

Oracle® Coherence

Release Notes for Oracle Coherence

Release 3.5.1

E15433-01

August 2009

Oracle Coherence Release Notes for Oracle Coherence, Release 3.5.1

E15433-01

Copyright © 2009, Oracle and/or its affiliates. All rights reserved.

Primary Author: Thomas Pfaeffle

Contributing Author: Noah Arliss, Mark Falco, Alex Gleyzer, Gene Gleyzer, Jason Howes, Adam Leftik, Rob Misek, Patrick Peralta

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	v
Audience.....	v
Documentation Accessibility	v
Related Documents	vi
Conventions	vi
1 Technical Changes and Enhancements	
Oracle Coherence for Java 3.5.1	1-1
Coherence*Web Enhancements and Fixes	1-1
Management Framework Fixes.....	1-1
Other Enhancements and Fixes.....	1-1
Oracle Coherence for .NET 3.5.1	1-2
Oracle Coherence for C++ 3.5.1	1-3
2 Documentation Errata	
Correction to "Coherence*Web Session Management Features"	2-1
Correction to "Read-Through, Write-Through, Write-Behind, and Refresh-Ahead Caching" ..	2-2
Correction to Managing Map Operations with Triggers	2-2
Correction to the cache-config.dtd	2-2
Addition to the address-provider Subelement Description	2-2
Addition to the backing-map-scheme Subelement Description	2-3
Sizing Considerations for Coherence Cluster JVMs	2-3
Basic Sizing Recommendations.....	2-3
Cache Servers.....	2-3
TCMP Clients.....	2-4
Extend Clients.....	2-4
Storage Ratios	2-4
Running with Large Heaps	2-4
Using Available System Memory	2-4
Additions to Log and Error Messages	2-5
Additions to Partitioned Cache Service Log Messages	2-5
Additions to TCMP Log Messages	2-6

Index

List of Tables

2-1	Changes Made to the <i>Write-Behind Caching</i> Description.....	2-2
2-2	Changes Made to the backing-map-scheme Description.....	2-3

Preface

This document describes changes and enhancements that have been made to the Oracle Coherence product since the 3.5 release.

Audience

This document is intended for users of Oracle Coherence.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at <http://www.fcc.gov/cgb/consumerfacts/trs.html>, and a list of phone numbers is available at <http://www.fcc.gov/cgb/dro/trsphonebk.html>.

Related Documents

For more information, see the following documents in the Oracle Coherence documentation set:

- *Getting Started with Oracle Coherence*
- *Developer's Guide for Oracle Coherence*
- *Client Guide for Oracle Coherence*
- *Tutorial for Oracle Coherence*
- *User's Guide for Oracle Coherence*Web*
- *Integration Guide for Oracle Coherence*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Technical Changes and Enhancements

This chapter describes the changes and enhancements made to the Oracle Coherence product for the 3.5.1 release. This document is accurate at the time of publication. Oracle updates the release notes periodically after the software release.

Oracle Coherence for Java 3.5.1

The following is a list of new features, improvements, and bug fixes in Oracle Coherence for Java 3.5.1

Coherence*Web Enhancements and Fixes

- Optimized the session persistence for concurrently running requests.
- Fixed a regression in the session reaper.
- Fixed a concurrency issue caused by multiple requests for the same session running in parallel when thread locking is disabled.

Management Framework Fixes

- Added clean-up logic to unregister `ConnectionMBeans` in the case of abnormal connection termination.
- Added clean-up logic to unregister `CacheMbeans` by the `CacheFactory.release()` method.
- Fixed a potential deadlock caused by concurrent cluster shutdown and `InvocationService` startup.
- Added ability to choose an `MBeanServer` programmatically.
- Added ability to execute parametrized operations on custom `DynamicMBeans`.

Other Enhancements and Fixes

- Reduced the communication flow during partition distribution.
- Added ability to specify the default serializer type for the `ConfigurablePofContext`.
- Fixed the `hashCode` implementation in the `SimplePofPath` that prevented `PofExtractor`-based indexing to operate properly.
- Fixed a regression in the ownership conflict resolution protocol caused by rapid service membership change with storage-disabled service senior.

- Optimized query performance for comparison filters (`GreaterFilter`, `LessFilter`, and so on)
- Improved handling of unreachable addresses in WKA list.
- Made `ConfigurableAddressProvider` more tolerant to DNS resolution failures.
- Added the `ContextClassLoader` initialization for `MapEvents` fired on the `EventDispatcher` thread.
- Added an ability to configure the `SymmetricEncryptionFilters` transformation.
- Fixed a bug causing the `ReadWriteBackingMap` to ignore explicitly specified expiry times when `write-behind` is enabled.
- Fixed a bug in the `LikeFilter` with search patterns starting with an escaped wildcard.
- Fixed a regression in the backup distribution algorithm that could result in less balanced distribution.
- Fixed a potential `ProxyService` deadlock caused by `TcpAcceptor` failure.
- Fixed a potential deadlock between `TcpInitiator` and `TcpReader` threads during cluster shutdown.
- Fixed a bug causing a potential `IllegalStateException` thrown by the `ContinuousQueryCache.release()` call.
- Fixed POF serialization for `BigDecimal` and `BigInteger` objects.
- Fixed a typo in the `cache-config.dtd`. See ["Correction to the cache-config.dtd"](#) on page 2-2 for more detailed information.

Oracle Coherence for .NET 3.5.1

The following is a list of new features, improvements, and bug fixes in Oracle Coherence for .NET 3.5.1:

- Improved the `SynchronizedDictionary` and `LocalCache` implementations to support concurrent read operations.
- Made `ConfigurableAddressProvider` more tolerant to DNS resolution failures.
- Removed unnecessary wrapping of exceptions when they are rethrown.
- Fixed a potential deadlock between `TcpInitiator` and `TcpReader` threads during cluster shutdown.
- Fixed POF serialization for `BigDecimal` and `BigInteger` objects.
- Fixed a bug in the `LikeFilter` with search patterns starting with an escaped wildcard.
- Fixed the `hashCode` implementation in the `SimplePofPath` that prevented `PofExtractor`-based indexing to operate properly.
- Ensured that all internal `ICollection` instances are synchronized on their `SyncRoot` rather than the object reference.

Oracle Coherence for C++ 3.5.1

The following is a list of new features, improvements, and bug fixes in Oracle Coherence for C++ 3.5.1:

- Made `ConfigurableAddressProvider` more tolerant to DNS resolution failures.
- Added a conversion helper for `time_t` to `RawDate/RawTime`.
- Fixed a potential deadlock between `TcpInitiator` and `TcpReader` threads during cluster shutdown.
- Fixed a bug in the `LikeFilter` with search patterns starting with an escaped wildcard.
- Fixed an NPE caused by running the `ConditionalPut` processor on the `LocalCache`.
- Fixed the local processing of the `NotFilter::evaluateEntry()`.

Documentation Errata

This chapter describes changes, enhancements, and corrections made to the Oracle Coherence documentation library for 3.5.1. The library can be found at the following URL:

http://download.oracle.com/docs/cd/E14526_01/index.htm

- [Correction to "Coherence*Web Session Management Features"](#)
- [Correction to "Read-Through, Write-Through, Write-Behind, and Refresh-Ahead Caching"](#)
- [Correction to Managing Map Operations with Triggers](#)
- [Correction to the cache-config.dtd](#)
- [Addition to the address-provider Subelement Description](#)
- [Addition to the backing-map-scheme Subelement Description](#)
- [Sizing Considerations for Coherence Cluster JVMs](#)
- [Additions to Log and Error Messages](#)

Correction to "Coherence*Web Session Management Features"

Web applications that use different sticky optimization and locking settings should not be intermixed within the same cluster. With that in mind, the following note has been added to the *Session Models* section of the *Coherence*Web Session Management Features* chapter in the *User's Guide for Oracle Coherence*Web*.

Note: In general, Web applications that are part of the same Coherence cluster must use the same session model type. Inconsistent configurations may result in deserialization errors.

The following note has been added to the *Session Locking Modes* section of the *Coherence*Web Session Management Features* chapter in the *User's Guide for Oracle Coherence*Web*.

Note: In general, Web applications that are part of the same Coherence cluster must use the same locking mode and sticky session optimizations setting. Inconsistent configurations may result in deadlock.

Correction to "Read-Through, Write-Through, Write-Behind, and Refresh-Ahead Caching"

This section describes the changes made to the *Read-Through, Write-Through, Write-Behind, and Refresh-Ahead Caching* chapter in *Getting Started with Oracle Coherence*. Corrected text appears in *italics*.

Table 2–1 illustrates the changes made to the *Write-Behind Caching* section.

Table 2–1 Changes Made to the Write-Behind Caching Description

Old Text	New Text
In the *Write-Behind* scenario, modified cache entries are asynchronously written to the datasource after a configurable delay, whether after 10 seconds, 20 minutes, a day or even a week or longer. For *Write-Behind* caching, Coherence maintains a write-behind queue of the data that needs to be updated in the datasource. When the application updates X in the cache, X is added to the write-behind queue (if it isn't there already; otherwise, it is replaced), and after the specified write-behind delay Coherence will call the CacheStore to update the underlying datasource with the latest state of X. Note that the write-behind delay is relative to the first of a series of modifications—in other words, the data in the datasource will never lag behind the cache by more than the write-behind delay.	In the *Write-Behind* scenario, modified cache entries are asynchronously written to the datasource after a configurable delay, whether after 10 seconds, 20 minutes, a day or even a week or longer. <i>Note that this only applies to cache inserts and updates—cache entries are removed synchronously from the datasource.</i> For *Write-Behind* caching, Coherence maintains a write-behind queue of the data that needs to be updated in the datasource. When the application updates X in the cache, X is added to the write-behind queue (if it isn't there already; otherwise, it is replaced), and after the specified write-behind delay Coherence will call the CacheStore to update the underlying datasource with the latest state of X. Note that the write-behind delay is relative to the first of a series of modifications—in other words, the data in the datasource will never lag behind the cache by more than the write-behind delay.

Correction to Managing Map Operations with Triggers

In *Managing Map Operations with Triggers* in the *Developer's Guide for Oracle Coherence*: the introduction to Example 4–1 incorrectly states that the `createMapTrigger` method would return a new `MapTriggerListener(new MyCustomTrigger())`;

Instead of `createMapTrigger`, the name of the correct method should be `createTriggerListener`.

Correction to the cache-config.dtd

The pre-3.5.1 `cache-config.dtd` was not well-formed in that it was missing a comma (in line 445) between the `thread-count?` and `task-hung-threshold?` attributes in the `proxy-scheme` element definition.

```
...
<!ELEMENT proxy-scheme
(scheme-name?, scheme-ref?, service-name?, thread-count?
task-hung-threshold?, task-timeout?, request-timeout?, acceptor-config?,
proxy-config?, autostart?)>
...
```

This has been fixed for the 3.5.1 release.

Addition to the address-provider Subelement Description

The following addition was made to the description of the `address-provider` subelement of `well-known-addresses` in the *Developer's Guide for Oracle Coherence*:

The calling component will attempt to obtain the full list upon node startup, the provider must return a terminating null address to indicate that all available addresses have been returned.

Addition to the backing-map-scheme Subelement Description

The following addition was made to the description of `backing-map-scheme` subelement of the `distributed-scheme` element in the *Developer's Guide for Oracle Coherence*. Added text is in *italics*.

Table 2–2 Changes Made to the backing-map-scheme Description

Old Text	New Text
Note that when using an overflow-based backing map it is important that the corresponding <code><backup-storage></code> be configured for overflow (potentially using the same scheme as the backing-map). See "Partitioned Cache with Overflow" for an example configuration.	When using an off-heap backing map it is important that the corresponding <code><backup-storage></code> be configured for off-heap (potentially using the same scheme as the backing-map). <i>Here, off-heap refers to any storage where some or all entries are stored outside of the JVMs garbage collected heap space. Examples include overflow-scheme, and external-scheme. See "Partitioned Cache with Overflow" for an example configuration.</i>

Sizing Considerations for Coherence Cluster JVMs

This section discusses JVM sizing considerations for Coherence cluster JVMs. The primary issue to consider when sizing your JVMs is achieving a balance of available RAM versus garbage collection (GC) pause times.

GC Pauses

Lengthy GC pause times can negatively impact the Coherence cluster as they are, for the most part, indistinguishable from node death. For this reason, it is very important that cluster nodes are sized and/or tuned to ensure that their GC times remain minimal. As a good rule of thumb, a node should spend less than 10% of its time paused in GC, normal GC times should be under 100ms, and maximum GC times should be around 1 second.

You can monitor GC activity in several ways; some standard mechanisms include:

- JVM switch `-verbose:gc`
- JVM switch `-Xloggc:` (similar to `verbose gc` but includes timestamps)
- Over JMX using tools such as jConsole

Basic Sizing Recommendations

If you are looking to just get things up and running with minimal effort, the following recommendations should suffice.

Cache Servers

The standard safe recommendation for Coherence cache servers is to run fixed-size heap of up to 1GB. Additionally, it is recommended to use an incremental garbage collector to minimize GC pause durations, and to also run the JVM in "server" mode to encourage optimizations for long running processes.

For example, the following command allows for good performance without the need for more elaborate JVM tuning:

```
java -server -Xms1g -Xmx1g -Xincgc -Xloggc: -cp coherence.jar  
com.tangosol.net.DefaultCacheServer
```

TCMP Clients

Coherence TCMP clients should be configured similarly to cache servers as long GCs could cause them to be misidentified as being dead.

Extend Clients

Coherence Extend clients are not, technically speaking, cluster members, and as such the effect of long GCs is less detrimental. For extend clients, it is recommended that you follow the existing guidelines as set forth by the application in which you are embedding Coherence.

Storage Ratios

There is a related question of how much data you can store within a cache server of a given size. The basic recommendation is to use up to one-third of the heap for primary cache storage. This leaves another one-third for backup storage, and the final one-third for "scratch space". Scratch space is then used for things such as holding classes, temporary objects, network transfer buffers, and GC compaction. You may instruct Coherence to limit primary storage on a per-cache basis by using the `<high-units>` element, and specifying `BINARY` in the `<unit-calculator>` element. These settings are automatically applied to backup storage as well.

Ideally, both the primary and backup storage will also fit within the JVMs tenured space (for HotSpot based JVMs). See HotSpot's *Tuning Garbage Collection* guide for details on sizing the collectors generations.

See the *Developer's Guide for Oracle Coherence* for more information on the `<high-units>` and `<unit-calculator>` elements

Running with Large Heaps

It is possible to run cache servers with larger heap sizes, although it becomes more important to monitor and tune the JVMs to minimize the GC pauses. It may also be necessary to alter the storage ratios such that the amount of scratch space is increased to facilitate faster GC compactions. Additionally it is recommended that you make use of an up to date JVM version such as HotSpot 1.6 as it includes significant improvements for managing large heaps.

Using Available System Memory

Running multiple identical cache server instances on a single machine enables you to use the available system memory. It is important to not overcommit the available resources. Namely if you have a machine with 16GB of RAM, it is not reasonable to attempt to dedicate 16GB of memory to your JVMs. Ultimately when all the machines processes are running you want to be in a state that swap space is not actively being used. In selecting the size and number of JVMs to run, it is important to realize that the JVM process will use more memory than is specified when configuring the heap size. The heap size settings specify the amount of heap which the JVM makes available to the application, but the JVM itself will also consume additional memory. The amount consumed will differ depending on the OS, and JVM settings. For instance a HotSpot JVM running on Linux configured with a 1GB heap will consume roughly 1.2GB of RAM. It is important that you externally measure the JVMs memory utilization to ensure that you don't over commit your RAM. Tools such as 'top', 'vmstat', and 'Task Manager' are useful in identifying how much RAM is actually being used.

Additions to Log and Error Messages

The following most frequently seen messages have been documented for the 3.5.1 release.

- [Additions to Partitioned Cache Service Log Messages](#)
- [Additions to TCMP Log Messages](#)

Additions to Partitioned Cache Service Log Messages

The following partitioned cache service messages have been documented for the 3.5.1 release.

validatePolls: This service timed-out due to unanswered handshake request. Manual intervention is required to stop the members that have not responded to this Poll

Parameters: none;

Severity: 1—Error

Cause: When a node joins a clustered service, it performs a handshake with each clustered node running the service. A missing handshake response prevents this node from joining the service. Most commonly, it is caused by an unresponsive (for example, deadlocked) service thread.

Action: Corrective action may require locating and shutting down the JVM running the unresponsive service. See My Oracle Support (Metalink) Note 845363.1:

<https://metalink.oracle.com/CSP/main/article?cmd=show&type=NOT&id=845363.1>

Error while starting cluster: com.tangosol.net.RequestTimeoutException: Timeout during service start: ServiceInfo(%s)

Parameters: %s—information on the service that could not be started;

Severity: 1—Error

Cause: When joining a service, every service in the cluster must respond to the join request. If one or more nodes have a service that does not respond within the timeout period, the join times out.

Action: See My Oracle Support (Metalink) Note 845363.1:

<https://metalink.oracle.com/CSP/main/article?cmd=show&type=NOT&id=845363.1>

Failed to restart services: com.tangosol.net.RequestTimeoutException: Timeout during service start: ServiceInfo(%s)

Parameters: %s—information on the service that could not be started

Severity: 1—Error

Cause: When joining a service, every service in the cluster must respond to the join request. If one or more nodes have a service that does not respond within the timeout period, the join times out.

Action: See My Oracle Support (Metalink) Note 845363.1:

<https://metalink.oracle.com/CSP/main/article?cmd=show&type=NOT&id=845363.1>

Additions to TCMP Log Messages

The following TCMP log messages have been documented for the 3.5.1 release.

This node appears to have partially lost the connectivity: it receives responses from MemberSet(%s1) which communicate with Member(%s2), but is not responding directly to this member; that could mean that either requests are not coming out or responses are not coming in; stopping cluster service.

Parameters:

%s1—set of members that can communicate with the member indicated in %s2;

%s2—member that can communicate with set of members indicated in %s1.

Severity: 1—Error

Cause: The communication link between this member and the member indicated by %s2 has been broken. However, the set of witnesses indicated by %s1 report no communication issues with %s2. It is therefore assumed that this node is in a state of partial failure, thus resulting in the shutdown of its cluster threads.

Action: Corrective action is not necessarily required, since the rest of the cluster presumably is continuing its operation and this node may recover and rejoin the cluster. On the other hand, it may warrant an investigation into root causes of the problem (especially if it is recurring with some frequency).

validatePolls: This senior encountered an overdue poll, indicating a dead member, a significant network issue or an Operating System threading library bug (e.g. Linux NPTL): Poll

Parameters: none.

Severity: 2—Warning

Cause: When a node joins a cluster, it performs a handshake with each cluster node. A missing handshake response prevents this node from joining the service. The log message following this one will indicate the corrective action taken by this node.

Action: If this message reoccurs, further investigation into the root cause may be warranted.

Received panic from senior Member(%s1) caused by Member(%s2)

Parameters:

%s1—the cluster senior member as known by this node;

%s2—a member claiming to be the senior member.

Severity: 1—Error

Cause: This occurs after a cluster is split into multiple cluster islands (usually due to a network link failure.) When a link is restored and the corresponding island seniors see each other, the panic protocol is initiated to resolve the conflict.

Action: If this issue occurs frequently, the root cause of the cluster split should be investigated.

A potential communication problem has been detected. A packet has failed to be delivered (or acknowledged) after %n1 seconds, although other packets were acknowledged by the same cluster member (Member(%s1)) to this member (Member(%s2)) as recently as %n2 seconds ago. Possible causes include network failure, poor thread scheduling (see FAQ if running on Windows), an

extremely overloaded server, a server that is attempting to run its processes using swap space, and unreasonably lengthy GC times.

Parameters:

%n1—The number of seconds a packet has failed to be delivered or acknowledged;

%s1—the recipient of the packets indicated in the message;

%s2 —the sender of the packets indicated in the message;

%n2—the number of seconds since a packet was delivered successfully between the two members indicated above.

Severity: 2—Warning

Cause: Possible causes are indicated in the text of the message.

Action: If this issue occurs frequently, the root cause should be investigated.

Node %s1 is not allowed to create a new cluster; WKA list: \[%s2\]

Parameters:

%s1—Address of node attempting to join cluster;

%s2—List of WKA addresses.

Severity: 1—Error

Cause: The cluster is configured to use WKA, and there are no nodes present in the cluster that are in the WKA list.

Action: Ensure that at least one node in the WKA list exists in the cluster, or add this node's address to the WKA list.

This member is configured with a compatible but different WKA list than the senior Member(%s). It is strongly recommended to use the same WKA list for all cluster members.

Parameters: %s—the senior node of the cluster.

Severity: 2—Warning

Cause: The WKA list on this node is different than the WKA list on the senior node.

Action: Ensure that every node in the cluster has the same WKA list.

UnicastUdpSocket failed to set receive buffer size to %n1 packets (%n2 bytes); actual size is %n3 packets (%n4 bytes). Consult your OS documentation regarding increasing the maximum socket buffer size. Proceeding with the actual value may cause sub-optimal performance.

Parameters:

%n1—the number of packets that will fit in the buffer that Coherence attempted to allocate;

%n2—the size of the buffer Coherence attempted to allocate;

%n3—the number of packets that will fit in the actual allocated buffer size;

%n4—the actual size of the allocated buffer.

Severity: 2—Warning

Cause: See *Performance Tuning* in the *Developer's Guide for Oracle Coherence*

Action: See *Performance Tuning* in the *Developer's Guide for Oracle Coherence*

Index

B

backup distribution algorithm, 1-2
BigDecimal, 1-2
BigInteger, 1-2

C

cache-config.dtd, 1-2
CacheFactory.release method, 1-1
CacheMbeans, 1-1
CacheStore, 2-2
concurrency issue, 1-1
concurrent read operations, 1-2
ConditionalPut, 1-3
ConfigurableAddressProvider, 1-2, 1-3
ConfigurablePofContext, 1-1
conflict resolution protocol, 1-1
ConnectionMBeans, 1-1
ContextClassLoader, 1-2
ContinuousQueryCache.release method, 1-2
createMapTrigger, 2-2
createTriggerListener, 2-2

D

deadlock, 1-1, 1-2, 1-3
default serializer type, 1-1
DNS resolution, 1-2
DNS resolution failures, 1-2, 1-3
DynamicMBeans, 1-1

E

EventDispatcher, 1-2
Extend Clients, 2-4

G

garbage collection (GC), 2-3
GreaterFilter, 1-2

H

hashCode, 1-1, 1-2

I

IllegalStateException, 1-2
indexing, 1-1
InvocationService startup, 1-1

J

JVM sizing, 2-3
 Cache Servers, 2-3
 Running with Large Heaps, 2-4
 Storage Ratios, 2-4
 Utilizing Available System Memory, 2-4

L

LessFilter, 1-2
LikeFilter, 1-2, 1-3
LocalCache, 1-2, 1-3
log and error messages, 2-5

M

MapEvents, 1-2
MBeanServer, 1-1
multiple requests, 1-1

N

NotFiter
 evaluateEntry, 1-3
NPE, 1-3

P

partition distribution, 1-1
POF serialization, 1-2
PofExtractor, 1-1, 1-2
proxy-scheme, 2-2
ProxyService, 1-2

Q

query performance, 1-2

R

RawDate, 1-3
RawTime, 1-3
ReadWriteBackingMap, 1-2

S

session persistence, 1-1
session reaper, 1-1
SimplePofPath, 1-1, 1-2
SymmetricEncryptionFilters, 1-2
SynchronizedDictionary, 1-2

T

task-hung-threshold, 2-2
TCMP Clients, 2-4
TcpAcceptor, 1-2
TcpInitiator, 1-2, 1-3
TcpReader, 1-2, 1-3
thread locking, 1-1
thread-count, 2-2
time_t, 1-3

U

unreachable addresses, 1-2

V

verbose
gc, 2-3

W

well-known-addresses, 2-2
wildcard, 1-2
WKA list, 1-2
write-behind, 1-2, 2-2

X

Xloggc, 2-3