` artifact as a dependency in your `pom.xml` as follows:

```
<dependencies>
  <dependency>
    <groupId>org.infinispan</groupId>
    <artifactId>infinispan-core</artifactId>
  </dependency>
</dependencies>
```

Chapter 3. Running Infinispan as an Embedded Library

Learn how to run Infinispan as an embedded data store in your project.

Procedure

- Initialize the default Cache Manager and add a cache definition as follows:

```
GlobalConfigurationBuilder global = GlobalConfigurationBuilder.
defaultClusteredBuilder();
DefaultCacheManager cacheManager = new DefaultCacheManager(global.build());
ConfigurationBuilder builder = new ConfigurationBuilder();
builder.clustering().cacheMode(CacheMode.DIST_SYNC);
cacheManager.administration().withFlags(CacheContainerAdmin.AdminFlag.VOLATILE).getOrC
reateCache("myCache", builder.build());
```

The preceding code initializes a default, clustered Cache Manager. Cache Managers contain your cache definitions and control cache lifecycles.

Infinispan does not provide default cache definitions so after initializing the default Cache Manager, you need to add at least one cache instance. This example uses the `ConfigurationBuilder` class to create a cache definition that uses the distributed, synchronous cache mode. You then call the `getOrCreateCache()` method that either creates a cache named "myCache" on all nodes in the cluster or returns it if it already exists.

Next steps

Now that you have a running Cache Manager with a cache created, you can add some more cache definitions, put some data into the cache, or configure Infinispan as needed.

Reference

- Configuring Infinispan Programmatically
- org.infinispan.configuration.global.GlobalConfigurationBuilder
- org.infinispan.manager.EmbeddedCacheManager
- org.infinispan.Cache

Chapter 4. Setting Up Infinispan Clusters

Infinispan requires a transport layer so nodes can automatically join and leave clusters. The transport layer also enables Infinispan nodes to replicate or distribute data across the network and perform operations such as re-balancing and state transfer.

4.1. Getting Started with Default Stacks

Infinispan uses JGroups protocol stacks so nodes can send each other messages on dedicated cluster channels.

Infinispan provides preconfigured JGroups stacks for `UDP` and `TCP` protocols. You can use these default stacks as a starting point for building custom cluster transport configuration that is optimized for your network requirements.

Procedure

1. Locate the default JGroups stacks, `default-jgroups-*.xml`, in the `default-configs` directory inside the `infinispan-core-12.0.0.Final.jar` file.
2. Do one of the following:
   - Use the `stack` attribute in your `infinispan.xml` file.

```
<infinispan>
  <cache-container default-cache="replicatedCache">
    <transport cluster="${infinispan.cluster.name}"
               stack="udp" ①
               node-name="${infinispan.node.name:}"/>
  </cache-container>
</infinispan>
```

① Uses `default-jgroups-udp.xml` for cluster transport.

   - Use the `addProperty()` method to set the JGroups stack file:

```
GlobalConfiguration globalConfig = new GlobalConfigurationBuilder().transport()
        .defaultTransport()
        .clusterName("qa-cluster")
        .addProperty("configurationFile", "default-jgroups-udp.xml") ①
        .build();
```

① Uses the `default-jgroups-udp.xml` stack for cluster transport.

Infinispan logs the following message to indicate which stack it uses:

```
[org.infinispan.CLUSTER] ISPN000078: Starting JGroups channel cluster with stack udp
```

4.1.1. Default JGroups Stacks

Learn about default JGroups stacks that configure cluster transport.

File name	Stack name	Description
`default-jgroups-udp.xml`	`udp`	Uses UDP for transport and UDP multicast for discovery. Suitable for larger clusters (over 100 nodes) or if you are using replicated caches or invalidation mode. Minimizes the number of open sockets.
`default-jgroups-tcp.xml`	`tcp`	Uses TCP for transport and the `MPING` protocol for discovery, which uses `UDP` multicast. Suitable for smaller clusters (under 100 nodes) *only if* you are using distributed caches because TCP is more efficient than UDP as a point-to-point protocol.
`default-jgroups-ec2.xml`	`ec2`	Uses TCP for transport and `S3_PING` for discovery. Suitable for Amazon EC2 nodes where UDP multicast is not available.
`default-jgroups-kubernetes.xml`	`kubernetes`	Uses TCP for transport and `DNS_PING` for discovery. Suitable for Kubernetes and Red Hat OpenShift nodes where UDP multicast is not always available.
`default-jgroups-google.xml`	`google`	Uses TCP for transport and `GOOGLE_PING2` for discovery. Suitable for Google Cloud Platform nodes where UDP multicast is not available.
`default-jgroups-azure.xml`	`azure`	Uses TCP for transport and `AZURE_PING` for discovery. Suitable for Microsoft Azure nodes where UDP multicast is not available.

Reference

- JGroups Protocols

4.1.2. TCP and UDP Ports for Cluster Traffic

Infinispan uses the following ports for cluster transport messages:

Default Port	Protocol	Description
`7800`	TCP/UDP	JGroups cluster bind port
`46655`	UDP	JGroups multicast

Cross-Site Replication

Infinispan uses the following ports for the JGroups RELAY2 protocol:

`7900`

For Infinispan clusters running on Kubernetes.

`7800`

If using UDP for traffic between nodes and TCP for traffic between clusters.

`7801`

If using TCP for traffic between nodes and TCP for traffic between clusters.

4.2. Customizing JGroups Stacks

Adjust and tune properties to create a cluster transport configuration that works for your network requirements.

Infinispan provides attributes that let you extend the default JGroups stacks for easier configuration. You can inherit properties from the default stacks while combining, removing, and replacing other properties.

Procedure

1. Create a new JGroups stack declaration in your `infinispan.xml` file.

```
<infinispan>
  <jgroups>
    <stack name="my-stack"> ①
    </stack>
  </jgroups>
</infinispan>
```

① Creates a custom JGroups stack named "my-stack".

2. Add the `extends` attribute and specify a JGroups stack to inherit properties from.

```
<infinispan>
  <jgroups>
    <stack name="my-stack" extends="tcp"> ①
    </stack>
  </jgroups>
</infinispan>
```

① Inherits from the default TCP stack.

3. Use the `stack.combine` attribute to modify properties for protocols configured in the inherited stack.
4. Use the `stack.position` attribute to define the location for your custom stack.

   For example, you might evaluate using a Gossip router and symmetric encryption with the default TCP stack as follows:

```
<jgroups>
  <stack name="my-stack" extends="tcp">
    <TCPGOSSIP initial_hosts=
"${jgroups.tunnel.gossip_router_hosts:localhost[12001]}"
               stack.combine="REPLACE"
               stack.position="MPING" /> ①
    <FD_SOCK stack.combine="REMOVE"/> ②
    <VERIFY_SUSPECT timeout="2000"/> ③
    <SYM_ENCRYPT sym_algorithm="AES"
                 keystore_name="mykeystore.p12"
                 keystore_type="PKCS12"
                 store_password="changeit"
                 key_password="changeit"
                 alias="myKey"
                 stack.combine="INSERT_AFTER"
                 stack.position="VERIFY_SUSPECT" /> ④
  </stack>
</jgroups>
```

① Uses the `TCPGOSSIP` protocol as the discovery mechanism instead of `MPING`.

② Removes the `FD_SOCK` protocol from the stack.

③ Modifies the timeout value for the `VERIFY_SUSPECT` protocol.

④ Adds the `SYM_ENCRYPT` protocol to the stack after the `VERIFY_SUSPECT` protocol.

5. Specify the stack name as the value for the `stack` attribute in the `transport` configuration.

```
<infinispan>
  <jgroups>
    <stack name="my-stack" extends="tcp">
     ...
    </stack>
    <cache-container name="default" statistics="true">
      <transport cluster="${infinispan.cluster.name}"
                 stack="my-stack" ①
                 node-name="${infinispan.node.name:}"/>
   </cache-container>
  </jgroups>
</infinispan>
```

① Configures Infinispan to use "my-stack" for cluster transport.

6. Check Infinispan logs to ensure it uses the stack.

```
[org.infinispan.CLUSTER] ISPN000078: Starting JGroups channel cluster with stack
my-stack
```

4.2.1. Inheritance Attributes

When you extend a JGroups stack, inheritance attributes let you adjust protocols and properties in the stack you are extending.

- `stack.position` specifies protocols to modify.
- `stack.combine` uses the following values to extend JGroups stacks:

Value	Description
`COMBINE`	Overrides protocol properties.
`REPLACE`	Replaces protocols.
`INSERT_AFTER`	Adds a protocol into the stack after another protocol. Does not affect the protocol that you specify as the insertion point. Protocols in JGroups stacks affect each other based on their location in the stack. For example, you should put a protocol such as `NAKACK2` after the `SYM_ENCRYPT` or `ASYM_ENCRYPT` protocol so that `NAKACK2` is secured.
`INSERT_BEFORE`	Inserts a protocols into the stack before another protocol. Affects the protocol that you specify as the insertion point.
`REMOVE`	Removes protocols from the stack.

4.3. Using JGroups System Properties

Pass system properties to Infinispan at startup to tune cluster transport.

Procedure

- Use `-D<property-name>=<property-value>` arguments to set JGroups system properties as required.

For example, set a custom bind port and IP address as follows:

```
$ java -cp ... -Djgroups.bind.port=1234 -Djgroups.bind.address=192.0.2.0
```

When you embed Infinispan clusters in clustered WildFly applications, JGroups system properties can clash or override each other.

For example, you do not set a unique bind address for either your Infinispan cluster or your WildFly application. In this case both Infinispan and your WildFly application use the JGroups default property and attempt to form clusters using the same bind address.

4.3.1. System Properties for JGroups Stacks

Set system properties that configure JGroups cluster transport stacks.

System Property	Description	Default Value	Required/Optional
`jgroups.bind.address`	Bind address for cluster transport.	`SITE_LOCAL`	Optional
`jgroups.bind.port`	Bind port for the socket.	`7800`	Optional
`jgroups.mcast_addr`	IP address for multicast, both discovery and inter-cluster communication. The IP address must be a valid "class D" address that is suitable for IP multicast.	`228.6.7.8`	Optional
`jgroups.mcast_port`	Port for the multicast socket.	`46655`	Optional
`jgroups.ip_ttl`	Time-to-live (TTL) for IP multicast packets. The value defines the number of network hops a packet can make before it is dropped.	2	Optional
`jgroups.thread_pool.min_threads`	Minimum number of threads for the thread pool.	0	Optional
`jgroups.thread_pool.max_threads`	Maximum number of threads for the thread pool.	200	Optional
`jgroups.join_timeout`	Maximum number of milliseconds to wait for join requests to succeed.	2000	Optional
`jgroups.thread_dumps_threshold`	Number of times a thread pool needs to be full before a thread dump is logged.	10000	Optional

Amazon EC3

The following system properties apply only to `default-jgroups-ec2.xml`:

System Property	Description	Default Value	Required/Optional
`jgroups.s3.access_key`	Amazon S3 access key for an S3 bucket.	No default value.	Optional
`jgroups.s3.secret_access_key`	Amazon S3 secret key used for an S3 bucket.	No default value.	Optional
`jgroups.s3.bucket`	Name of the Amazon S3 bucket. The name must exist and be unique.	No default value.	Optional

Kubernetes

The following system properties apply only to `default-jgroups-kubernetes.xml`:

System Property	Description	Default Value	Required/Optional
`jgroups.dns.query`	Sets the DNS record that returns cluster members.	No default value.	Required

Google Cloud Platform

The following system properties apply only to `default-jgroups-google.xml`:

System Property	Description	Default Value	Required/Optional
`jgroups.google.bucket_name`	Name of the Google Compute Engine bucket. The name must exist and be unique.	No default value.	Required

Reference

- JGroups System Properties
- JGroups Protocol List

4.4. Using Inline JGroups Stacks

You can insert complete JGroups stack definitions into `infinispan.xml` files.

Procedure

- Embed a custom JGroups stack declaration in your `infinispan.xml` file.

```
<infinispan>
  <jgroups> ①
    <stack name="prod"> ②
      <TCP bind_port="7800" port_range="30" recv_buf_size="20000000" send_buf_size
="640000"/>
      <MPING break_on_coord_rsp="true"
             mcast_addr="${jgroups.mping.mcast_addr:228.2.4.6}"
             mcast_port="${jgroups.mping.mcast_port:43366}"
             num_discovery_runs="3"
             ip_ttl="${jgroups.udp.ip_ttl:2}"/>
      <MERGE3 />
      <FD_SOCK />
      <FD_ALL timeout="3000" interval="1000" timeout_check_interval="1000" />
      <VERIFY_SUSPECT timeout="1000" />
      <pbcast.NAKACK2 use_mcast_xmit="false" xmit_interval="100"
xmit_table_num_rows="50"
                      xmit_table_msgs_per_row="1024"
xmit_table_max_compaction_time="30000" />
      <UNICAST3 xmit_interval="100" xmit_table_num_rows="50"
xmit_table_msgs_per_row="1024"
                xmit_table_max_compaction_time="30000" />
      <pbcast.STABLE stability_delay="200" desired_avg_gossip="2000" max_bytes="1M"
/>
      <pbcast.GMS print_local_addr="false" join_timeout=
"${jgroups.join_timeout:2000}" />
      <UFC max_credits="4m" min_threshold="0.40" />
      <MFC max_credits="4m" min_threshold="0.40" />
      <FRAG3 />
    </stack>
  </jgroups>
  <cache-container default-cache="replicatedCache">
    <transport stack="prod" /> ③
    ...
  </cache-container>
</infinispan>
```

① Contains one or more JGroups stack definitions.

② Defines a custom JGroups stack named "prod".

③ Configures Infinispan to use "prod" for cluster transport.

4.5. Using External JGroups Stacks

Reference external files that define custom JGroups stacks in `infinispan.xml` files.

Procedure

1. Put custom JGroups stack files on the application classpath.

   Alternatively you can specify an absolute path when you declare the external stack file.

2. Reference the external stack file with the `stack-file` element.

```
<infinispan>
  <jgroups>
     <stack-file name="prod-tcp" path="prod-jgroups-tcp.xml"/> ①
  </jgroups>
  <cache-container default-cache="replicatedCache">
    <transport stack="prod-tcp" /> ②
    <replicated-cache name="replicatedCache"/>
  </cache-container>
  ...
</infinispan>
```

① Creates a stack named "prod-tcp" that uses the "prod-jgroups-tcp.xml" definition.

② Configures Infinispan to use "prod-tcp" for cluster transport.

4.6. Cluster Discovery Protocols

Infinispan supports different protocols that allow nodes to automatically find each other on the network and form clusters.

There are two types of discovery mechanisms that Infinispan can use:

- Generic discovery protocols that work on most networks and do not rely on external services.
- Discovery protocols that rely on external services to store and retrieve topology information for Infinispan clusters.
  For instance the DNS_PING protocol performs discovery through DNS server records.

Running Infinispan on hosted platforms requires using discovery mechanisms that are adapted to network constraints that individual cloud providers impose.

Reference

- JGroups Discovery Protocols

4.6.1. PING

PING, or UDPPING is a generic JGroups discovery mechanism that uses dynamic multicasting with the UDP protocol.

When joining, nodes send PING requests to an IP multicast address to discover other nodes already in the Infinispan cluster. Each node responds to the PING request with a packet that contains the address of the coordinator node and its own address. C=coordinator's address and A=own address. If no nodes respond to the PING request, the joining node becomes the coordinator node in a new cluster.

PING configuration example

```
<config>
  <PING num_discovery_runs="3"/>
  ...
</config>
```

Reference

- JGroups PING

4.6.2. TCPPING

TCPPING is a generic JGroups discovery mechanism that uses a list of static addresses for cluster members.

With TCPPING, you manually specify the IP address or hostname of each node in the Infinispan cluster as part of the JGroups stack, rather than letting nodes discover each other dynamically.

TCPPING configuration example

```
<config>
  <TCP bind_port="7800" />
  <TCPPING timeout="3000"
           initial_hosts=
"${jgroups.tcpping.initial_hosts:hostname1[port1],hostname2[port2]}"
          port_range="0" ①
          num_initial_members="3"/>
  ...
</config>
```

① For reliable discovery, Red Hat recommends `port-range=0`.

Reference

- JGroups TCPPING

4.6.3. MPING

MPING uses IP multicast to discover the initial membership of Infinispan clusters.

You can use MPING to replace TCPPING discovery with TCP stacks and use multicasing for discovery instead of static lists of initial hosts. However, you can also use MPING with UDP stacks.

MPING configuration example

```
<config>
  <MPING mcast_addr="${jgroups.mcast_addr:228.6.7.8}"
         mcast_port="${jgroups.mcast_port:46655}"
         num_discovery_runs="3"
         ip_ttl="${jgroups.udp.ip_ttl:2}"/>
  ...
</config>
```

Reference

- JGroups MPING

4.6.4. TCPGOSSIP

Gossip routers provide a centralized location on the network from which your Infinispan cluster can retrieve addresses of other nodes.

You inject the address (`IP:PORT`) of the Gossip router into Infinispan nodes as follows:

1. Pass the address as a system property to the JVM; for example, `-DGossipRouterAddress="10.10.2.4[12001]"`.
2. Reference that system property in the JGroups configuration file.

Gossip router configuration example

```
<config>
  <TCP bind_port="7800" />
  <TCPGOSSIP timeout="3000"
             initial_hosts="${GossipRouterAddress}"
             num_initial_members="3" />
  ...
</config>
```

Reference

- JGroups Gossip Router

4.6.5. JDBC_PING

JDBC_PING uses shared databases to store information about Infinispan clusters. This protocol supports any database that can use a JDBC connection.

Nodes write their IP addresses to the shared database so joining nodes can find the Infinispan cluster on the network. When nodes leave Infinispan clusters, they delete their IP addresses from the shared database.

JDBC_PING configuration example

```
<config>
  <JDBC_PING connection_url="jdbc:mysql://localhost:3306/database_name"
             connection_username="user"
             connection_password="password"
             connection_driver="com.mysql.jdbc.Driver"/>
  ...
</config>
```

Add the appropriate JDBC driver to the classpath so Infinispan can use JDBC_PING.

Reference

- JDBC_PING
- JDBC_PING Wiki

4.6.6. DNS_PING

JGroups DNS_PING queries DNS servers to discover Infinispan cluster members in Kubernetes environments such as OKD and Red Hat OpenShift.

DNS_PING configuration example

```
<config>
  <dns.DNS_PING dns_query="myservice.myproject.svc.cluster.local" />
  ...
</config>
```

Reference

- JGroups DNS_PING
- DNS for Services and Pods (Kubernetes documentation for adding DNS entries)

4.7. Using Custom JChannels

Construct custom JGroups JChannels as in the following example:

```
GlobalConfigurationBuilder global = new GlobalConfigurationBuilder();
JChannel jchannel = new JChannel();
// Configure the jchannel as needed.
JGroupsTransport transport = new JGroupsTransport(jchannel);
global.transport().transport(transport);
new DefaultCacheManager(global.build());
```

Infinispan cannot use custom JChannels that are already connected.

Reference

JGroups JChannel

4.8. Encrypting Cluster Transport

Secure cluster transport so that nodes communicate with encrypted messages. You can also configure Infinispan clusters to perform certificate authentication so that only nodes with valid identities can join.

4.8.1. Infinispan Cluster Security

To secure cluster traffic, you configure Infinispan nodes to encrypt JGroups message payloads with secret keys.

Infinispan nodes can obtain secret keys from either:

- The coordinator node (asymmetric encryption).
- A shared keystore (symmetric encryption).

Retrieving secret keys from coordinator nodes

You configure asymmetric encryption by adding the `ASYM_ENCRYPT` protocol to a JGroups stack in your Infinispan configuration. This allows Infinispan clusters to generate and distribute secret keys.

When using asymmetric encryption, you should also provide keystores so that nodes can perform certificate authentication and securely exchange secret keys. This protects your cluster from man-in-the-middle (MitM) attacks.

Asymmetric encryption secures cluster traffic as follows:

1. The first node in the Infinispan cluster, the coordinator node, generates a secret key.
2. A joining node performs certificate authentication with the coordinator to mutually verify identity.
3. The joining node requests the secret key from the coordinator node. That request includes the public key for the joining node.
4. The coordinator node encrypts the secret key with the public key and returns it to the joining node.
5. The joining node decrypts and installs the secret key.
6. The node joins the cluster, encrypting and decrypting messages with the secret key.

Retrieving secret keys from shared keystores

You configure symmetric encryption by adding the `SYM_ENCRYPT` protocol to a JGroups stack in your Infinispan configuration. This allows Infinispan clusters to obtain secret keys from keystores that you provide.

1. Nodes install the secret key from a keystore on the Infinispan classpath at startup.
2. Node join clusters, encrypting and decrypting messages with the secret key.

Comparison of asymmetric and symmetric encryption

`ASYM_ENCRYPT` with certificate authentication provides an additional layer of encryption in comparison with `SYM_ENCRYPT`. You provide keystores that encrypt the requests to coordinator nodes for the secret key. Infinispan automatically generates that secret key and handles cluster traffic, while letting you specify when to generate secret keys. For example, you can configure clusters to generate new secret keys when nodes leave. This ensures that nodes cannot bypass certificate authentication and join with old keys.

`SYM_ENCRYPT`, on the other hand, is faster than `ASYM_ENCRYPT` because nodes do not need to exchange keys with the cluster coordinator. A potential drawback to `SYM_ENCRYPT` is that there is no configuration to automatically generate new secret keys when cluster membership changes. Users are responsible for generating and distributing the secret keys that nodes use to encrypt cluster traffic.

4.8.2. Configuring Cluster Transport with Asymmetric Encryption

Configure Infinispan clusters to generate and distribute secret keys that encrypt JGroups messages.

Procedure

1. Create a keystore with certificate chains that enables Infinispan to verify node identity.
2. Place the keystore on the classpath for each node in the cluster.

   For Infinispan Server, you put the keystore in the $ISPN_HOME directory.

3. Add the `SSL_KEY_EXCHANGE` and `ASYM_ENCRYPT` protocols to a JGroups stack in your Infinispan configuration, as in the following example:

```
<infinispan>
    <jgroups>
         <stack name="encrypt-tcp" extends="tcp"> ①
           <SSL_KEY_EXCHANGE keystore_name="mykeystore.jks" ②
                             keystore_password="changeit" ③
                             stack.combine="INSERT_AFTER"
                             stack.position="VERIFY_SUSPECT"/> ④
           <ASYM_ENCRYPT asym_keylength="2048" ⑤
                    asym_algorithm="RSA" ⑥
                    change_key_on_coord_leave = "false" ⑦
                    change_key_on_leave = "false" ⑧
                    use_external_key_exchange = "true" ⑨
                    stack.combine="INSERT_BEFORE"
                    stack.position="pbcast.NAKACK2"/> ⑩
         </stack>
    </jgroups>
    <cache-container name="default" statistics="true">
      <transport cluster="${infinispan.cluster.name}"
                 stack="encrypt-tcp" ⑪
                 node-name="${infinispan.node.name:}"/>
   </cache-container>
</infinispan>
```

① Creates a secure JGroups stack named "encrypt-tcp" that extends the default TCP stack for Infinispan.

② Names the keystore that nodes use to perform certificate authentication.

③ Specifies the keystore password.

④ Uses the `stack.combine` and `stack.position` attributes to insert `SSL_KEY_EXCHANGE` into the default TCP stack after the `VERIFY_SUSPECT` protocol.

⑤ Specifies the length of the secret key that the coordinator node generates. The default value is `2048`.

⑥ Specifies the cipher engine the coordinator node uses to generate secret keys. The default value is `RSA`.

⑦ Configures Infinispan to generate and distribute a new secret key when the coordinator node changes.

⑧ Configures Infinispan to generate and distribute a new secret key when nodes leave.

⑨ Configures Infinispan nodes to use the `SSL_KEY_EXCHANGE` protocol for certificate authentication.

⑩ Uses the `stack.combine` and `stack.position` attributes to insert `ASYM_ENCRYPT` into the default TCP stack before the `pbcast.NAKACK2` protocol.

⑪ Configures the Infinispan cluster to use the secure JGroups stack.

Verification

When you start your Infinispan cluster, the following log message indicates that the cluster is using

the secure JGroups stack:

```
[org.infinispan.CLUSTER] ISPN000078: Starting JGroups channel cluster with stack
<encrypted_stack_name>
```

Infinispan nodes can join the cluster only if they use `ASYM_ENCRYPT` and can obtain the secret key from the coordinator node. Otherwise the following message is written to Infinispan logs:

```
[org.jgroups.protocols.ASYM_ENCRYPT] <hostname>: received message without encrypt
header from <hostname>; dropping it
```

Reference

The example `ASYM_ENCRYPT` configuration in this procedure shows commonly used parameters. Refer to JGroups documentation for the full set of available parameters.

- JGroups 4 Manual
- JGroups 4.2 Schema

4.8.3. Configuring Cluster Transport with Symmetric Encryption

Configure Infinispan clusters to encrypt JGroups messages with secret keys from keystores that you provide.

Procedure

1. Create a keystore that contains a secret key.
2. Place the keystore on the classpath for each node in the cluster.

   For Infinispan Server, you put the keystore in the $ISPN_HOME directory.
3. Add the `SYM_ENCRYPT` protocol to a JGroups stack in your Infinispan configuration, as in the following example:

```
<infinispan>
    <jgroups>
         <stack name="encrypt-tcp" extends="tcp"> ①
           <SYM_ENCRYPT keystore_name="myKeystore.p12" ②
                        keystore_type="PKCS12" ③
                        store_password="changeit" ④
                        key_password="changeit" ⑤
                        alias="myKey" ⑥
                        stack.combine="INSERT_AFTER"
                        stack.position="VERIFY_SUSPECT"/> ⑦
         </stack>
    </jgroups>
    <cache-container name="default" statistics="true">
      <transport cluster="${infinispan.cluster.name}"
                 stack="encrypt-tcp" ⑧
                 node-name="${infinispan.node.name:}"/>
   </cache-container>
</infinispan>
```

① Creates a secure JGroups stack named "encrypt-tcp" that extends the default TCP stack for Infinispan.

② Names the keystore from which nodes obtain secret keys.

③ Specifies the keystore type. JGroups uses JCEKS by default.

④ Specifies the keystore password.

⑤ Specifies the secret key password.

⑥ Specifies the secret key alias.

⑦ Uses the `stack.combine` and `stack.position` attributes to insert `SYM_ENCRYPT` into the default TCP stack after the `VERIFY_SUSPECT` protocol.

⑧ Configures the Infinispan cluster to use the secure JGroups stack.

Verification

When you start your Infinispan cluster, the following log message indicates that the cluster is using the secure JGroups stack:

```
[org.infinispan.CLUSTER] ISPN000078: Starting JGroups channel cluster with stack
<encrypted_stack_name>
```

Infinispan nodes can join the cluster only if they use `SYM_ENCRYPT` and can obtain the secret key from the shared keystore. Otherwise the following message is written to Infinispan logs:

```
[org.jgroups.protocols.SYM_ENCRYPT] <hostname>: received message without encrypt
header from <hostname>; dropping it
```

Reference

The example `SYM_ENCRYPT` configuration in this procedure shows commonly used parameters. Refer to JGroups documentation for the full set of available parameters.

- JGroups 4 Manual
- JGroups 4.2 Schema