



Red Hat Data Grid 8.2

Data Grid Guide to Cross-Site Replication

Back up data across global Data Grid clusters

Red Hat Data Grid 8.2 Data Grid Guide to Cross-Site Replication

Back up data across global Data Grid clusters

Legal Notice

Copyright © 2021 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Data Grid can form global clusters to replicate your data across geographic locations. This guide explains how to configure backup locations for caches and perform cross-site operations. You can also learn how to monitor and troubleshoot global clusters.

Table of Contents

RED HAT DATA GRID	4
DATA GRID DOCUMENTATION	5
DATA GRID DOWNLOADS	6
MAKING OPEN SOURCE MORE INCLUSIVE	7
CHAPTER 1. DATA GRID CROSS-SITE REPLICATION	8
1.1. CROSS-SITE REPLICATION	8
1.1.1. Site Masters	8
1.2. ADDING BACKUPS TO CACHES	8
1.3. BACKUP STRATEGIES	9
Synchronous strategy	9
Asynchronous strategy	9
Synchronous vs asynchronous backups	10
1.4. AUTOMATICALLY TAKING BACKUPS OFFLINE	10
1.5. STATE TRANSFER	11
1.5.1. Automatic State Transfer	11
1.6. CLIENT CONNECTIONS ACROSS SITES	12
Active/Passive	12
Active/Active	12
Backup strategies	13
1.6.1. Concurrent Writes and Conflicting Entries	13
Conflict resolution algorithms	14
1.7. EXPIRATION WITH CROSS-SITE REPLICATION	15
CHAPTER 2. CONFIGURING DATA GRID FOR CROSS-SITE REPLICATION	16
2.1. CONFIGURING CLUSTER TRANSPORT FOR CROSS-SITE REPLICATION	16
2.1.1. JGroups RELAY2 Stacks	16
2.1.2. Custom JGroups RELAY2 Stacks	17
2.2. ADDING BACKUP LOCATIONS TO CACHES	18
2.3. BACKING UP TO CACHES WITH DIFFERENT NAMES	19
2.4. CONFIGURING CROSS-SITE STATE TRANSFER	19
2.5. CUSTOMIZING THE CONFLICT RESOLUTION ALGORITHM	20
2.6. VERIFYING CROSS-SITE VIEWS	21
2.7. CONFIGURING HOT ROD CLIENTS FOR CROSS-SITE REPLICATION	21
CHAPTER 3. PERFORMING CROSS-SITE REPLICATION OPERATIONS	23
3.1. PERFORMING CROSS-SITE OPERATIONS WITH THE CLI	23
3.1.1. Bringing Backup Locations Offline and Online	23
3.1.2. Configuring Cross-Site State Transfer Modes	23
3.1.3. Pushing State to Backup Locations	24
3.2. PERFORMING CROSS-SITE OPERATIONS WITH THE REST API	24
3.2.1. Getting Status of All Backup Locations	24
3.2.2. Getting Status of Specific Backup Locations	25
3.2.3. Taking Backup Locations Offline	25
3.2.4. Bringing Backup Locations Online	25
3.2.5. Pushing State to Backup Locations	26
3.2.6. Canceling State Transfer	26
3.2.7. Getting State Transfer Status	26
3.2.8. Clearing State Transfer Status	26
3.2.9. Modifying Take Offline Conditions	26

3.2.10. Canceling State Transfer from Receiving Sites	27
3.2.11. Getting Status of Backup Locations	27
3.2.12. Taking Backup Locations Offline	28
3.2.13. Bringing Backup Locations Online	28
3.2.14. Retrieving the State Transfer Mode	28
3.2.15. Setting the State Transfer Mode	28
3.2.16. Starting State Transfer	28
3.2.17. Canceling State Transfer	28
3.3. PERFORMING CROSS-SITE OPERATIONS WITH JMX	29
3.3.1. Configuring Data Grid to Register JMX MBeans	29
3.3.2. Performing Cross-Site Operations	29
CHAPTER 4. MONITORING AND TROUBLESHOOTING GLOBAL DATA GRID CLUSTERS	31
4.1. ENABLING DATA GRID STATISTICS	31
4.2. CONFIGURING DATA GRID METRICS	31
4.2.1. Collecting Data Grid Metrics	32
4.3. CONFIGURING DATA GRID TO REGISTER JMX MBEANS	32
4.3.1. JMX MBeans for Cross-Site Replication	33
4.4. COLLECTING LOGS AND TROUBLESHOOTING CROSS-SITE REPLICATION	33
4.4.1. Cross-Site Log Messages	34

RED HAT DATA GRID

Data Grid is a high-performance, distributed in-memory data store.

Schemaless data structure

Flexibility to store different objects as key-value pairs.

Grid-based data storage

Designed to distribute and replicate data across clusters.

Elastic scaling

Dynamically adjust the number of nodes to meet demand without service disruption.

Data interoperability

Store, retrieve, and query data in the grid from different endpoints.

DATA GRID DOCUMENTATION

Documentation for Data Grid is available on the Red Hat customer portal.

- [Data Grid 8.2 Documentation](#)
- [Data Grid 8.2 Component Details](#)
- [Supported Configurations for Data Grid 8.2](#)
- [Data Grid 8 Feature Support](#)
- [Data Grid Deprecated Features and Functionality](#)

DATA GRID DOWNLOADS

Access the [Data Grid Software Downloads](#) on the Red Hat customer portal.



NOTE

You must have a Red Hat account to access and download Data Grid software.

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).


CHAPTER 1. DATA GRID CROSS-SITE REPLICATION

Cross-site replication allows you to back up data from one Data Grid cluster to another. Learn the concepts to understand how Data Grid cross-site replication works before you configure your clusters.


1.1. CROSS-SITE REPLICATION

Data Grid clusters running in different locations can discover and communicate with each other.

Sites are typically data centers in various geographic locations. Cross-site replication bridges Data Grid clusters in sites to form global clusters, as in the following diagram:



LON is a datacenter in London, England.
NYC is a datacenter in New York City, USA.



NOTE

Data Grid can form global clusters across two or more sites.

For example, configure a third Data Grid cluster running in San Francisco, **SFO**, as backup location for **LON** and **NYC**.

1.1.1. Site Masters

Site masters are the nodes in Data Grid clusters that are responsible for sending and receiving requests from backup locations.

If a node is not a site master, it must forward backup requests to a local site master. Only site masters can send requests to backup locations.

For optimal performance, you should configure all nodes as site masters. This increases the speed of backup requests because each node in the cluster can backup to remote sites directly without having to forward backup requests to site masters.



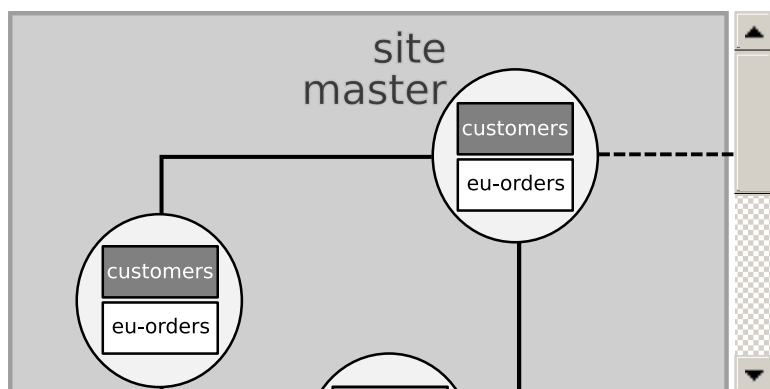
NOTE

Diagrams in this document illustrate Data Grid clusters with one site master because this is the default for the JGroups RELAY2 protocol. Likewise, a single site master is easier to illustrate because each site master in a cluster communicates with each site master in the remote cluster.

1.2. ADDING BACKUPS TO CACHES

Name remote sites as backup locations in your cache definitions.

For example, the following diagram shows three caches, "customers", "eu-orders", and "us-orders":



- In **LON**, "customers" names **NYC** as a backup location.
- In **NYC**, "customers" names **LON** as a backup location.
- "eu-orders" and "us-orders" do not have backups and are local to the respective cluster.

1.3. BACKUP STRATEGIES

Data Grid replicates data between clusters at the same time that writes to caches occur. For example, if a client writes "k1" to **LON**, Data Grid backs up "k1" to **NYC** at the same time.

To back up data to a different cluster, Data Grid can use either a synchronous or asynchronous strategy.

Synchronous strategy

When Data Grid replicates data to backup locations, it writes to the cache on the local cluster and the cache on the remote cluster concurrently. With the synchronous strategy, Data Grid waits for both write operations to complete before returning.

You can control how Data Grid handles writes to the cache on the local cluster if backup operations fail. Data Grid can do the following:

- Ignore the failed backup and silently continue the write to the local cluster.
- Log a warning message or throw an exception and continue the write to the local cluster.
- Handle failed backup operations with custom logic.

Synchronous backups also support two-phase commits with caches that participate in optimistic transactions. The first phase of the backup acquires a lock. The second phase commits the modification.



IMPORTANT

Two-phase commit with cross-site replication has a significant performance impact because it requires two round-trips across the network.

Asynchronous strategy

When Data Grid replicates data to backup locations, it does not wait until the operation completes before writing to the local cache.

Asynchronous backup operations and writes to the local cache are independent of each other. If backup operations fail, write operations to the local cache continue and no exceptions occur. When this happens Data Grid also retries the write operation until the remote cluster disconnects from the cross-site view.

Synchronous vs asynchronous backups

Synchronous backups offer the strongest guarantee of data consistency across sites. If **strategy=sync**, when **cache.put()** calls return you know the value is up to date in the local cache and in the backup locations.

The trade-off for this consistency is performance. Synchronous backups have much greater latency in comparison to asynchronous backups.

Asynchronous backups, on the other hand, do not add latency to client requests so they have no performance impact. However, if **strategy=async**, when **cache.put()** calls return you cannot be sure of that the value in the backup location is the same as in the local cache.

1.4. AUTOMATICALLY TAKING BACKUPS OFFLINE

You can configure backup locations to go offline automatically when the remote sites become unavailable. This prevents Data Grid nodes from continuously attempting to replicate data to offline backup locations, which results in error messages and consumes resources.

Timeout for backup operations

Backup configurations include timeout values for operations to replicate data. If operations do not complete before the timeout occurs, Data Grid records them as failures.

```
<!-- Operations to replicate data to NYC are recorded as failures if they do not complete after 10 seconds. -->
<backup site="NYC" strategy="ASYNC" timeout="10000">
  <!-- Configuration for backup locations goes here. -->
</backup>
```

Number of failures

You can specify the number of **consecutive** failures that can occur before backup locations go offline.

For example, the following configuration for **NYC** sets five as the number of failed operations before it goes offline:

```
<!-- If a cluster attempts to replicate data to NYC and five consecutive operations fail, NYC automatically goes offline. -->
<backup site="NYC" strategy="ASYNC" timeout="10000">
  <take-offline after-failures="5"/>
</backup>
```

Time to wait

You can also specify how long to wait before taking sites offline when backup operations fail. If a backup request succeeds before the wait time runs out, Data Grid does not take the site offline.

```
<!-- If a cluster attempts to replicate data to NYC and there are more than five consecutive failures and 15 seconds elapse after the first failed operation, NYC automatically goes offline. -->
<backup site="NYC" strategy="ASYNC" timeout="10000">
```

```
<take-offline after-failures="5"
  min-wait="15000"/>
</backup>
```

TIP

Set a negative or zero value for the **after-failures** attribute if you want to use only a minimum time to wait to take sites offline.

```
<take-offline after-failures="-1" min-wait="10000"/>
```

1.5. STATE TRANSFER

State transfer is an administrative operation that synchronizes data between sites.

For example, **LON** goes offline and **NYC** starts handling client requests. When you bring **LON** back online, the Data Grid cluster in **LON** does not have the same data as the cluster in **NYC**.

To ensure the data is consistent between **LON** and **NYC**, you can push state from **NYC** to **LON**.

- State transfer is bidirectional. For example, you can push state from **NYC** to **LON** or from **LON** to **NYC**.
- Pushing state to offline sites brings them back online.
- State transfer overwrites only data that exists on both sites, the originating site and the receiving site. Data Grid does not delete data. For example, "k2" exists on **LON** and **NYC**. "k2" is removed from **NYC** while **LON** is offline. When you bring **LON** back online, "k2" still exists at that location. If you push state from **NYC** to **LON**, the transfer does not affect "k2" on **LON**.

TIP

To ensure contents of the cache are identical after state transfer, remove all data from the cache on the receiving site before pushing state.

Use the **clear()** method or the **clearcache** command from the CLI.

- State transfer does not overwrite updates to data that occur after you initiate the push. For example, "k1,v1" exists on **LON** and **NYC**. **LON** goes offline so you push state transfer to **LON** from **NYC**, which brings **LON** back online. Before state transfer completes, a client puts "k1,v2" on **LON**.

In this case the state transfer from **NYC** does not overwrite "k1,v2" because that modification happened after you initiated the push.

1.5.1. Automatic State Transfer

By default you must manually perform cross-site state transfer operations with the CLI or via JMX or REST.

However, when using the asynchronous backup strategy, Data Grid can automatically perform cross-site state transfer operations. When it detects that a backup location has come back online, and the network connection is stable, Data Grid initiates bi-directional state transfer between backup locations. For

example, Data Grid simultaneously transfers state from **LON** to **NYC** and **NYC** to **LON**.



NOTE

To avoid temporary network disconnects triggering state transfer operations, there are two conditions that backup locations must meet to go offline. The status of a backup location must be offline and it must not be included in the cross-site view with JGroups RELAY2.

Reference

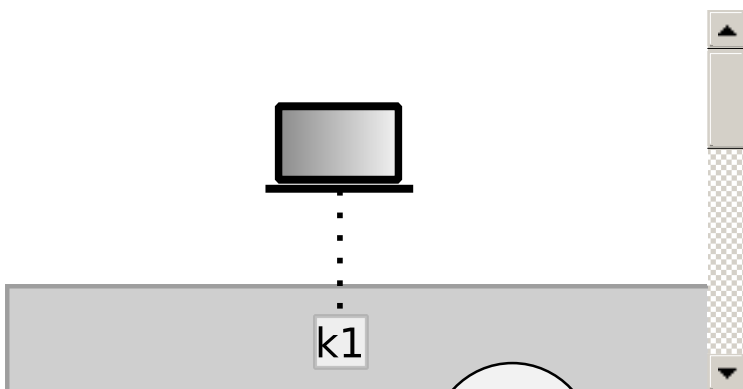
- [org.infinispan.Cache.clear\(\)](#)
- [Clearing Caches with the CLI](#)
- [Clearing Caches with the REST API](#)
- [Configure Cross-Site State Transfer](#)

1.6. CLIENT CONNECTIONS ACROSS SITES

Clients can write to Data Grid clusters in either an Active/Passive or Active/Active configuration.

Active/Passive

The following diagram illustrates Active/Passive where Data Grid handles client requests from one site only:




In the preceding image:

1. Client connects to the Data Grid cluster at **LON**.
2. Client writes "k1" to the cache.
3. The site master at **LON**, "n1", sends the request to replicate "k1" to the site master at **NYC**, "nA".

With Active/Passive, **NYC** provides data redundancy. If the Data Grid cluster at **LON** goes offline for any reason, clients can start sending requests to **NYC**. When you bring **LON** back online you can synchronize data with **NYC** and then switch clients back to **LON**.

Active/Active

The following diagram illustrates Active/Active where Data Grid handles client requests at two sites:



In the preceding image:

1. Client A connects to the Data Grid cluster at **LON**.
2. Client A writes "k1" to the cache.
3. Client B connects to the Data Grid cluster at **NYC**.
4. Client B writes "k2" to the cache.
5. Site masters at **LON** and **NYC** send requests so that "k1" is replicated to **NYC** and "k2" is replicated to **LON**.

With Active/Active both **NYC** and **LON** replicate data to remote caches while handling client requests. If either **NYC** or **LON** go offline, clients can start sending requests to the online site. You can then bring offline sites back online, push state to synchronize data, and switch clients as required.

Backup strategies



IMPORTANT

With an Active/Active configuration, you should always use an asynchronous backup strategy (**strategy=async**).

If multiple clients attempt to write to the same entry concurrently, and the backup strategy is synchronous (**strategy=sync**), then deadlocks occur. However you can use the synchronous backup strategy with an Active/Passive configuration if both sites access different data sets, in which case there is no risk of deadlocks from concurrent writes.

1.6.1. Concurrent Writes and Conflicting Entries

Conflicting entries can occur with Active/Active site configurations if clients write to the same entries at the same time but at different sites.

For example, client A writes to "k1" in **LON** at the same time that client B writes to "k1" in **NYC**. In this case, "k1" has a different value in **LON** than in **NYC**. After replication occurs, there is no guarantee which value for "k1" exists at which site.

To ensure data consistency, Data Grid uses a vector clock algorithm to detect conflicting entries during backup operations, as in the following illustration:

LON NYC

```

k1=(n/a)  0,0   0,0
k1=2      1,0 --> 1,0 k1=2
k1=3      1,1 <-- 1,1 k1=3
k1=5      2,1   1,2 k1=8
           --> 2,1 (conflict)
(conflict) 1,2 <--

```

Vector clocks are timestamp metadata that increment with each write to an entry. In the preceding example, **0,0** represents the initial value for the vector clock on "k1".

A client puts "k1=2" in **LON** and the vector clock is **1,0**, which Data Grid replicates to **NYC**. A client then puts "k1=3" in **NYC** and the vector clock updates to **1,1**, which Data Grid replicates to **LON**.

However if a client puts "k1=5" in **LON** at the same time that a client puts "k1=8" in **NYC**, Data Grid detects a conflicting entry because the vector value for "k1" is not strictly greater or less between **LON** and **NYC**.

When it finds conflicting entries, Data Grid uses the Java `compareTo(String anotherString)` method to compare site names. To determine which key takes priority, Data Grid selects the site name that is lexicographically less than the other. Keys from a site named **AAA** take priority over keys from a site named **AAB** and so on.

Following the same example, to resolve the conflict for "k1", Data Grid uses the value for "k1" that originates from **LON**. This results in "k1=5" in both **LON** and **NYC** after Data Grid resolves the conflict and replicates the value.

TIP

Prepend site names with numbers as a simple way to represent the order of priority for resolving conflicting entries; for example, **1LON** and **2NYC**.



NOTE

Data Grid performs conflict resolution with the asynchronous backup strategy (`strategy=async`) only. However, you should not use the synchronous backup strategy with an Active/Active configuration because concurrent writes result in deadlocks.

Conflict resolution algorithms

Data Grid provides different algorithms for resolving conflicts in addition to an `XSiteEntryMergePolicy` SPI that lets you implement custom conflict resolution strategies.

Apart from the default conflict resolution strategy, which uses lexicographical comparison, you can use Data Grid conflict resolution algorithms to:

- Always remove conflicting entries.
- Keep write operations when write/remove conflicts occur.
- Remove entries when write/remove conflicts occur.

Find all available algorithms and their descriptions in the `org.infinispan.xsite.spi.XSiteMergePolicy` enum.

Additional resources

- [java.lang.String#compareTo\(\)](#)
- [org.infinispan.xsite.spi.XSiteMergePolicy](#)
- [org.infinispan.xsite.spi.XSiteEntryMergePolicy](#)
- [Customizing the Conflict Resolution Strategy](#)

1.7. EXPIRATION WITH CROSS-SITE REPLICATION

Expiration removes cache entries based on time. Data Grid provides two ways to configure expiration for entries:

lifespan

Sets the maximum amount of time that entries can exist.

When you set **lifespan** with cross-site replication, Data Grid clusters expire entries independently of remote sites.

max-idle

Specifies how long entries can exist based on read or write operations in a given time period.

When you set a **max-idle** with cross-site replication, Data Grid clusters send touch commands to coordinate idle timeout values with remote sites.



NOTE

Using maximum idle expiration in cross-site deployments can impact performance because the additional processing to keep **max-idle** values synchronized means some operations take longer to complete.

CHAPTER 2. CONFIGURING DATA GRID FOR CROSS-SITE REPLICATION

Configuring Data Grid to replicate data across sites, you first set up cluster transport so Data Grid clusters can discover each other and site masters can communicate. You then add backup locations to cache definitions in your Data Grid configuration.

2.1. CONFIGURING CLUSTER TRANSPORT FOR CROSS-SITE REPLICATION

Add JGroups RELAY2 to your transport layer so that Data Grid clusters can communicate with backup locations.

Procedure

1. Open **infinispan.xml** for editing.
2. Add the RELAY2 protocol to a JGroups stack.
3. Configure Data Grid cluster transport to use the stack.

```
<infinispan>
  <jgroups>
    <!-- Extends the default UDP stack. -->
    <stack name="xsite" extends="udp">
      <!-- Adds RELAY2 for cross-site replication. -->
      <!-- Names the local site as LON. -->
      <!-- Specifies 1000 nodes as the maximum number of site masters. -->
      <relay.RELAY2 site="LON" xmlns="urn:org:jgroups" max_site_masters="1000"/>
      <!-- Uses the default TCP stack for inter-cluster communication. -->
      <!-- Names all sites that act as backup locations. -->
      <remote-sites default-stack="tcp">
        <remote-site name="LON"/>
        <remote-site name="NYC"/>
      </remote-sites>
    </stack>
  </jgroups>
  <cache-container name="default" statistics="true">
    <!-- Use the "xsite" stack for cluster transport. -->
    <transport cluster="{cluster.name}" stack="xsite"/>
  </cache-container>
</infinispan>
```

4. Save and close **infinispan.xml**.

Reference

- [JGroups RELAY2 Stacks](#)
- [Data Grid Configuration Schema](#)

2.1.1. JGroups RELAY2 Stacks

Data Grid clusters use JGroups RELAY2 for inter-cluster discovery and communication.

```
<jgroups>
  <!-- Uses the default JGroups UDP stack for intra-cluster traffic. -->
  <stack name="xsite" extends="udp">
    <!-- Adds RELAY2 to the stack for inter-cluster transport. -->
    <!-- Names the local site. Data in caches from the local site is replicated to backup locations. -->
    <!-- Configures a maximum of 1000 site masters for the local cluster. -->
    <relay.RELAY2 xmlns="urn:org:jgroups"
      site="LON"
      max_site_masters="1000"/>
    <!-- Specifies all site names and uses the default JGroups TCP stack for inter-cluster transport. -->
  >
  <remote-sites default-stack="tcp">
    <!-- Names all sites that participate in cross-site replication. -->
    <remote-site name="LON"/>
    <remote-site name="NYC"/>
  </remote-sites>
</stack>
</jgroups>
```

TIP

Set **max_site_masters** => the number of nodes in the Data Grid cluster for optimal performance with backup requests.

2.1.2. Custom JGroups RELAY2 Stacks

The following configuration adds a custom RELAY2 stack that extends the default TCP stack:

```
<jgroups>
  <!-- Uses TCPPING instead of MPING for discovery. -->
  <stack name="relay-global" extends="tcp">
    <MPING stack.combine="REMOVE"/>
    <TCPPING initial_hosts="192.0.2.0[7800]"
      stack.combine="INSERT_AFTER"
      stack.position="TCP"/>
  </stack>
  <!-- Extends the default UDP stack with RELAY2. -->
  <!-- Specifies RELAY2 properties. -->
  <stack name="xsite" extends="udp">
    <relay.RELAY2 site="LON" xmlns="urn:org:jgroups"
      max_site_masters="10"
      can_become_site_master="true"/>
    <remote-sites default-stack="relay-global">
      <remote-site name="LON"/>
      <remote-site name="NYC"/>
    </remote-sites>
  </stack>
</jgroups>
```

You can also reference externally defined JGroups stack files, for example:

```
<infinispan>
```

```

<jgroups>
  <stack-file name="relay-global" path="jgroups-relay.xml"/>
</jgroups>
<cache-container default-cache="replicatedCache">
  <transport stack="relay-global" />
  <replicated-cache name="replicatedCache"/>
</cache-container>
</infinispan>

```

Where **jgroups-relay.xml** references a JGroups stack file such as:

```

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

  <!-- Use TCP for inter-cluster transport. -->
  <TCP bind_addr="127.0.0.1"
    bind_port="7200"
    port_range="30"

    thread_pool.min_threads="0"
    thread_pool.max_threads="8"
    thread_pool.keep_alive_time="5000"
  />

  <!-- Use TCPPING for inter-cluster discovery. -->
  <TCPPING timeout="3000"
    initial_hosts="127.0.0.1[7200]"
    port_range="3"
    ergonomics="false"/>

  <!-- Provide other JGroups stack configuration as required. -->
</config>

```

Reference


- [Setting Up Data Grid Clusters](#)
- [JGroups RELAY2](#)
- [Relaying between multiple sites \(RELAY2\)](#)

2.2. ADDING BACKUP LOCATIONS TO CACHES

Specify the names of remote sites so Data Grid can back up data to those locations.

Procedure

1. Add the **backups** element to your cache definition.
2. Specify the name of each remote site with the **backup** element.
As an example, in the **LON** configuration, specify **NYC** as the remote site.
3. Repeat the preceding steps so that each site is a backup for all other sites. For example, you cannot add **LON** as a backup for **NYC** without adding **NYC** as a backup for **LON**.

**NOTE**

Cache configurations can be different across sites and use different backup strategies. Data Grid replicates data based on cache names.

Example "customers" configuration in LON

```
<replicated-cache name="customers">
  <backups>
    <backup site="NYC" strategy="ASYNC" />
  </backups>
</replicated-cache>
```

Example "customers" configuration in NYC

```
<distributed-cache name="customers">
  <backups>
    <backup site="LON" strategy="SYNC" />
  </backups>
</distributed-cache>
```

Reference

- [Data Grid Configuration Schema](#)

2.3. BACKING UP TO CACHES WITH DIFFERENT NAMES

By default, Data Grid replicates data between caches that have the same name.

Procedure

- Use **backup-for** to replicate data from a remote site into a cache with a different name on the local site.

For example, the following configuration backs up the "customers" cache on **LON** to the "eu-customers" cache on **NYC**.

```
<distributed-cache name="eu-customers">
  <backups>
    <backup site="LON" strategy="SYNC" />
  </backups>
  <backup-for remote-cache="customers" remote-site="LON" />
</distributed-cache>
```

2.4. CONFIGURING CROSS-SITE STATE TRANSFER

Change cross-site state transfer settings to optimize performance and specify whether operations happen manually or automatically.

Procedure

- Use **<state-transfer>** to configure state transfer operations.

For example, the following configuration automatically performs cross-site state transfer with custom settings:

```
<distributed-cache name="eu-customers">
  <backups>
    <backup site="LON" strategy="ASYNC">
      <state-transfer chunk-size="64"
        timeout="30000"
        max-retries="30"
        wait-time="2000"
        mode="AUTO"/>
    </backup>
  </backups>
</distributed-cache>
```

2.5. CUSTOMIZING THE CONFLICT RESOLUTION ALGORITHM

Configure Data Grid to use a different algorithm to resolve conflicting entries between backup locations.

Procedure

1. Open your Data Grid configuration for editing.
2. Specify one of the Data Grid algorithms or a custom implementation with the **merge-policy** attribute for the **backups** element.

Using Data Grid algorithms

Find all Data Grid algorithms and their descriptions in the **org.infinispan.xsite.spi.XSiteMergePolicy** enum.

The following example configuration uses the **ALWAYS_REMOVE** algorithm that deletes conflicting entries from both sites:

```
<distributed-cache name="eu-customers">
  <backups merge-policy="ALWAYS_REMOVE">
    <backup site="LON" strategy="ASYNC"/>
  </backups>
</distributed-cache>
```

Using custom implementations

1. Create a custom **XSiteEntryMergePolicy** implementation.

```
public interface XSiteEntryMergePolicy<K, V> {
    CompletionStage<SiteEntry<V>> merge(K key, SiteEntry<V> localEntry, SiteEntry<V>
remoteEntry);
}
```

2. Specify the fully qualified class name as the value of the **merge-policy** attribute, as in the following example:

```
<distributed-cache name="eu-customers">
  <backups merge-policy="org.mycompany.MyCustomXSiteEntryMergePolicy">
```



```

    <backup site="LON" strategy="ASYNC"/>
  </backups>
</distributed-cache>

```

Additional resources

- [org.infinispan.xsite.spi.XSiteEntryMergePolicy](#)
- [org.infinispan.xsite.spi.XSiteMergePolicy](#)
- [org.infinispan.xsite.spi.SiteEntry](#)
- [Data Grid Configuration Schema](#)

2.6. VERIFYING CROSS-SITE VIEWS

After you configure Data Grid for cross-site replication, you should verify that Data Grid clusters successfully form cross-site views.

Procedure

- Check log messages for **ISPN000439: Received new x-site view** messages.

For example, if the Data Grid cluster in **LON** has formed a cross-site view with the Data Grid cluster in **NYC**, it provides the following messages:

```

INFO [org.infinispan.XSITE] (jgroups-5,${server.hostname}) ISPN000439: Received new x-site view:
[NYC]
INFO [org.infinispan.XSITE] (jgroups-7,${server.hostname}) ISPN000439: Received new x-site view:
[NYC, LON]

```

2.7. CONFIGURING HOT ROD CLIENTS FOR CROSS-SITE REPLICATION

Configure Hot Rod clients to use Data Grid clusters at different sites.

hotrod-client.properties

```

# Servers at the active site
infinispan.client.hotrod.server_list = LON_host1:11222,LON_host2:11222,LON_host3:11222

# Servers at the backup site
infinispan.client.hotrod.cluster.NYC =
NYC_hostA:11222,NYC_hostB:11222,NYC_hostC:11222,NYC_hostD:11222

```

ConfigurationBuilder

```

ConfigurationBuilder builder = new ConfigurationBuilder();
builder.addServers("LON_host1:11222;LON_host2:11222;LON_host3:11222")
    .addCluster("NYC")

    .addClusterNodes("NYC_hostA:11222;NYC_hostB:11222;NYC_hostC:11222;NYC_hostD:11222")

```

TIP

Use the following methods to switch Hot Rod clients to the default cluster or to a cluster at a different site:

- **`RemoteCacheManager.switchToDefaultCluster()`**
- **`RemoteCacheManager.switchToCluster(${site.name})`**

Reference

- [org.infinispan.client.hotrod.configuration package description](#)
- [org.infinispan.client.hotrod.configuration.ConfigurationBuilder](#)
- [org.infinispan.client.hotrod.RemoteCacheManager](#)

CHAPTER 3. PERFORMING CROSS-SITE REPLICATION OPERATIONS

Bring sites online and offline. Transfer cache state to remote sites.

3.1. PERFORMING CROSS-SITE OPERATIONS WITH THE CLI

The Data Grid command line interface lets you remotely connect to Data Grid servers, manage sites, and push state transfer to backup locations.

Prerequisites

- Start the Data Grid CLI.
- Connect to a running Data Grid cluster.

3.1.1. Bringing Backup Locations Offline and Online

Take backup locations offline manually and bring them back online.

Procedure

1. Create a CLI connection to Data Grid.
2. Check if backup locations are online or offline with the **site status** command:

```
[//containers/default]> site status --cache=cacheName --site=NYC
```



NOTE

--site is an optional argument. If not set, the CLI returns all backup locations.

3. Manage backup locations as follows:
 - Bring backup locations online with the **bring-online** command:

```
[//containers/default]> site bring-online --cache=customers --site=NYC
```

- Take backup locations offline with the **take-offline** command:

```
[//containers/default]> site take-offline --cache=customers --site=NYC
```

For more information and examples, run the **help site** command.

3.1.2. Configuring Cross-Site State Transfer Modes

You can configure cross-site state transfer operations to happen automatically when Data Grid detects that backup locations come online. Alternatively you can use the default mode, which is to manually perform state transfer through the CLI or via JMX or REST.

Procedure

1. Create a CLI connection to Data Grid.
2. Use the **site** command to configure state transfer modes, as in the following examples:

- Retrieve the current state transfer mode.

```
[//containers/default]> site state-transfer-mode get --cache=cacheName --site=NYC
"MANUAL"
```

- Configure automatic state transfer operations for a cache and backup location.

```
[//containers/default]> site state-transfer-mode set --cache=cacheName --site=NYC --mode=AUTO
```

TIP

Run the **help site** command for more information and examples.

3.1.3. Pushing State to Backup Locations

Transfer cache state to remote backup locations.

Procedure

1. Create a CLI connection to Data Grid.
2. Use the **site** command to push state transfer, as in the following example:

```
[//containers/default]> site push-site-state --cache=cacheName --site=NYC
```

For more information and examples, run the **help site** command.

Reference

[Data Grid Command Line Interface](#)

3.2. PERFORMING CROSS-SITE OPERATIONS WITH THE REST API

Data Grid servers provide a REST API that allows you to perform cross-site operations.

3.2.1. Getting Status of All Backup Locations

Retrieve the status of all backup locations with **GET** requests.

```
GET /v2/caches/{cacheName}/x-site/backups/
```

Data Grid responds with the status of each backup location in JSON format, as in the following example:

```
{
  "NYC": "online",
  "LON": "offline"
}
```

Table 3.1. Returned Status

Value	Description
online	All nodes in the local cluster have a cross-site view with the backup location.
offline	No nodes in the local cluster have a cross-site view with the backup location.
mixed	Some nodes in the local cluster have a cross-site view with the backup location, other nodes in the local cluster do not have a cross-site view. The response indicates status for each node.

3.2.2. Getting Status of Specific Backup Locations

Retrieve the status of a backup location with **GET** requests.

```
GET /v2/caches/{cacheName}/x-site/backups/{siteName}
```

Data Grid responds with the status of each node in the site in JSON format, as in the following example:

```
{
  "NodeA":"offline",
  "NodeB":"online"
}
```

Table 3.2. Returned Status

Value	Description
online	The node is online.
offline	The node is offline.
failed	Not possible to retrieve status. The remote cache could be shutting down or a network error occurred during the request.

3.2.3. Taking Backup Locations Offline

Take backup locations offline with **POST** requests and the **?action=take-offline** parameter.

```
POST /v2/caches/{cacheName}/x-site/backups/{siteName}?action=take-offline
```

3.2.4. Bringing Backup Locations Online

Bring backup locations online with the **?action=bring-online** parameter.

```
POST /v2/caches/{cacheName}/x-site/backups/{siteName}?action=bring-online
```

3.2.5. Pushing State to Backup Locations

Push cache state to a backup location with the **?action=start-push-state** parameter.

```
POST /v2/caches/{cacheName}/x-site/backups/{siteName}?action=start-push-state
```

3.2.6. Canceling State Transfer

Cancel state transfer operations with the **?action=cancel-push-state** parameter.

```
POST /v2/caches/{cacheName}/x-site/backups/{siteName}?action=cancel-push-state
```

3.2.7. Getting State Transfer Status

Retrieve status of state transfer operations with the **?action=push-state-status** parameter.

```
GET /v2/caches/{cacheName}/x-site/backups?action=push-state-status
```

Data Grid responds with the status of state transfer for each backup location in JSON format, as in the following example:

```
{
  "NYC":"CANCELED",
  "LON":"OK"
}
```

Table 3.3. Returned Status

Value	Description
SENDING	State transfer to the backup location is in progress.
OK	State transfer completed successfully.
ERROR	An error occurred with state transfer. Check log files.
CANCELLING	State transfer cancellation is in progress.

3.2.8. Clearing State Transfer Status

Clear state transfer status for sending sites with the **?action=clear-push-state-status** parameter.

```
POST /v2/caches/{cacheName}/x-site/local?action=clear-push-state-status
```

3.2.9. Modifying Take Offline Conditions

Sites go offline if certain conditions are met. Modify the take offline parameters to control when backup locations automatically go offline.

Procedure

1. Check configured take offline parameters with **GET** requests and the **take-offline-config** parameter.

```
GET /v2/caches/{cacheName}/x-site/backups/{siteName}/take-offline-config
```

The Data Grid response includes **after_failures** and **min_wait** fields as follows:

```
{
  "after_failures": 2,
  "min_wait": 1000
}
```

2. Modify take offline parameters in the body of **PUT** requests.

```
PUT /v2/caches/{cacheName}/x-site/backups/{siteName}/take-offline-config
```

If the operation successfully completes, the service returns **204 (No Content)**.

3.2.10. Canceling State Transfer from Receiving Sites

If the connection between two backup locations breaks, you can cancel state transfer on the site that is receiving the push.

Cancel state transfer from a remote site and keep the current state of the local cache with the **?action=cancel-receive-state** parameter.

```
POST /v2/caches/{cacheName}/x-site/backups/{siteName}?action=cancel-receive-state
```

3.2.11. Getting Status of Backup Locations

Retrieve the status of all backup locations from Cache Managers with **GET** requests.

```
GET /rest/v2/cache-managers/{cacheManagerName}/x-site/backups/
```

Data Grid responds with status in JSON format, as in the following example:

```
{
  "SFO-3":{
    "status":"online"
  },
  "NYC-2":{
    "status":"mixed",
    "online":[
      "CACHE_1"
    ],
    "offline":[
      "CACHE_2"
    ]
  }
}
```

Table 3.4. Returned Status

Value	Description
online	All nodes in the local cluster have a cross-site view with the backup location.
offline	No nodes in the local cluster have a cross-site view with the backup location.
mixed	Some nodes in the local cluster have a cross-site view with the backup location, other nodes in the local cluster do not have a cross-site view. The response indicates status for each node.

3.2.12. Taking Backup Locations Offline

Take backup locations offline with the **?action=take-offline** parameter.

```
POST /rest/v2/cache-managers/{cacheManagerName}/x-site/backups/{siteName}?action=take-offline
```

3.2.13. Bringing Backup Locations Online

Bring backup locations online with the **?action=bring-online** parameter.

```
POST /rest/v2/cache-managers/{cacheManagerName}/x-site/backups/{siteName}?action=bring-online
```

3.2.14. Retrieving the State Transfer Mode

Check the state transfer mode with **GET** requests.

```
GET /rest/v2/caches/{cacheName}/x-site/backups/{site}/state-transfer-mode
```

3.2.15. Setting the State Transfer Mode

Configure the state transfer mode with the **?action=set** parameter.

```
POST /rest/v2/caches/{cacheName}/x-site/backups/{site}/state-transfer-mode?action=set&mode={mode}
```

3.2.16. Starting State Transfer

Push state of all caches to remote sites with the **?action=start-push-state** parameter.

```
POST /rest/v2/cache-managers/{cacheManagerName}/x-site/backups/{siteName}?action=start-push-state
```

3.2.17. Canceling State Transfer

Cancel ongoing state transfer operations with the **?action=cancel-push-state** parameter.

```
POST /rest/v2/cache-managers/{cacheManagerName}/x-site/backups/{siteName}?action=cancel-push-state
```

3.3. PERFORMING CROSS-SITE OPERATIONS WITH JMX

Data Grid provides JMX tooling to perform cross-site operations such as pushing state transfer and bringing sites online.

3.3.1. Configuring Data Grid to Register JMX MBeans

Data Grid can register JMX MBeans that you can use to collect statistics and perform administrative operations. You must enable statistics separately to JMX otherwise Data Grid provides **0** values for all statistic attributes.

Procedure

Modify your cache container configuration to enable JMX in one of the following ways:

- Declarative: Add the `<jmx enabled="true" />` element to the cache container.
- Programmatic: Call the `.jmx().enable()` method.

Declarative

```
<cache-container>
  <jmx enabled="true" />
</cache-container>
```

Programmatic

```
GlobalConfiguration globalConfig = new GlobalConfigurationBuilder()
    .jmx().enable()
    .build();
```

3.3.2. Performing Cross-Site Operations

Perform cross-site operations via JMX clients.

Prerequisites

- Configure Data Grid to register JMX MBeans

Procedure

1. Connect to Data Grid with any JMX client.
2. Invoke operations from the following MBeans:
 - **XSiteAdmin** provides cross-site operations for caches.
 - **GlobalXSiteAdminOperations** provides cross-site operations for Cache Managers.

For example, to bring sites back online, invoke **bringSiteOnline(siteName)**.

See the *Data Grid JMX Components* documentation for details about available cross-site operations.

Reference

- [XSiteAdmin MBean](#)
- [GlobalXSiteAdminOperations MBean](#)

CHAPTER 4. MONITORING AND TROUBLESHOOTING GLOBAL DATA GRID CLUSTERS

Data Grid provides statistics for cross-site replication operations via JMX or the **/metrics** endpoint for Data Grid server.

Cross-site replication statistics are available at cache level so you must explicitly enable statistics for your caches. Likewise, if you want to collect statistics via JMX you must configure Data Grid to register MBeans.

Data Grid also includes an **org.infinispan.XSITE** logging category so you can monitor and troubleshoot common issues with networking and state transfer operations.

4.1. ENABLING DATA GRID STATISTICS

Configure Data Grid to export statistics for Cache Managers and caches.

Procedure

Modify your configuration to enable Data Grid statistics in one of the following ways:

- Declarative: Add the **statistics="true"** attribute.
- Programmatic: Call the **.statistics()** method.

Declarative

```
<!-- Enables statistics for the Cache Manager. -->
<cache-container statistics="true">
  <!-- Enables statistics for the named cache. -->
  <local-cache name="mycache" statistics="true"/>
</cache-container>
```

Programmatic

```
GlobalConfiguration globalConfig = new GlobalConfigurationBuilder()
  //Enables statistics for the Cache Manager.
  .cacheContainer().statistics(true)
  .build();

Configuration config = new ConfigurationBuilder()
  //Enables statistics for the named cache.
  .statistics().enable()
  .build();
```

4.2. CONFIGURING DATA GRID METRICS

Configure Data Grid to export gauges and histograms via the **metrics** endpoint.

Procedure

- Turn gauges and histograms on or off in the **metrics** configuration as appropriate.

Declarative

```

<!-- Computes and collects statistics for the Cache Manager. -->
<cache-container statistics="true">
  <!-- Exports collected statistics as gauge and histogram metrics. -->
  <metrics gauges="true" histograms="true" />
</cache-container>

```

Programmatic

```

GlobalConfiguration globalConfig = new GlobalConfigurationBuilder()
  //Computes and collects statistics for the Cache Manager.
  .statistics().enable()
  //Exports collected statistics as gauge and histogram metrics.
  .metrics().gauges(true).histograms(true)
  .build();

```

4.2.1. Collecting Data Grid Metrics

Collect Data Grid metrics with monitoring tools such as Prometheus.

Prerequisites

- Enable statistics. If you do not enable statistics, Data Grid provides **0** and **-1** values for metrics.
- Optionally enable histograms. By default Data Grid generates gauges but not histograms.

Procedure

- Get metrics in Prometheus (OpenMetrics) format:

```
$ curl -v http://localhost:11222/metrics
```

- Get metrics in MicroProfile JSON format:

```
$ curl --header "Accept: application/json" http://localhost:11222/metrics
```

Next steps

Configure monitoring applications to collect Data Grid metrics. For example, add the following to **prometheus.yml**:

```

static_configs:
  - targets: ['localhost:11222']

```

Reference

- [Prometheus Configuration](#)
- [Enabling Data Grid Statistics](#)

4.3. CONFIGURING DATA GRID TO REGISTER JMX MBEANS

Data Grid can register JMX MBeans that you can use to collect statistics and perform administrative operations. You must enable statistics separately to JMX otherwise Data Grid provides **0** values for all statistic attributes.

Procedure

Modify your cache container configuration to enable JMX in one of the following ways:

- Declarative: Add the `<jmx enabled="true" />` element to the cache container.
- Programmatic: Call the `.jmx().enable()` method.

Declarative

```
<cache-container>
  <jmx enabled="true" />
</cache-container>
```

Programmatic

```
GlobalConfiguration globalConfig = new GlobalConfigurationBuilder()
    .jmx().enable()
    .build();
```

4.3.1. JMX MBeans for Cross-Site Replication

Data Grid provides JMX MBeans for cross-site replication that let you gather statistics and perform remote operations.

The `org.infinispan:type=Cache` component provides the following JMX MBeans:

- **XSiteAdmin** exposes cross-site operations that apply to specific cache instances.
- **StateTransferManager** provides statistics for state transfer operations.
- **InboundInvocationHandler** provides statistics and operations for asynchronous and synchronous cross-site requests.

The `org.infinispan:type=CacheManager` component includes the following JMX MBean:

- **GlobalXSiteAdminOperations** exposes cross-site operations that apply to all caches in a cache container.

For details about JMX MBeans along with descriptions of available operations and statistics, see the *Data Grid JMX Components* documentation.

Reference

[Data Grid JMX Components](#)

4.4. COLLECTING LOGS AND TROUBLESHOOTING CROSS-SITE REPLICATION

Diagnose and resolve issues related to Data Grid cross-site replication. Use the Data Grid Command Line Interface (CLI) to adjust log levels at run-time and perform cross-site troubleshooting.

Procedure

1. Open a terminal in **\$RHDG_HOME**.
2. Create a Data Grid CLI connection.
3. Adjust run-time logging levels to capture DEBUG messages if necessary.
For example, the following command enables DEBUG log messages for the org.infinispan.XSITE category:

```
[//containers/default]> logging set --level=DEBUG org.infinispan.XSITE
```

You can then check the Data Grid log files for cross-site messages in the **`\${rhdg.server.root}/log** directory.

4. Use the **site** command to view status for backup locations and perform troubleshooting.

For example, check the status of the "customers" cache that uses "LON" as a backup location:

```
[//containers/default]> site status --cache=customers
{
  "LON" : "online"
}
```

Another scenario for using the Data Grid CLI to troubleshoot is when the network connection between backup locations is broken during a state transfer operation.

If this occurs, Data Grid clusters that receive state transfer continually wait for the operation to complete. In this case you should cancel the state transfer to the receiving site to return it to a normal operational state.

For example, cancel state transfer for "NYC" as follows:

```
[//containers/default]> site cancel-receive-state --cache=mycache --site=NYC`
```

Reference

- [Data Grid Server Troubleshooting](#)
- [Working with Data Grid Server Logs](#)

4.4.1. Cross-Site Log Messages

Find user actions for log messages related to cross-site replication.

Log level	Identifier	Message	Description
-----------	------------	---------	-------------

Log level	Identifier	Message	Description
DEBUG	ISPN000400	Node null was suspected	Data Grid prints this message when it cannot reach backup locations. Ensure that sites are online and check network status.
INFO	ISPN000439	Received new x-site view: <code>#{site.name}</code>	Data Grid prints this message when sites join and leave the global cluster.
INFO	ISPN100005	Site <code>#{site.name}</code> is online.	Data Grid prints this message when a site comes online.
INFO	ISPN100006	Site <code>#{site.name}</code> is offline.	Data Grid prints this message when a site goes offline. If you did not take the site offline manually, this message could indicate a failure has occurred. Check network status and try to bring the site back online.
WARN	ISPN000202	Problems backing up data for cache <code>#{cache.name}</code> to site <code>#{site.name}</code> :	Data Grid prints this message when issues occur with state transfer operations along with the exception. If necessary adjust Data Grid logging to get more fine-grained logging messages.
WARN	ISPN000289	Unable to send X-Site state chunk to <code>#{site.name}</code> .	Indicates that Data Grid cannot transfer a batch of cache entries during a state transfer operation. Ensure that sites are online and check network status.
WARN	ISPN000291	Unable to apply X-Site state chunk.	Indicates that Data Grid cannot apply a batch of cache entries during a state transfer operation. Ensure that sites are online and check network status.
WARN	ISPN000322	Unable to re-start x-site state transfer to site <code>#{site.name}</code>	Indicates that Data Grid cannot resume a state transfer operation to a backup location. Ensure that sites are online and check network status.

Log level	Identifier	Message	Description
ERROR	ISPN000477	Unable to perform operation <code>\$\$\$operation.name</code> for site <code>\$\$\$site.name</code>	Indicates that Data Grid cannot successfully complete an operation on a backup location. If necessary adjust Data Grid logging to get more fine-grained logging messages.
FATAL	ISPN000449	XSite state transfer timeout must be higher or equals than 1 (one).	Results when the value of the timeout attribute is 0 or a negative number. Specify a value of at least 1 for the timeout attribute in the state transfer configuration for your cache definition.
FATAL	ISPN000450	XSite state transfer waiting time between retries must be higher or equals than 1 (one).	Results when the value of the wait-time attribute is 0 or a negative number. Specify a value of at least 1 for the wait-time attribute in the state transfer configuration for your cache definition.
FATAL	ISPN000576	Cross-site Replication not available for local cache.	Cross-site replication does not work with the local cache mode. Either remove the backup configuration from the local cache definition or use a distributed or replicated cache mode.