


---

# WebSphere eXtreme Scale Version 8.6

## Near cache invalidation



This presentation covers the new near cache invalidation feature added in WebSphere® eXtreme Scale version 8.6. This presentation assumes you are familiar with WebSphere eXtreme Scale software or the WebSphere DataPower® XC10 appliance and have some knowledge of WebSphere eXtreme Scale configuration model and WebSphere eXtreme Scale APIs.

---

## Table of contents

- Data invalidation
- Near cache invalidation overview
- Near cache invalidation configuration and usage details
- Troubleshooting
- Summary

This presentation covers data invalidation in general and what mechanisms already exist for removing stale data, the new cache invalidation feature, and some troubleshooting pointers.

## Data invalidation

- Administrative invalidation
  - use the web console or the xscmd utility to invalidate data based on the key
- Event-based invalidation
  - Event queue – stores the data change events
  - Event publisher - publishes the data change events to the event queue
  - Event consumer - consumes data change events

There are several ways to remove stale data from cache. One is administrative invalidation. You can use the web console or the xscmd utility to invalidate data based on the key. You can filter the cache data with a regular expression and then invalidate the data based on the regular expression. You can use the query interfaces in the monitoring console and in the xscmd utility to retrieve small sets of keys from a map and invalidate sets of data. If you are using the web console to query and invalidate data, configure the monitoring console first. If you are using xscmd to query and invalidate data, set up the xscmd utility.

Event based invalidation is another way of removing stale data. Sparse and complete caches can be invalidated or updated using an event generator such as Java Message Service (JMS). Invalidation using JMS can be manually tied to any process that updates the back-end using a database trigger. A JMS ObjectGridEventListener plug-in is provided in eXtreme Scale that can notify clients when the server cache changes. This type of notification decreases the amount of time that the client can see stale data. Event-based invalidation normally consists of three components; Event queue, Event publisher, and Event consumer. An Event queue stores the data change events. It can be a JMS queue, a database, an in-memory queue, or any kind of manifest as long as it can manage the data change events. An Event publisher publishes the data change events to the event queue. An event publisher is typically an application you create or an eXtreme Scale plug-in implementation. The event publisher knows when the data is changed or it changes the data itself. When a transaction commits, events are generated for the changed data and the event publisher publishes these events to the event queue. An Event consumer consumes data change events. An event consumer typically is an application to ensure the target grid data is updated with the latest change from other grids. This event consumer interacts with the event queue to get the latest data change and applies the data changes in the target grid. The event consumers can use eXtreme Scale APIs to invalidate stale data or update the grid with the latest data. For example, JMSObjectGridEventListener has an option for a client-server model, in which the event queue is a designated JMS destination. All server processes are event publishers. When a transaction commits, the server gets the data changes and publishes them to the designated JMS destination. All the client processes are event consumers. They receive the data changes from the designated JMS destination and apply the changes to the client's near cache.

## Data invalidation


- Programmatic invalidation – manual programming using WebSphere eXtreme Scale APIs
- Global index invalidation - Enabled through `dynacache-remote-deployment.xml` and `dynacache-remote-global-index-objectgrid.xml` descriptors
- **Near cache invalidation** - configure an asynchronous invalidation that is triggered by an update, delete, or invalidation operation

Another way to handle data invalidation is to do Programmatic invalidation. You can do manual programming using WebSphere eXtreme Scale APIs to handle stale data. You can also use programmatic invalidation with other techniques to determine when to invalidate the data. For example, this invalidation method uses event-based invalidation mechanisms to receive the data change events, and then uses APIs to invalidate the stale data.

Global index invalidation is also another way to handle data invalidation. You can optionally enable global index invalidation to improve invalidation efficiency in a large, partitioned environment; for example, more than 40 partitions. Without the global index feature, the dynamic cache template and dependency invalidation send agents to all partitions and therefore, you might experience slower performance. With global index, invalidation agents are sent only to applicable partitions that contain cache entries to invalidate and related template and dependency index data. In an environment with 100 partitions as an example, you might experience performance that is three times as fast when you remove stale cache data. Global index invalidation is enabled through `dynacache-remote-deployment.xml` and `dynacache-remote-global-index-objectgrid.xml`.

Near cache invalidation is new in WebSphere eXtreme Scale version 8.6. If you are using a near cache, you can configure an asynchronous invalidation that is triggered each time an update, delete, or invalidation operation is run against the data grid.

## Near cache



- local copy of the stored data in the eXtreme Scale Grid
- Independent ObjectGrid on each client serving as a cache for the remote, server-side cache
- Very fast – provides in-memory access to a subset of the entire cached data set that is stored remotely
- Not partitioned and contains data from any of the remote eXtreme Scale partitions
- Fast response time because all access to the data is local
- Disadvantage – stale data
  - Use when response time is important and stale data can be tolerated

5

Near cache invalidation

© 2013 IBM Corporation

Near cache is a local copy of the stored data in the eXtreme Scale grid. It is an independent ObjectGrid® on each client, serving as a cache for the remote, server-side cache. Near cache is very fast because it provides in-memory access to a subset of the entire cached data set that is stored remotely in the eXtreme Scale servers. It is not partitioned and contains data from any of the remote eXtreme Scale partitions.

The advantage of near cache is that it gives fast response time because all access to the data is local. Looking for the data in the near cache first saves a trip to the grid of servers, thus making even the remote data locally accessible. The disadvantage of the near cache is that it increases the chance of stale data because the near cache at each tier can be out of sync with the current data in the data grid. It also relies on an evictor to invalidate data to avoid running out of memory. Near cache is recommended for use when response time is important and stale data can be tolerated.

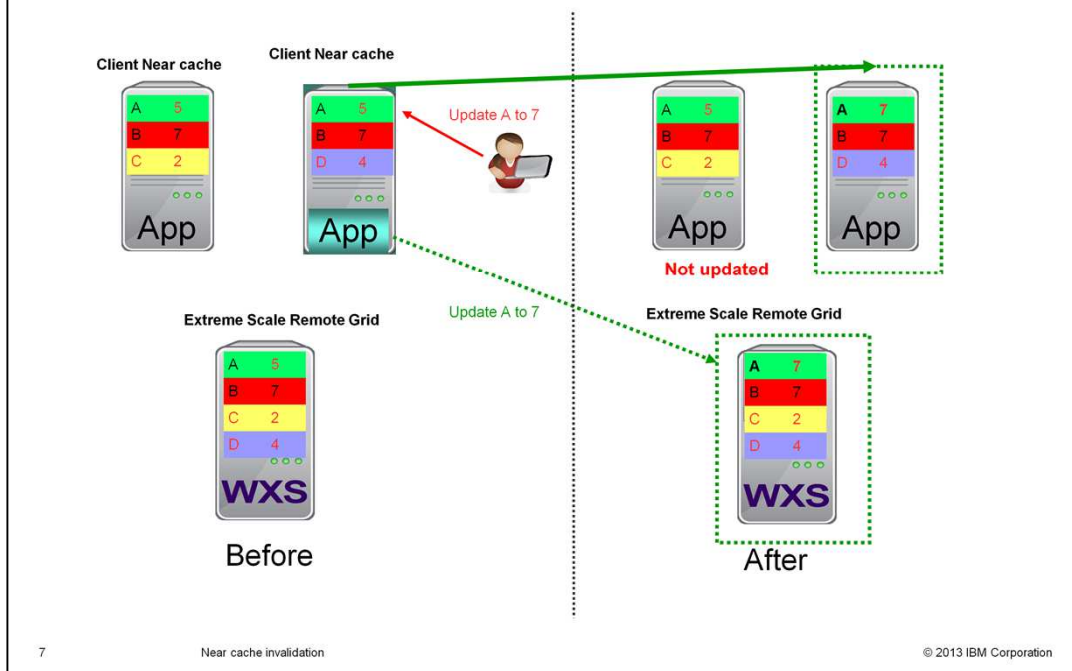
## Near cache invalidation key concepts

- Near cache – a local copy of the data stored in an eXtreme Scale Grid
  - This copy can be out of date with respect to the data in the grid
- Near cache invalidation – removing stale data from the near cache as quickly as possible instead of using time to live eviction
  - As a WebSphere eXtreme Scale developer you want your client JVM to have a relatively current near cache to avoid remote calls to the grid while limiting exposure to stale data

In a distributed WebSphere eXtreme Scale environment, the client side has a near cache by default when using the optimistic locking strategy or when locking is disabled. As you have seen in the previous slide, near cache is a local copy of the stored data in the eXtreme Scale grid. This data can become stale from time to time thus the need for near cache invalidation feature. Near cache invalidation feature has the ability to remove stale data from the near cache as quickly as possible instead of using TimeToLive eviction.

As a developer you want your client JVM to have a current near cache to avoid remote calls to the grid while limiting exposure to stale data. You can configure near cache invalidation to remove stale data from the near cache as quickly as possible. When an update, deletion, or invalidation operation is run against the remote data grid, an asynchronous invalidation gets triggered in the near cache. This mechanism works more quickly than the other option of using time-to-live (TTL) eviction in the near cache.

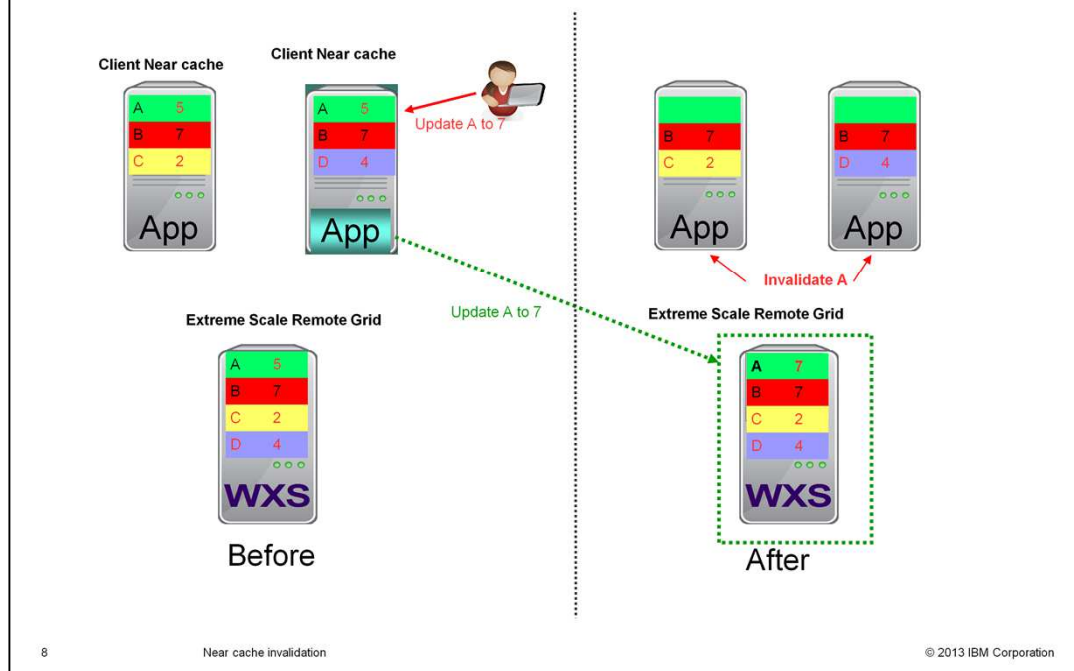
## Example: No near cache invalidation feature



Here is an example of how data can become stale.

Starting from the left, suppose you have a grid and some clients. These clients have near caches, and so have cached a few values locally. As an example, suppose someone updates key A from 5 to 7 on one of the clients. The remote grid gets updated to 7, and the client that did the update now has a value of 7 as shown on the right. However, for the other client that had cached the key A value, the cache still has a value of 5 and is not aware of the update. This is the problem that the near cache invalidation feature solves.

### Example: With near cache invalidation feature



With the near cache invalidation feature in place, if someone updates the key A value from 5 to 7 on one of the clients, the remote grid is updated with a value of 7 and the remote grid will issue an invalidation message to all the clients to invalidate their key A. This means that now those clients will need to go to the remote Grid to access the key A since it is no longer available in the near cache. The result is that stale data is removed from the near cache as quickly as possible.



## Near cache invalidation options today

- Time To Live evictor – remove cache entry after X minutes
  - Based on create time, last access time, or last update time
- JMSObjectGridEventListener
  - Existing eXtreme Scale plug-in
  - Synchronizes client caches
  - Sends all updates, not just invalidation
  - Requires external JMS bus configuration/deployment
- Roll your own
  - Trigger near cache invalidations based on some event (for example, a database update)

There are several cache invalidation options available. One of them is to configure the Timetolive evictor. You remove cache entries after a certain period of time based on when the cache was created, or when last accessed or updated. This way you are guaranteed the near cache entry will not hold an entry for more than five minutes.

The second invalidation option is the JMSObjectGridEventListener. This is a plug-in in eXtreme scale that synchronizes client caches and sends all updates in addition to invalidation. This listener is a JMS-based implementation of ObjectGridEventListener. It is designed to support client side near cache invalidation and peer-to-peer replication mechanisms.

The client invalidation mechanism can be used in distributed ObjectGrid environment to ensure client near cache data is synchronized with servers or other clients. Without this function, the client near cache can hold stale data. However, even with this JMS-based client invalidation mechanism, you have to realize that there is a timing window for updating the client near cache. There is a delay for the ObjectGrid runtime to publish updates. The downside of this plug-in is that it is network-intensive and requires external JMS bus configuration and deployment.

The third invalidation option is to create your own near cache invalidations triggered by some event; for example, a database update.

## Near cache invalidation feature

New!

- Triggered by server side update, remove, invalidate, or clear
  - Affected keys are sent to interested clients
  - Affected keys are removed from client's near cache
- Does not completely prevent stale data in near cache
  - Invalidation notice is asynchronous
- No changes to client code required
  - Configured through objectgrid.xml
- Must use optimistic or no locking
  - Near caches are not available with pessimistic locking
- Near cache evictor still recommended
  - Ensure near cache does not grow unbounded

The near cache invalidation feature allows for stale data removal from the near cache as quickly as possible. When a server side update occurs, affected keys are sent to the interested clients and the keys are removed from the clients' near cache. This helps lower the occurrence of stale data in the near cache. The feature does not require any changes to the client code instead it is configured through the objectgrid.xml. The near cache invalidation feature does not completely prevent stale data in near cache. It might take time before invalidation notice is received by clients since invalidation notice is asynchronous. The update can occur while taking some time for the clients to get the notification to remove the stale data. Also note that near caches are not available with pessimistic locking. They only use optimistic or no locking mechanism. This is because near caches work best where fast response is needed and data is not changed frequently. Another point to note is that the near cache evictor is still recommended to prevent near cache from growing unboundedly in terms of size. Use eviction policies like leastRecentlyUsed eviction policy, most recently used policy, and so on.

## How to configure the new near cache invalidation feature

- Use IBM eXtreme IO
- Near cache must be enabled
  - Run `BackingMap.isNearCacheEnabled()` to check
- On the `ObjectGrid.xml` file:
  - Set `nearCacheInvalidationEnabled` attribute to `true`
  - Set `lockStrategy` attribute to `NONE` or `OPTIMISTIC`

```
<backingMap name="TestMap1"  
nearCacheInvalidationEnabled="true" lockStrategy="OPTIMISTIC"
```

- Restart the container servers and clients

IBM eXtreme IO is a non-platform-specific network transport protocol that replaces the ORB for remote data transport. This transport protocol enables you to use clients other than Java. See the IBM WebSphere eXtreme Scale version 8.6 information center for more information on this transport. Use eXtreme IO with near cache.

You can check to see if you are using the near cache by running the `BackingMap.isNearCacheEnabled` method.

To enable the near cache invalidation feature you need to set two attributes in the backing map of the `objectGrid.xml` file. Set the `nearCacheInvalidationEnabled` attribute to `true`, and set the `lockStrategy` attribute to `NONE` or `OPTIMISTIC`. The settings in `objectgrid.xml` apply to all clients, unless you override the settings.

Save your XML file and you are ready to restart your container servers and clients.

## Cache invalidation - ObjectGrid configuration xml

```
<?xml version="1.0" encoding="UTF-8" ?>
- <objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd" xmlns="http://ibm.com/ws/objectgrid/config">
- <objectGrids>
- <objectGrid name="Grid">
  <backingMap name="TestMap1" nearCacheInvalidationEnabled="true" lockStrategy="OPTIMISTIC" copyMode="COPY_TO_BYTES" />
  <backingMap name="TestMapNoInvalidation" lockStrategy="OPTIMISTIC" copyMode="COPY_TO_BYTES" />
</objectGrid>
</objectGrids>
</objectGridConfig>
```

Here is an example of an ObjectGrid xml file configured for near cache invalidation. An ObjectGrid descriptor XML file is used to define the structure of the ObjectGrid that is used by the application. It includes a list of backing map configurations. These backing maps store the cache data. Take note of the two map elements; TestMap1 that has nearcacheInvalidation enabled and TestMapNoInvalidation with no near cache invalidation. You can have some maps with this enablement and some maps without. Note the lock strategy must be optimistic.

## Cache invalidation – deployment.xml

```
<?xml version="1.0" encoding="UTF-8" ?>
- <deploymentPolicy xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/deploymentPolicy ../deploymentPolicy.xsd"
  xmlns="http://ibm.com/ws/objectgrid/deploymentPolicy">
- <objectgridDeployment objectgridName="Grid">
  - <mapSet name="TestMapSet1" developmentMode="true" numInitialContainers="1" numberOfPartitions="5" minSyncReplicas="0"
    maxSyncReplicas="0" maxAsyncReplicas="1">
    <map ref="TestMap1" />
    <map ref="TestMapNoInvalidation" />
  </mapSet>
</objectgridDeployment>
</deploymentPolicy>
```

Here is the corresponding deployment descriptor xml file. This file defines the objectgridDeployment element for the ObjectGrid. Note the number of partitions for the mapSet. The minSyncReplicas and maxSyncReplicas attributes of mapSet are to specify the minimum and maximum number of synchronous replicas for each partition in the mapSet. These are optional attributes and the default for each is 0.

## Cache invalidation - example

```

UpdateClient
map1.update("key1map1", "value" + valueCounter);
map2.update("key1map2", "value" + valueCounter);
Object v1 = map1.get("key1map1");
Object v2 = map2.get("key1map2");
System.out.println("client "+clientName+" map1(invalidation) value is now: " + v1);
System.out.println("client "+clientName+" map2(no invalidation) value is now: " + v2);

output:
client updater map1(invalidation) value is now: value7
client updater map2(no invalidation) value is now: value7
client updater map1(invalidation) value is now: value8
client updater map2(no invalidation) value is now: value8
client updater map1(invalidation) value is now: value9
client updater map2(no invalidation) value is now: value9
client updater map1(invalidation) value is now: value10
client updater map2(no invalidation) value is now: value10
client updater map1(invalidation) value is now: value11
client updater map2(no invalidation) value is now: value11

ReadClient
Object v1 = map1.get("key1map1");
Object v2 = map2.get("key1map2");
System.out.println("client "+clientName+" map1(invalidation) value is now: " + v1);
System.out.println("client "+clientName+" map2(no invalidation) value is now: " + v2);

output:
client reader map1(invalidation) value is now: value8
client reader map2(no invalidation) value is now: value8
client reader map1(invalidation) value is now: value9
client reader map2(no invalidation) value is now: value8
client reader map1(invalidation) value is now: value10
client reader map2(no invalidation) value is now: value8
client reader map1(invalidation) value is now: value11
client reader map2(no invalidation) value is now: value8

```

Here is a cache invalidation example that contains an update client that is updating keys in two different maps. One of those maps has invalidation enabled and one of them does not. The update client is running in a loop updating the keys - one in each map. The result is that it updates both key value pairs.

The read client is reading the values of the keys from the same maps; one with invalidation the other without.

Initially both keys had values of 8 but as the loop iterates, the client with invalidation sees the new value update from 8 to 9 and then to 10. This is because the value will keep getting "invalidated" by the updater client, so it will keep reading the latest value from the grid, rather than using a stale value from the near cache. Then the client with no invalidation caches the first value it reads and never sees an updated value. It only sees the old value of 8.

## Near cache sizing

- Near cache size is unlimited by default
  - This can cause the client JVM to run out of memory
- A client side configuration override can be used to configure an evictor
  - Evictor will only run against the near cache
  - Will limit the maximum size of the cache
  - LRU evictor is typically recommended
- The client side override can be configured in two ways
  - Programmatically
  - Using a client-side objectgrid.xml override

By default there is no restriction on the maximum size of a near cache. This means that a client can easily run out of memory by storing a large number of items in the near cache. To prevent this, an evictor can be configured to run against the near cache only. This evictor will have no effect on the data stored in the remote data grid, but it will control the maximum size of the near cache. The mechanism for accomplishing this is a client side override, which allows a client to specify configuration information that will only apply to the client process but not to the remote data grid configuration. These overrides can be specified programmatically or using an XML file provided by the client during initialization.

## Configuring a near cache evictor programmatically

```
ObjectGridConfiguration companyGridConfig = ObjectGridConfigFactory
    .createObjectGridConfiguration("CompanyGrid");
BackingMapConfiguration customerMapConfig = ObjectGridConfigFactory
    .createBackingMapConfiguration("Customer");

Plugin evictorPlugin = ObjectGridConfigFactory.createPlugin(PluginType.EVICTOR,
    "com.ibm.websphere.objectgrid.plugins.builtins.LRUEvictor");

ConfigProperty
numQueues=ObjectGridConfigFactory.createConfigProperty(ConfigPropertyType.INT_PRIM,
    "numberOfLRUQueues", "53");
evictorPlugin.addConfigProperty(numQueues);
ConfigProperty
maxSize=ObjectGridConfigFactory.createConfigProperty(ConfigPropertyType.INT_PRIM,
    "maxSize", "1000");
evictorPlugin.addConfigProperty(maxSize);
customerMapConfig.addPlugin(evictorPlugin);
companyGridConfig.addBackingMapConfiguration(customerMapConfig);

ClientClusterContext client = ogManager.connect(catalogServerEndpoints, null, null);
ObjectGrid companyGrid = ogManager.getObjectGrid(client,
    objectGridName, companyGridConfig);
```

This example code illustrates the process of defining a client side configuration override that enables a Least Recently Used evictor. First, a new configuration object for the grid and map are created. Next, a new configuration plug-in of type Evictor is created, specifying the LRU evictor class name. The evictor is then configured with 53 queues and a maximum size of 1000 items per queue, resulting in an LRU evictor that will enforce a maximum size of 53,000 entries for the near cache supporting this particular map. Finally, the configured evictor is added to the configuration object and passed to the Object Grid Manager during the connection step.



## Configuring a near cache evictor in XML

```
<?xml version="1.0" encoding="UTF-8"?>
<objectGridConfig xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://ibm.com/ws/objectgrid/config ../objectGrid.xsd"
  xmlns="http://ibm.com/ws/objectgrid/config">
  <objectGrids>
    <objectGrid name="CompanyGrid">
      <backingMap name="Customer" nearCacheEnabled="true" pluginCollectionRef="customerPlugins"/>
    </objectGrid>
  </objectGrids>
  <backingMapPluginCollections>
    <backingMapPluginCollection id="customerPlugins">
      <!-- Limit the near cache size to 53*1000=53,000 entries using an LRU algorithm -->
      <bean id="Evictor" className="com.ibm.websphere.objectgrid.plugins.builtins.LRUEvictor">
        <property name="numberOfLRUQueues" type="int" value="53" description="set number of LRU
          queues" />
        <property name="maxSize" type="int" value="1000" description="set max size for each LRU
          queue" />
      </backingMapPluginCollection>
    </backingMapPluginCollections>
  </objectGridConfig>

ObjectGridManager ogManager = ObjectGridManagerFactory.ObjectGridManager(); ClientClusterContext clientClusterContext =
ogManager.connect("MyServer1.company.com:2809", null, new URL( "file:xml/companyGridClientSide.xml"));
```

Here a client side override XML file is illustrated. The file copies the grid and map definitions from the server side objectgrid xml file, but adds a configuration definition for a Least Recently Used evictor on the Customer map. This file is then provided by the client during connection to the grid, as shown in the sample code. This will result in an LRU evictor being configured to control the size of the near cache for the Customer map.

## Troubleshooting: sample error and warning messages (1 of 2)

- CWOBJ7661I: Initializing near cache invalidation for grid {0} map {1}.
  - Confirms NCI is configured/enabled. Seen on client and server. Includes partition on server side.
- CWOBJ7662E: Client near-cache invalidation is enabled for grid {0} map {1}, but not eXtremeIO is not enabled.
  - Error, NCI requires XIO (cannot use with ORB transport)
- CWOBJ7663W: Near Cache invalidation notification has been disconnected for grid {0}, map {1}, partition {2}, the near cache has been disabled.
  - The client has not received messages from the publisher for this partition recently, as a result it may be disconnected and proactively disabled the nearcache to avoid stale data.
  - This disables the entire nearcache, not just content for this partition
- CWOBJ7664I: Near Cache invalidation notification has been re-enabled for grid {0}, map {1}.
  - The client re-established a connection to all partition publishers

Here are some troubleshooting error and warning messages. Some of the messages shown in this slide include a message confirming near cache invalidation is configured and enabled, near cache invalidation disabled, near cache invalidation re-enabled, and messages for failures with the feature. You can look up more of these messages in the information center.

## Troubleshooting: sample error and warning messages (2 of 2)

- CWOBJ7652W: The subscriber did not receive one or more invalidation messages for the near cache invalidation topic {0}
  - The client will attempt to resynchronize the near cache with the server, in the meantime stale data may be seen in the nearcache
- CWOBJ7657E: Client near-cache invalidation is enabled, but not supported for backingMap, {0} and objectGrid, {1}.
  - Config issue, ensure map is not a pessimistic locking map, does not have near cache disabled
- CWOBJ7658E: A failure occurred installing plug-ins which support the Near Cache Invalidation and Continuous Query features for objectGrid: {0}. The error is {1}.
  - Significant install issue, likely jars/classes not found during container jvm startup. NCI will not work.

This slide shows more troubleshooting warning and error messages, including a warning that an invalidation message did not arrive, and error messages resulting from configuration and installation issues.

## Troubleshooting: possible cache invalidation issues

- **Scenario 1: Client is not reading data from the near cache**
  - Ensure they are using optimistic or none for locking
  - Near cache is not disabled
  - Bucket size is not zero
  - Ensure no evictor configured that might be purging the near cache prematurely
- **Scenario 2: Client is reading stale data from the nearcache**
  - Ensure NCI is properly enabled (init messages)
  - Client has not marked NCI disabled
    - It takes a few minutes to mark NCI disabled, stale data can be seen in the meantime
- **Scenario 3: Client JVM goes out of memory.**
  - Likely cause: near cache grew unbounded. Not an NCI problem
  - When using the nearcache, clients should configure a client-side evictor (typically LRU) to cap the maximum size of the nearcache. NCI will not control the upper limit of the near cache size

Here are some sample scenarios of possible near cache invalidation issues. In Scenario 1), the client is not reading data from the near cache. Make sure you are using optimistic or none for locking strategy, make sure near cache is enabled and the bucket size is not zero. Also make sure no evictor is configured that might be purging the near cache prematurely.

The second scenario is where client is reading stale data from the near cache. Make sure that near cache invalidation is enabled. You can do so by checking the init messages. Also make sure that the client process has not marked near cache invalidation as disabled. Note that It takes a few minutes to mark near cache invalidation disabled, and stale data can be seen in the meantime. Near cache invalidation does not guarantee that there is no stale data, and that invalidation messages take time to propagate.

In scenario three, the client JVM runs out of memory. In this scenario, the likely cause is that the near cache grew unbounded. When using the near cache, configure a client-side evictor (typically least recently used) to cap the maximum size of the near cache. Near cache invalidation will not control the upper limit of the near cache size.

## Troubleshooting: Near cache side effect

- Reads from near cache content do not touch server
  - This means last-access TTL maps will have server entries expire prematurely
- New feature in 8.6 to address, backing map attribute
  - nearCacheLastAccessTTLSyncEnabled=true
- Reads from near cache will drive a last access time update on the server
  - Prevents server side entries from hitting TTL expiration while being accessed by clients
- Independent of NCI feature/enablement

Something to consider as a side effect of near cache is that reads from near cache content do not make it to the server. This means that last-access TTL maps will have server entries expire prematurely. The new feature addresses this with a backing map attribute. The attribute `nearCacheLastAccessTTLSyncEnabled` is set to `true`. Note that the backing map can be considered as an in-memory cache of committed data for an individual map.

Also note that reads from near cache will drive a last access time update on the server. This prevents server side entries from hitting TTL expiration while being accessed by clients.

## Troubleshooting: Near cache invalidation trace

- Requires client and server trace:
  - ObjectGridPubSub=all:ObjectGridCacheInvalidator=all
- Normally need trace enabled before the issue occurs
  - If key in near cache did not get invalidated
    - too late for trace
    - Solution: enable trace, then perform an update on the key in question to trigger an invalidation. This will show where invalidation got lost

Trace can be configured for the near cache invalidation feature. This feature requires client and server trace to be turned on. To turn on trace, enter:

```
ObjectGridPubSub=all:ObjectGridCacheInvalidator=all
```

You should therefore enable trace before an issue arises. If you have a key in your near cache that did not get invalidated, it is probably too late to turn on trace and determine why it did not get invalidated if trace was not enabled. The best practice is to enable trace, then perform an update on the key in question to trigger an invalidation. This will show if any invalidation was sent and where it got lost.

---

## Summary

- Near cache invalidation feature is a great addition to other cache invalidation options present today
- Easy to configure without code changes (Configuration only in the XML)
- Faster removal of stale data from near caches compared with other invalidation options

In summary, you can configure near cache invalidation to remove stale data from the near cache as quickly as possible. No code changes are required. When an update, deletion, or invalidation operation is run against the remote data grid, an asynchronous invalidation is triggered in the near cache. This mechanism works more quickly than the other option of using time-to-live (TTL) eviction in the near cache. For troubleshooting information, see the WebSphere eXtreme Scale version 8.6 information center.

## Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

[mailto:iea@us.ibm.com?subject=Feedback\\_about\\_XS86\\_NearCache.ppt](mailto:iea@us.ibm.com?subject=Feedback_about_XS86_NearCache.ppt)

This module is also available in PDF format at: [../XS86\\_NearCache.pdf](..../XS86_NearCache.pdf)

You can help improve the quality of IBM Education Assistant content by providing feedback.





## Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, DataPower, ObjectGrid, and WebSphere are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2013. All rights reserved.