

**DO NOT EDIT. This is not the
active copy.**

Java™ Caching API

The Java Caching API is an API for interacting with caching systems from Java programs

Second Public Review Draft

JSR107 Expert Group

Specification Leads:
Greg Luck, Terracotta
Brian Oliver, Oracle Corporation

6 August 2013

Comments to: jsr107@googlegroups.com

License

ORACLE IS WILLING TO LICENSE THIS SPECIFICATION TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS LICENSE AGREEMENT ("AGREEMENT"). PLEASE READ THE TERMS AND CONDITIONS OF THIS AGREEMENT CAREFULLY. BY DOWNLOADING THIS SPECIFICATION, YOU ACCEPT THE TERMS AND CONDITIONS OF THIS AGREEMENT. IF YOU ARE NOT WILLING TO BE BOUND BY THEM, SELECT THE "DECLINE" BUTTON AT THE BOTTOM OF THIS PAGE AND THE DOWNLOADING PROCESS WILL NOT CONTINUE.

Specification: JSR-107 JCACHE - Java Temporary Caching API ("Specification")

Version: Second Public Draft

Status: Public Draft Review

Release: 27 June 2013

Copyright 2013 Oracle America, Inc.
500 Oracle Parkway, Redwood City, California 94065, U.S.A.

Copyright 2013 Greg Luck.

All rights reserved.

NOTICE

The Specification is protected by copyright and the information described therein may be protected by one or more U.S. patents, foreign patents, or pending applications. Except as provided under the following license, no part of the Specification may be reproduced in any form by any means without the prior written authorization of Oracle America, Inc. ("Oracle") and its licensors, if any. Any use of the Specification and the information described therein will be governed by the terms and conditions of this Agreement.

Subject to the terms and conditions of this license, including your compliance with Paragraphs 1 and 2 below, Oracle hereby grants you a fully-paid, non-exclusive, non-transferable, limited license (without the right to sublicense) under Oracle's intellectual property rights to:

1. Review the Specification for the purposes of evaluation. This includes: (i) developing implementations of the Specification for your internal, non-commercial use; (ii) discussing the Specification with any third party; and (iii) excerpting brief portions of the Specification in oral or written communications which discuss the Specification provided that such excerpts do not in the aggregate constitute a significant portion of the Technology.
2. Distribute implementations of the Specification to third parties for their testing and evaluation use, provided that any such implementation:
 - (i) does not modify, subset, superset or otherwise extend the Licensor Name Space, or include any public or protected packages, classes, Java interfaces, fields or methods within the Licensor Name Space other than those required/authorized by the Specification or Specifications being implemented;
 - (ii) is clearly and prominently marked with the word "UNTESTED" or "EARLY ACCESS" or "INCOMPATIBLE" or "UNSTABLE" or "BETA" in any list of available builds and in proximity to every link initiating its download, where the list or link is under Licensee's control; and
 - (iii) includes the following notice:
"This is an implementation of an early-draft specification developed under the Java Community

Process (JCP) and is made available for testing and evaluation purposes only. The code is not compatible with any specification of the JCP."

The grant set forth above concerning your distribution of implementations of the specification is contingent upon your agreement to terminate development and distribution of your "early draft" implementation as soon as feasible following final completion of the specification. If you fail to do so, the foregoing grant shall be considered null and void.

No provision of this Agreement shall be understood to restrict your ability to make and distribute to third parties applications written to the Specification.

Other than this limited license, you acquire no right, title or interest in or to the Specification or any other Oracle intellectual property, and the Specification may only be used in accordance with the license terms set forth herein. This license will expire on the earlier of: (a) two (2) years from the date of Release listed above; (b) the date on which the final version of the Specification is publicly released; or (c) the date on which the Java Specification Request (JSR) to which the Specification corresponds is withdrawn. In addition, this license will terminate immediately without notice from Oracle if you fail to comply with any provision of this license. Upon termination, you must cease use of or destroy the Specification.

"Licensor Name Space" means the public class or interface declarations whose names begin with "java", "javax", "com.oracle" or their equivalents in any subsequent naming convention adopted by Oracle through the Java Community Process, or any recognized successors or replacements thereof

TRADEMARKS

No right, title, or interest in or to any trademarks, service marks, or trade names of Oracle or Oracle's licensors is granted hereunder. Oracle, the Oracle logo, Java are trademarks or registered trademarks of Oracle USA, Inc. in the U.S. and other countries.

DISCLAIMER OF WARRANTIES

THE SPECIFICATION IS PROVIDED "AS IS" AND IS EXPERIMENTAL AND MAY CONTAIN DEFECTS OR DEFICIENCIES WHICH CANNOT OR WILL NOT BE CORRECTED BY ORACLE. ORACLE MAKES NO REPRESENTATIONS OR WARRANTIES, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT THAT THE CONTENTS OF THE SPECIFICATION ARE

SUITABLE FOR ANY PURPOSE OR THAT ANY PRACTICE OR IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADE SECRETS OR OTHER RIGHTS. This document does not represent any commitment to release or implement any portion of the Specification in any product.

THE SPECIFICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION THEREIN; THESE CHANGES WILL BE INCORPORATED INTO NEW VERSIONS OF THE SPECIFICATION, IF ANY. ORACLE MAY MAKE IMPROVEMENTS AND/OR CHANGES TO THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THE SPECIFICATION AT ANY TIME. Any use of such changes in the Specification will be governed by the then-current license for the applicable version of the Specification.

LIMITATION OF LIABILITY

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ORACLE OR ITS LICENSORS BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION, LOST REVENUE, PROFITS OR DATA, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL OR PUNITIVE

DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF OR RELATED TO ANY FURNISHING, PRACTICING, MODIFYING OR ANY USE OF THE SPECIFICATION, EVEN IF ORACLE AND/OR ITS LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

You will hold Oracle (and its licensors) harmless from any claims based on your use of the Specification for any purposes other than the limited right of evaluation as described above, and from any claims that later versions or releases of any Specification furnished to you are incompatible with the Specification provided to you under this license.

RESTRICTED RIGHTS LEGEND

If this Software is being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), then the Government's rights in the Software and accompanying documentation shall be only as set forth in this license; this is in accordance with 48 C.F.R. 227.7201 through 227.7202-4 (for Department of Defense (DoD) acquisitions) and with 48 C.F.R. 2.101 and 12.212 (for non-DoD acquisitions).

REPORT

You may wish to report any ambiguities, inconsistencies or inaccuracies you may find in connection with your evaluation of the Specification ("Feedback"). To the extent that you provide Oracle with any Feedback, you hereby: (i) agree that such Feedback is provided on a non-proprietary and non-confidential basis, and (ii) grant Oracle a perpetual, non-exclusive, worldwide, fully paid-up, irrevocable license, with the right to sublicense through multiple levels of sublicensees, to incorporate, disclose, and use without limitation the Feedback for any purpose related to the Specification and future versions, implementations, and test suites thereof.

GENERAL TERMS

Any action related to this Agreement will be governed by California law and controlling U.S. federal law. The U.N. Convention for the International Sale of Goods and the choice of law rules of any jurisdiction will not apply.

The Specification is subject to U.S. export control laws and may be subject to export or import regulations in other countries. Licensee agrees to comply strictly with all such laws and regulations and acknowledges that it has the responsibility to obtain such licenses to export, re-export or import as may be required after delivery to Licensee.

This Agreement is the parties' entire agreement relating to its subject matter. It supersedes all prior or contemporaneous oral or written communications, proposals, conditions, representations and warranties and prevails over any conflicting or additional terms of any quote, order, acknowledgment, or other communication between the parties relating to its subject matter during the term of this Agreement. No modification to this Agreement will be binding, unless in writing and signed by an authorized representative of each party.

Contents

[Introduction](#)

[Overview](#)

[What is Caching?](#)

[Objectives](#)

[Non-Objectives](#)

[Java SE and Java EE support](#)

[Package](#)

[Optional Features](#)

[Specification Status](#)

[Document Conventions](#)

[Expert Group Members](#)

[Acknowledgements](#)

[Fundamentals](#)

[Core Concepts](#)

[Caches and Maps](#)

[Consistency](#)

[Default Consistency](#)

[Transactional Consistency](#)

[Further Consistency Modes](#)

[Cache Topologies](#)

[A Simple Example](#)

[CacheManagers](#)

[Acquiring a default CacheManager](#)

[Configuring Caches](#)

[Cache Names and Cache Scoping](#)

[Acquiring Caches](#)

[Cache and CacheManager Lifecycle](#)

[Closing Caches](#)

[Destroying Caches](#)

[Closing CacheManagers](#)

[ClassLoading](#)

[Caches](#)

[Cache Type-Safety](#)

[Compile-time Type-Safety](#)

[Example 1](#)

[Runtime Type-Safety](#)

[Example 2](#)

[Expiry Policies](#)

[Example 1](#)

[Integration](#)

[Cache Loading](#)

- [Example 1](#)
- [Read-Through Caching](#)
- [Write-Through Caching](#)
- [Cache Entry Listeners](#)
 - [Events and Event Types](#)
 - [CacheEntryListeners](#)
 - [Registration of Listeners](#)
 - [Invocation of Listeners](#)
- [Entry Processors](#)
- [Caching Providers](#)
 - [CacheManager Identity and Configuration](#)
- [Transactions](#)
 - [All or nothing](#)
 - [XA Transactions](#)
 - [Example 1](#)
 - [Enlistment](#)
 - [Recovery](#)
 - [Local Transactions](#)
 - [Recovery](#)
- [Isolation Levels](#)
 - [READ_COMMITTED](#)
 - [READ_UNCOMMITTED](#)
 - [SERIALIZABLE](#)
 - [REPEATABLE_READ](#)
- [Caching Annotations](#)
 - [Annotations](#)
 - [@CacheDefaults](#)
 - [Example 1](#)
 - [@CacheResult](#)
 - [@CachePut](#)
 - [@CacheRemoveEntry](#)
 - [@CacheRemoveAll](#)
 - [@CacheKey](#)
 - [@CacheValue](#)
 - [Example 2](#)
 - [Cache Resolution](#)
 - [Cache Name](#)
 - [CacheResolverFactory](#)
 - [CacheResolver](#)
 - [Key Generation](#)
 - [Annotation Support Classes](#)
 - [CacheMethodDetails](#)
 - [CacheInvocationContext](#)
 - [CacheKeyInvocationContext](#)
 - [CacheInvocationParameter](#)

[GeneratedCacheKey](#)
[Annotations Interactions](#)
[Annotation Inheritance and Ordering](#)
[Multiple Annotations](#)
[Transactions](#)
[Management](#)
[Enabling and Disabling](#)
[MXBean Definitions](#)
[Accessing Management Information](#)
[Example 1](#)
[Statistics Effects of Cache Operations](#)
[Portability Recommendations](#)
[Glossary](#)
[Bibliography](#)
[Appendix A - Revision History](#)
[Early Draft 1](#)
[Public Review Draft](#)
[Second Public Review Draft](#)

Introduction

This specification describes the objectives and functionality of the Java Caching Application Programming Interface (“API”).

The Java Caching API provides a common way for Java programs to create, access, update and remove entries from caches.

Overview

Caching is a proven technique for dramatically increasing the performance and scalability of applications.

Caching involves keeping a temporary copy of information in a low-latency structure for some period of time so that future requests for the same information may be performed faster.

Applications that repetitively make use of information which is either expensive to create or access will typically benefit from caching. For example consider a servlet that creates a web page containing information obtained from multiple databases, network servers and expensive computations; this information might be reusable for the creation of later web pages, and if so, using caching to reuse previously created information can reduce page construction time.

The Java Caching API provides a common way for applications to use and adopt caching thus allowing developers to focus on application development and avoid the burden of implementing caches themselves. This specification defines caching terminology, semantics and a corresponding set of Java interfaces.

Caching products that implement the Java Caching API do so by supplying a Caching Provider that implements the Java Caching API interfaces.

What is Caching?

The term Caching is ubiquitous in computing. In the context of application design it is often used to describe the technique whereby application developers utilize a separate in-memory or low-latency data-structure, a Cache, to temporarily store, or cache, a copy of or reference to information that an application may reuse at some later point in time, thus alleviating the cost to re-access or re-create it.

In the context of the Java Caching API the term Caching describes the technique whereby Java developers use a Caching Provider to temporarily cache Java objects.

It is often assumed that information from a database is being cached. This however is not a requirement of caching. Fundamentally any information that is expensive or time consuming to produce or access can be stored in a cache. Some common use cases are:

- *client side caching of Web service calls*
- *caching of expensive computations such as rendered images*
- *caching of data*
- *servlet response caching*

- *caching of domain object graphs*

Objectives

The Java Caching API's objectives are to:

- provide applications with caching functionality, in particular the ability to cache Java objects;
- define a common set of caching concepts and facilities;
- minimize the number of concepts Java developers need to learn to adopt caching;
- maximize the portability of applications that use caching between caching implementations;
- support both in-process and distributed cache implementations;
- support caching Java objects by-value and optionally by-reference;
- define runtime cache annotations in accordance with JSR-175: A Metadata Facility for the Java Programming Language; so that Java developers making use of optionally provided annotation processors may declaratively specify application caching requirements; and
- define the semantics of optionally supported transactional caches in the context of local and XA transactions as defined by JTA 2.0^[10].

Non-Objectives

The Java Caching API does not address:

- Resource and Memory Constraint Configuration - While many caching implementations provide support for constraining the amount of resources caches may use at runtime, this specification does not define how this functionality is configured or represented. This specification does however define a standard mechanism for developers to specify how long information should be available to be cached.
- Cache Storage and Topology - This specification does not specify how caching implementations store or represent information that is cached.
- Administration - This specification does not specify how caches are administered. It does define mechanisms to programmatically configure caches and investigate cache statistics via Java Management Extensions (JMX).
- Security - This specification does not specify how cache content may be secured or how access and operations on caches can be controlled.
- External Resource Synchronization - This specification does not specify how an application or caching implementations should keep caches and external resource content synchronized.

While developers may utilize read-through and write-through techniques as provided by the specification, these techniques are only effective when a cache is the only application mutating an external resource. Outside of this scenario cache synchronization can't be guaranteed.

Java SE and Java EE support

The Java Caching API is designed to be suitable for use by applications using the Standard and Enterprise Editions, versions 6 or newer.

A caching implementation:

- may choose to only work on a higher version of Java.
- may support its use by applications using Java EE, however this specification does not specify any standard for how that may be done.

Package

The top level package name for the Java Caching API is `javax.cache`.

Optional Features

All features in this specification are mandatory except for those enumerated in the `OptionalFeature` enum:

- `transactions`
- `storeByReference`

If implemented, these must be implemented exactly as described in this specification.

A developer may determine which of the optional features have been implemented by a caching provider using the capabilities API. Given a `CachingProvider` instance, call `cachingProvider.isSupported(OptionalFeature feature)`.

Some optional features only make sense only in some contexts. For example, `storeByReference` is generally not supported or supportable by distributed caching topologies.

Optional features allow for a caching implementation to support the specification without necessarily supporting all the features, and allows end users and frameworks to discover what the features are so they can dynamically configure appropriate usage.

Specification Status

This version is the second public review draft.

The latest API can be found online at:

<https://github.com/jsr107/jsr107spec>

The reference implementation can be obtained from:

<https://github.com/jsr107/RI>

Finally the TCK can be obtained from:

<https://github.com/jsr107/jsr107tck>

The expert group seeks feedback from the community on any aspect of this specification, please send comments to:

jsr-107-comments@jcp.org

or, for a publicly readable forum to:

jsr107@googlegroups.com

Document Conventions

The regular Arial (11 point) font is used for information that is normative for this specification.

The italic Arial (11 point) font is used for paragraphs that contain non-normative information, such as notes describing typical use, or notes clarifying the text with prescriptive specification.

The Courier New (11 Point) font is used for inline code descriptions. Java code, examples and sample data fragments also use the Courier New font. There are formatted as below (in 10 point font):

```
package com.example.hello;
public class Hello {
    public static void main(String args[] {
        System.out.println("Hello Worlds");
    }
}
```

In addition, the keywords 'MUST', 'MUST NOT', 'REQUIRED', 'SHALL', 'SHALL NOT', 'SHOULD', 'SHOULD NOT', 'RECOMMENDED', 'MAY', and 'OPTIONAL' in this document are to be interpreted as described in RFC 2119.

Expert Group Members

This specification is being developed under the Java Community Process v2.9.

Leading experts throughout the entire Java community have come together to build this Java caching standard.

The following are expert group members:

- Greg Luck
- Brian Oliver, Oracle
- Cameron Purdy, Oracle
- Manik Surtani, Red Hat, Inc.
- Nikita Ivanov, Grid Gain
- Chris Berry
- Jon Stevens
- Rick Hightower
- Ben Cotton, Credit Suisse
- David Mossakowski, Citigroup
- Bongjae Chang
- Steve Millidge
- Gabe Montero, IBM
- Eric Dalquist
- Pete Muir, Red Hat, Inc.
- William Newport, Goldman Sachs
- Ryan Gardner, Dealer.com
- Chris Dennis, Terracotta, Inc.

The following are official observers:

- Linda DeMichiel, Oracle

Acknowledgements

During the course of the JSR we have received many excellent suggestions on the JSR mailing lists. Thanks to those people.

Fundamentals

Core Concepts

The Java Caching API defines five core interfaces: `CachingProvider`, `CacheManager`, `Cache`, `Entry` and `Expiry`.

A `CachingProvider` defines the mechanism to establish, configure, acquire, manage and control zero or more `CacheManagers`. An application may access and use zero or more `CachingProviders` at runtime.

A `CacheManager` defines the mechanism to establish, configure, acquire, manage and control zero or more uniquely named `Caches` all within the context of the said `CacheManager`. A `CacheManager` is owned by a single `CachingProvider`.

A `Cache` is a Map-like data-structure that permits the temporary storage of Key-based Values. A `Cache` is owned by a single `CacheManager`.

An `Entry` is a single key-value pair stored by a `Cache`.

Each entry stored by a cache has a defined duration called the Expiry Duration, during which they may be accessed, updated and removed. Once this duration has passed, the entry is said to be Expired. Once expired, entries are no longer available to be accessed, updated or removed, just as if they never existed in a cache. The function that defines the expiry duration for entries is called an Expiry Policy.

Caches and Maps

While Caches and Maps share somewhat similar APIs, Caches are not Maps and Maps are not Caches. The following section outlines the main differences.

Like Map-based data-structures:

- Cache values are stored and accessed through an associated key.
- Each key may only be associated with a single value in a Cache.
- Custom key classes must implement or override both the `Object.hashCode` and `Object.equals` methods.

There is no requirement for a caching implementation to call the `Object.hashCode` and `Object.equals` methods but it cannot be assumed they won't be required or called.

Unlike Map-based data-structures:

- Cache keys and values must not be null.

Any attempt to use null for keys or values will result in a NullPointerException being thrown, regardless of the use.

- Entries may expire.

The process of ensuring Entries are no longer available to an application because they are no longer considered valid is called "expiry"

While entries have a defined expiry duration, entries may in fact be expired early, eagerly and at any time as part of a caching implementation requirement to carefully manage system resources. There is no guarantee that a recently stored entry will be accessible any time in the future, regardless of the expiry duration configured.

- Entries may be evicted.

Caches are typically not configured to store an entire data set. Instead they are often used to store a small, often used subset of the entire dataset.

To maintain that the size of a Cache doesn't consume resources without bound, a Cache implementation may define a policy to constrain the amount of resources a Cache may use at runtime by removing certain entries when a resource limit is exceeded.

The process of removing entries from a Cache to when the Cache has exceeded a resource limit is called "eviction". When an Entry is removed from a Cache due to resource constraints, it is said to be "evicted".

While the specification does not define the capacity of a cache, a sensible implementation will define mechanisms to represent desired capacity limits, together with suitable strategies to choose and evict entries once that capacity has been reached. For example: the LRU eviction strategy attempts to evict Least-Recently-Used entries.

The reason that capacity is not defined in the specification is because:

- implementations may use tiered storage and define capacity per tier, rather than as an overall capacity. Consequently a single capacity number would be ambiguous

- implementations may define capacity in terms of bytes rather than entry count on each tier

- the relative cost of entries in terms of memory used is directly related to the internal representation of the implementation of an Entry at runtime.

- Implementations may require Keys and Values to be serializable in some manner for some Cache Topologies.

- Caches may be configured to control how entries are stored, by default Store-By-Value with an option of Store-By-Reference.
- Implementations may optionally allow operations to be performed on a Cache as part of a JTA Transaction.

Consistency

Consistency refers to the guarantees that exist around concurrent cache mutation and visibility of said mutations when multiple threads are using a cache.

Default Consistency

In default consistency, consistency for most cache operations is described as if there exists a locking mechanism on each key. If a cache operation gets an exclusive read and write lock on a key, then all subsequent operations on that key will block until that lock is released. The consequences are that operations performed by a thread happen-before read or mutation operations performed by another thread, including threads in different Java Virtual Machines.

This can be understood as a pessimistic locking approach. Lock, mutate and unlock.

The operations which follow a different convention are the “CAS” or Compare and Swap cache operations, which only apply a mutation where the current state matches the argument given. Multiple threads calling these methods are free to compete to apply these changes i.e. as if they share a lock. These are:

- `boolean putIfAbsent(K key, V value);`
- `boolean remove(K key, V oldValue);`
- `boolean replace(K key, V oldValue, V newValue);`
- `boolean replace(K key, V value);`
- `V getAndReplace(K key, V value);`

This can be understood as an optimistic locking approach. If my understanding of the current state is correct, then apply the mutation, otherwise fail. They are known as CAS, after the CPU instructions that also operate in this way.

As these methods must interact with other cache operations acting as if they had an exclusive lock, the CAS methods cannot write new values without acting as if they also had an exclusive lock.

As a result, while the CAS methods can allow a higher level of concurrency they will be held up by the non-CAS methods.

Transactional Consistency

Where a cache is transactional it will take on the semantics of the configured transaction Isolation Level.

Transactional consistency is enabled by calling `setTransactions(IsolationLevel level,`

Mode mode) on MutableConfiguration.

See the Transactions chapter for more details.

Further Consistency Modes

An implementation may support additional consistency models.

Examples are last one wins, eventual consistency, and user control over locking.

Cache Topologies

While the specification does not mandate particular cache topologies, it is cognizant that Cache entries may well be stored locally and/or distributed across multiple processes. Implementation may choose to support neither, one, both and/or other topologies.

This notion is expressed in the specification in a number of ways:

- Most mutative methods have a version with a void or low cost return types. So while Map has `V put(K key, V value)` Cache has `void put(K key, V value)`.

Versions with a more expensive return type are also provided. An example is the `V getAndPut(K key, V value)` method on Cache. It returns the old value like map does.

- By having creation semantics which do not assume in-process. Configuration is `Serializable` so that it can be sent over the network. Developers may define implementations of `CacheEntryListener`, `ExpiryPolicy`, `CacheEntryFilter`, `CacheWriter` and `CacheLoader` and associate them with a cache. To support distributed topologies, a developer defines a Factory for their creation rather than the instance. The Factory interface is `Serializable`.
- Use of `Iterable` throughout for return types and parameters that might be large. Methods which return an entire collection such as the Map method `keySet()` can be problematic. A Cache may be so large that a key set may not fit in available memory and it might also be very network inefficient. Cache, the listener methods on `CacheEntryListener`'s subinterfaces, and the batch methods on `CacheLoader` and `CacheWriter` all use `Iterable`.
- No assumption is made as to where implementations of `CacheEntryListener`, `ExpiryPolicy`, `CacheEntryFilter`, `CacheWriter` and `CacheLoader` are instantiated and executed.

In a distributed implementation these may all reside close to the data rather than in process with the application.

- `CachingProvider.getCacheManager(URI uri, ClassLoader classLoader)` returns a `CacheManager` with a specific `ClassLoader` and URI. *This enables implementations to instantiate multiple instances.*

A Simple Example

This simple example creates a default `CacheManager`, configures a cache on it called “simpleCache” with a key type of `String` and a value type of `Integer` and an expiry of one hour and then performs a put and a get.

```
//resolve a cache manager
CachingProvider cachingProvider = Caching.getCachingProvider();
CacheManager cacheManager = cachingProvider.getCacheManager();

//configure the cache
MutableConfiguration<String, Integer> config =
    new MutableConfiguration<String, Integer>();
config.setStoreByValue(false)
    .setTypes(String.class, Integer.class)
    .setExpiryPolicyFactory(AccessedExpiryPolicy.factoryOf(ONE_HOUR))
    .setStatisticsEnabled(true);

//create the cache
cacheManager.createCache("simpleCache", config);

//get the cache
Cache<String, Integer> cache = cacheManager.getCache("simpleCache",
    String.class, Integer.class);

//cache operations
String key = "key";
Integer value1 = 1;
cache.put("key", value1);
Integer value2 = cache.get(key);
assertEquals(value1, value2);
cache.remove(key);
assertNull(cache.get(key));
```

Where the default `CachingProvider` and default `CacheManager` are being used, there is a static convenience method for getting a `Cache`, `Caching.getCache`:

```
//get the cache
Cache<String, Integer> cache = Caching.getCache("simpleCache",
    String.class, Integer.class);
```

CacheManagers

CacheManagers are a core concept of the Java Caching API. It is through CacheManagers that developers interact with caches.

A `CacheManager` provides:

- A means of establishing and configuring uniquely named caches.
- A means of acquiring a cache given its uniquely configured name.
- A means of scoping uniquely named caches; caches of the same name but originating from different Cache Managers are considered different caches.
- A means of closing a cache so that it is no longer managed.
- A means of destroying a cache including all of its contents.
- The `ClassLoader` that caches will use, should they require it, for resolving and loading application classes.
- A means of iterating over the currently managed caches.
- A means to close the `CacheManager` together with all of the currently managed caches.
- A means to enable and disable statistics gathering for caches.
- A means to enable and disable JMX management of caches.
- A means of acquiring `CachingProvider` specific properties defined for the `CacheManager`.
- A means of acquiring a `UserTransaction` for the managed caches, assuming transactions are supported by the caching provider.
- A means of querying the capabilities and optional features supported by the `CachingProvider`.

The `CacheManager` interface is defined as follows:

```
/**
 * A {@link CacheManager} provides a means of establishing, configuring,
 * acquiring, closing and destroying uniquely named {@link Cache}s.
 * <p/>
 * {@link Cache}s produced and owned by a {@link CacheManager} typically share
 * common infrastructure, for example, a common {@link ClassLoader} and
```

```

* implementation specific {@link Properties}.
* <p/>
* Implementations of {@link CacheManager} may additionally provide and share
* external resources between the {@link Cache}s being managed, for example,
* the content of the managed {@link Cache}s may be stored in the same cluster.
* <p/>
* By default {@link CacheManager} instances are typically acquired through the
* use of a {@link CachingProvider}. Implementations however may additionally
* provide other mechanisms to create, acquire, manage and configure
* {@link CacheManager}s, including:
* <ul>
*   <li>making use of {@link java.util.ServiceLoader}s,</li>
*   <li>permitting the use of the <code>new</code> operator to create a
*     concrete implementation, </li>
*   <li>providing the construction through the use of one or more
*     builders, and</li>
*   <li>through the use of dependency injection.</li>
* </ul>
* <p/>
* The default {@link CacheManager} however can always be acquired using the
* default configured {@link CachingProvider} obtained by the {@link Caching}
* class. For example:
* <code>
*     CachingProvider provider = Caching.getCachingProvider();
*     CacheManager manager = provider.getCacheManager();
* </code>
* <p/>
* Within a Java process {@link CacheManager}s and the {@link Cache}s they
* manage are scoped and uniquely identified by a {@link URI}, the meaning of
* which is implementation specific. To obtain the default {@link URI},
* {@link ClassLoader} and {@link Properties} for an implementation, consult the
* {@link CachingProvider} class.
*
* @author Greg Luck
* @author Yannis Cosmadopoulos
* @author Brian Oliver
* @since 1.0
*
* @see Caching
* @see CachingProvider
* @see Cache
*/
public interface CacheManager extends Closeable {

    /**
     * Get the {@link CachingProvider} that created and is responsible for
     * the {@link CacheManager}.
     *
     * @return the CachingProvider or <code>null</code> if the {@link CacheManager}
     *         was created without using a {@link CachingProvider}
    */

```

```

 */
CachingProvider getCachingProvider();

/**
 * Get the URI of the {@link CacheManager}.
 *
 * @return the URI of the {@link CacheManager}
 */
URI getURI();

/**
 * Get the Properties that were used to create this {@link CacheManager}.
 *
 * @return the Properties used to create the {@link CacheManager}
 */
Properties getProperties();

/**
 * Ensures that a named {@link Cache} is being managed by the
 * {@link CacheManager}.
 * <p/>
 * If such a {@link Cache} is unknown to the {@link CacheManager}, one is
 * created according to the provided
 * {@link javax.cache.configuration.Configuration} after which it becomes
 * managed by the {@link CacheManager}.
 * <p/>
 * If such a {@link Cache} is known to the {@link CacheManager},
 * a CacheException is thrown.
 * <p/>
 * {@link javax.cache.configuration.Configuration}s provided to this method are
 * always validated within the context of the {@link CacheManager}.
 * <p/>
 * For example: Attempting to use a
 * {@link javax.cache.configuration.Configuration} requiring transactional
 * support with an implementation that does not support
 * transactions will result in an {@link UnsupportedOperationException}.
 * <p/>
 * Implementers should be aware that the
 * {@link javax.cache.configuration.Configuration} may be used to configure
 * other {@link Cache}s.
 * There's no requirement on the part of a developer to call this method for
 * each {@link Cache} an application may use. Implementations may support
 * the use of declarative mechanisms to pre-configure {@link Cache}s, thus
 * removing the requirement to configure them in an application. In such
 * circumstances a developer may simply call either the {@link #getCache(String)}
 * or {@link #getCache(String, Class, Class)} methods to acquire a
 * pre-configured {@link Cache}.
 *
 * @param cacheName the name of the {@link Cache}

```

```

* @param configuration the {@link javax.cache.configuration.Configuration}
*
*         to use if the {@link Cache} is known
* @throws IllegalStateException      if the {@link CacheManager}
*
*         {@link #isClosed()}
* @throws CacheException            if there was an error configuring the
*
*         {@link Cache},
*
*         which includes trying to create a
*
*         cache which already exists.
* @throws IllegalArgumentException    if the configuration is invalid
* @throws UnsupportedOperationException if the configuration specifies
*
*         an unsupported feature
* @throws NullPointerException      if the cache configuration or name
*
*         is null
*/
<K, V> Cache<K, V> createCache(String cacheName,
                               Configuration<K, V> configuration)
                               throws IllegalArgumentException;

/**
 * Looks up a managed {@link Cache} given it's name.
 * <p/>
 * This method must be used for {@link Cache}s that were configured with
 * runtime key and value types. Use {@link #getCache(String)} for
 * {@link Cache}s where these were not specified.
 * <p/>
 * Implementations must ensure that the key and value types are the same as
 * those configured for the {@link Cache} prior to returning from this method.
 * <p/>
 * Implementations may further perform type checking on cache mutation and
 * throw a {@link ClassCastException} if said checks fail.
 * <p/>
 * Implementations that support declarative mechanisms for pre-configuring
 * {@link Cache}s may return a pre-configured {@link Cache} instead of
 * <code>null</code>.
 *
 * @param cacheName the name of the managed {@link Cache} to acquire
 * @param keyType    the expected {@link Class} of the key
 * @param valueType  the expected {@link Class} of the value
 * @return the Cache or null if it does exist or can't be pre-configured
 * @throws IllegalStateException if the CacheManager is {@link #isClosed()}
 * @throws IllegalArgumentException if the specified key and/or value types are
 *
 *         incompatible with the configured cache.
 */
<K, V> Cache<K, V> getCache(String cacheName, Class<K> keyType, Class<V>
valueType);

/**
 * Looks up a managed {@link Cache} given it's name.
 * <p/>
 * This method must be used for {@link Cache}s that were not configured with

```

```

* runtime key and value types. Use {@link #getCache(String, Class, Class)} to
* acquire {@link Cache}s that were configured with specific runtime types.
* <p/>
* Implementations must check that no key and value types were specified
* when the cache was configured. If either the keyType or valueType of the
* configured cache are not their defaults then a {@link IllegalArgumentException}
* is thrown.
* <p/>
* Implementations that support declarative mechanisms for pre-configuring
* {@link Cache}s may return a pre-configured {@link Cache} instead of
* <code>null</code>.
*
* @param cacheName the name of the cache to look for
* @return the Cache or null if it does exist or can't be pre-configured
* @throws IllegalStateException if the CacheManager is {@link #isClosed()}
* @throws IllegalArgumentException if the {@link Cache} was configured with
* specific types, this method cannot be used
* @see #getCache(String, Class, Class)
*/
<K, V> Cache<K, V> getCache(String cacheName);

/**
* Obtains an {@link Iterable} over the names of {@link Cache}s managed by the
* {@link CacheManager}.
* <p/>
* {@link java.util.Iterator}s returned by the {@link Iterable} are immutable.
* Any modification of the {@link java.util.Iterator}, including remove, will
* raise an {@link IllegalStateException}. If the {@link Cache}s managed by
* the {@link CacheManager} change, the {@link Iterable} and
* associated {@link java.util.Iterator}s are not affected.
*
* @return an {@link Iterable} over the names of managed {@link Cache}s.
*/
Iterable<String> getCacheNames();

/**
* Destroys a specifically named and managed {@link Cache}. Once destroyed
* a new {@link Cache} of the same name but with a different {@link Configuration}
* may be configured.
* <p/>
* This is equivalent to the following sequence of method calls:
* <ol>
* <li>{@link javax.cache.Cache#clear()}</li>
* <li>{@link javax.cache.Cache#close()}</li>
* </ol>
* followed by allowing the name of the {@link Cache} to be used for other
* {@link Cache} configurations.
* <p/>
* From the time this method is called, the specified {@link Cache} is not
* available for operational use. An attempt to call an operational method on

```



```

* the {@link Cache} will throw an {@link IllegalStateException}.
*
* @param cacheName the cache name
* @throws IllegalStateException if the {@link Cache} is {@link #isClosed()}
* @throws NullPointerException if cacheName is null
*/
void destroyCache(String cacheName);

/**
 * Obtains a UserTransaction for transactional {@link Cache}s managed
 * by the {@link CacheManager}.
 *
 * @return the UserTransaction
 * @throws UnsupportedOperationException if JTA is not supported
 */
UserTransaction getUserTransaction();

/**
 * Enables or disables statistics gathering for a managed {@link Cache} at
 * runtime.
 * <p/>
 * Each cache's statistics object must be registered with an ObjectName that
 * is unique and has the following type and attributes:
 * <p/>
 * Type:
 * <code>javax.cache:type=CacheStatistics</code>
 * <p/>
 * Required Attributes:
 * <ul>
 * <li>CacheManager the name of the CacheManager
 * <li>Cache the name of the Cache
 * </ul>
 *
 * @param cacheName the name of the cache to register
 * @param enabled true to enable statistics, false to disable.
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws NullPointerException if cacheName is null
 */
void enableStatistics(String cacheName, boolean enabled);

/**
 * Controls whether management is enabled. If enabled the
 * {@link javax.cache.management.CacheMBean} for each cache is registered in
 * the platform MBean server. The platform MBeanServer is obtained using
 * {@link java.lang.management.ManagementFactory#getPlatformMBeanServer()}
 * <p/>
 * Management information includes the name and configuration information for
 * the cache.
 * <p/>
 * Each cache's management object must be registered with an ObjectName that

```

```

* is unique and has the following type and attributes:
* <p/>
* Type:
* <code>javax.cache:type=Cache</code>
* <p/>
* Required Attributes:
* <ul>
* <li>CacheManager the name of the CacheManager
* <li>Cache the name of the Cache
* </ul>
*
* @param cacheName the name of the cache to register
* @param enabled true to enable management, false to disable.
*/
void enableManagement(String cacheName, boolean enabled);

/**
* Closes the {@link CacheManager}.
* <p/>
* For each {@link Cache} managed by the {@link CacheManager}, the
* {@link javax.cache.Cache#close()} method will be invoked, in no guaranteed
* order.
* <p/>
* If a {@link javax.cache.Cache#close()} call throws an exception, the
* exception will be ignored.
* <p/>
* After executing this method, the {@link #isClosed()} method will return
* <code>>true</code>.
* <p/>
* All attempts to close a previously closed {@link CacheManager} will be
* ignored.
*/
void close();

/**
* Determines whether the {@link CacheManager} instance has been closed. A
* {@link CacheManager} is considered closed if;
* <ol>
* <li>the {@link #close()} method has been called</li>
* <li>the associated {@link #getCachingProvider()} has been closed, or</li>
* <li>the {@link CacheManager} has been closed using the associated
*     {@link #getCachingProvider()}</li>
* </ol>
* <p/>
* This method generally cannot be called to determine whether the
* {@link CacheManager} is valid or invalid. A typical client can determine
* that a {@link CacheManager} is invalid by catching any exceptions that
* might be thrown when an operation is attempted.
*
* @return true if this {@link CacheManager} instance is closed; false if it

```

```

    *           is still open
    */
boolean isClosed();

/**
 * Provides a standard mechanism to access the underlying concrete caching
 * implementation to provide access to further, proprietary features.
 * <p/>
 * If the provider's implementation does not support the specified class,
 * the {@link IllegalArgumentException} is thrown.
 *
 * @param clazz the proprietary class or interface of the underlying concrete
 *              {@link CacheManager}. It is this type which is returned.
 * @return an instance of the underlying concrete {@link CacheManager}
 * @throws IllegalArgumentException if the caching provider doesn't support the
specified class.
 */
<T> T unwrap(java.lang.Class<T> clazz);
}

```

Acquiring a default CacheManager

To ease adoption of the Java Caching API developers may acquire a default `CacheManager` from a default `CachingProvider` by using the `Caching` helper class. For example:

```

//acquire the default CachingProvider
CachingProvider provider = Caching.getCachingProvider();

//acquire the default CacheManager
CacheManager manager = provider.getCacheManager();

```

To acquire non-default or alternative configurations of `CacheManagers`, for example with custom `ClassLoaders` or caching implementor properties, developers should use one of the overloaded `CachingProvider` `getCacheManager` methods.

How to configure a `CachingProvider` is covered in the section on `CachingProviders`.

Configuring Caches

There are two approaches for configuring caches with `CacheManagers`:

- `CacheManagers` must allow applications to programmatically configure caches at runtime through the `CacheManager.createCache` method.
- `CacheManagers` may optionally provide mechanisms to declaratively configure caches for applications thus avoiding the need for applications to use the `createCache` method.

The mechanism(s) by which a `CacheManager` may allow the declarative definition of caches for an application is implementation dependent. One approach is to have a XML configuration file which configures a `CacheManager` and the `Caches` in it.

CacheManagers have the responsibility to validate Cache configurations that are provided by applications. Should a Cache configuration be invalid for a CacheManager, attempting to create the Cache will throw an IllegalArgumentException.

To ease configuration of caches the Java Caching API provides a concrete implementation of the javax.cache.configuration.Configuration interface called javax.cache.configuration.MutableConfiguration.

Interfaces and Classes related to cache configuration are defined in the javax.cache.configuration package.

Caching implementations may choose to provide additional implementations of the Configuration interface in order to provide implementation specific configuration.

To simplify programmatic configuration when using the MutableConfiguration class all setter methods return the MutableConfiguration instance thus allowing the class to be used in a fluent manner.

As applications may establish MutableConfiguration instances that are invalid for specific implementations of CacheManagers, MutableConfiguration instances do not perform validation.

Commonly used constructors and setter methods of the MutableConfiguration class are defined as follows:

```
/**
 * Constructs a default {@link MutableConfiguration}.
 */
public MutableConfiguration()

/**
 * A copy-constructor for a {@link MutableConfiguration}.
 *
 * @param configuration the {@link Configuration} from which to copy
 */
public MutableConfiguration(Configuration<K, V> configuration)

/**
 * Sets the expected type of keys and values for a {@link javax.cache.Cache}
 * configured with this {@link Configuration}. Setting both to <code>>null</code>
 * means type-safety checks are not required.
 *
 * @param keyType the expected key type
 * @param valueType the expected value type
 * @return the {@link MutableConfiguration} to permit fluent-style method calls
 */
public MutableConfiguration<K, V> setTypes(Class<K> keyType, Class<V> valueType)
```

```

/**
 * Add a configuration for a {@link javax.cache.event.CacheEntryListener}.
 *
 * @param listenerFactory the {@link javax.cache.event.CacheEntryListener}
 *                        {@link Factory}
 * @param filterFactory   the optional
 *                        {@link javax.cache.event.CacheEntryEventFilter}
 *                        {@link Factory}
 * @param isOldValueRequired if the old value is required for events with this
 *                        listenerFactory
 * @param isSynchronous    if the listenerFactory should block the thread
 *                        causing the event
 * @return the {@link MutableConfiguration} to permit fluent-style method calls
 */
public MutableConfiguration<K, V> addCacheEntryListenerConfiguration(
    Factory<? extends CacheEntryListener<? super K, ? super V>> listenerFactory,
    Factory<? extends CacheEntryEventFilter<? super K, ? super V>> filterFactory,
    boolean isOldValueRequired,
    boolean isSynchronous)

/**
 * Add a configuration for a {@link javax.cache.event.CacheEntryListener}.
 *
 * @param configuration the {@link CacheEntryListenerConfiguration}
 * @return the {@link MutableConfiguration} to permit fluent-style method calls
 */
public MutableConfiguration<K, V> addCacheEntryListenerConfiguration(
    CacheEntryListenerConfiguration<K, V> configuration)

/**
 * Set the {@link CacheLoader} factory.
 *
 * @param factory the {@link CacheLoader} {@link Factory}
 * @return the {@link MutableConfiguration} to permit fluent-style method calls
 */
public MutableConfiguration<K, V> setCacheLoaderFactory(Factory<? extends
    CacheLoader<K, V>> factory)

/**
 * Set the {@link CacheWriter} factory.
 *
 * @param factory the {@link CacheWriter} {@link Factory}
 * @return the {@link MutableConfiguration} to permit fluent-style method calls
 */
public MutableConfiguration<K, V> setCacheWriterFactory(Factory<? extends
    CacheWriter<? super K, ? super V>> factory) {

/**

```

```

* Set the {@link Factory} for the {@link ExpiryPolicy}. If null
* is specified the default {@link ExpiryPolicy} is used.
*
* @param factory the {@link ExpiryPolicy} {@link Factory}
* @return the {@link MutableConfiguration} to permit fluent-style method calls
*/
public MutableConfiguration<K, V> setExpiryPolicyFactory(Factory<? extends
    ExpiryPolicy<? super K, ? super V>> factory) {

/**
* Set the Transaction {@link IsolationLevel} and {@link Mode},
* which also sets {@link #isTransactionsEnabled()} to true.
*
* @param level the {@link IsolationLevel}
* @param mode the {@link Mode}
* @return the {@link MutableConfiguration} to permit fluent-style method calls
*/
public MutableConfiguration<K, V> setTransactions(IsolationLevel level,
    Mode mode) {

/**
* Set if read-through caching should be used.
* <p/>
* It is an invalid configuration to set this to true without specifying a
* {@link CacheLoader} {@link Factory}.
*
* @param isReadThrough true if read-through is required
* @return the {@link MutableConfiguration} to permit fluent-style method calls
*/
public MutableConfiguration<K, V> setReadThrough(boolean isReadThrough)

/**
* Set if write-through caching should be used.
* <p/>
* It is an invalid configuration to set this to true without specifying a
* {@link CacheWriter} {@link Factory}.
*
* @param isWriteThrough true if write-through is required
* @return the {@link MutableConfiguration} to permit fluent-style method calls
*/
public MutableConfiguration<K, V> setWriteThrough(boolean isWriteThrough)

/**
* Set if a configured cache should use store-by-value or store-by-reference
* semantics.
*
* @param isStoreByValue true if store-by-value is required,
* false for store-by-reference

```

```

    * @return the {@link MutableConfiguration} to permit fluent-style method calls
    */
public MutableConfiguration<K, V> setStoreByValue(boolean isStoreByValue)

/**
 * Sets whether statistics gathering is enabled on a cache.
 * <p/>
 * Statistics may be enabled or disabled at runtime via
 * {@link javax.cache.CacheManager#enableStatistics(String, boolean)}.
 *
 * @param enabled true to enable statistics, false to disable.
 * @return the {@link MutableConfiguration} to permit fluent-style method calls
 */
public MutableConfiguration<K, V> setStatisticsEnabled(boolean enabled)

/**
 * Sets whether management is enabled on a cache.
 * <p/>
 * Management may be enabled or disabled at runtime via
 * {@link javax.cache.CacheManager#enableManagement(String, boolean)}.
 *
 * @param enabled true to enable statistics, false to disable.
 * @return the {@link MutableConfiguration} to permit fluent-style method calls
 */
public MutableConfiguration<K, V> setManagementEnabled(boolean enabled)

```

Each of the configuration options provided by the `MutableConfiguration` class and thus the `javax.cache.configuration.Configuration` interface are discussed in depth in later sections of this document.

The following defaults are provided by a new instance of the `MutableConfiguration` class.

Configuration Option	Type	Default Value(s)
Key Type	<code>Class<?></code>	<code>null</code>
Value Type	<code>Class<?></code>	<code>null</code>
Cache Loader Factory	<code>Factory<CacheLoader<K, V>></code>	<code>null</code>
Cache Writer Factory	<code>Factory<CacheWriter<? super K, ? super V>></code>	<code>null</code>
Expiry Policy Factory	<code>Factory<ExpiryPolicy<? super K, ? super V>></code>	a factory producing an <code>EternalExpiryPolicy</code>
Read Through Enabled	<code>boolean</code>	<code>false</code>
Write Through Enabled	<code>boolean</code>	<code>false</code>

Cache Entry Listener Configuration	Iterable<CacheEntryListenerConfiguration<? super K, ? super V>>	an empty iteration
Statistics Enabled	boolean	false
Management Enabled	boolean	false
Transactions Enabled	boolean	false
Isolation Level	IsolationLevel	IsolationLevel.NONE
Transaction Mode	Mode	Mode.NONE

Cache Names and Cache Scoping

Caches are identified by their uniquely configured name in the scope of the `CacheManager` that was used to create or initially access them. This means that Caches with the same unique name which were created with different `CacheManagers`, with different URIs, are considered different Caches. In such circumstances the configuration and content of said Caches may be different.

While Cache names are represented as Strings, there are some restrictions and recommended naming conventions for portability. These are as follows:

- Cache Names starting with `java.` or `javax.` are reserved as internal platform Caches. These Caches will not be returned when attempting to iterate over the available Cache names using a `CacheManager`.
- Cache Names should not contain forward slashes (`/`) or full-colons (`:`) as these are often used within Java EE environments for JNDI-based lookups.
- Cache Names may use Unicode characters.

Acquiring Caches

There are two approaches for acquiring caches with `CacheManagers`:

- When a type-safe `Cache` is required, which is one that ensures that the correct and expected types of cache entries are used, an application should use the following `CacheManager` method:

```
<K, V> Cache<K, V> getCache(String cacheName,
                             Class<K> keyType,
                             Class<V> valueType)
```

- When an application is explicitly taking responsibility for cache entry type-safety, the following `CacheManager` method should be used:

```
<K, V> Cache<K, V> getCache(String cacheName);
```


For more information on Cache type-safety see the section on Cache Type-Safety.

A simple example of how to acquire a Cache from a CacheManager:

```
Cache<String, Integer> cache = cacheManager.getCache(
    "simpleCache", String.class, Integer.class);
```

Cache and CacheManager Lifecycle

All Cache and CacheManager instances operate in one of two possible states; opened or closed. When open, instances may be used operationally to make requests. For example; creating, updating, removing an entry or configuring, acquiring, closing, removing a cache and so on. When closed, any operational use of these instances will throw an `IllegalStateException`.

Closing Caches

Closing a Cache via a call to the `Cache.close()` method signals to the CacheManager that produced or owns the said Cache that it should no longer be managed. At this point in time the CacheManager:

- must close and release all resources being coordinated on behalf of the Cache by the CacheManager. This includes calling the `close` method on configured `CacheLoader`, `CacheWriter`, registered `CacheEntryListeners` and `ExpiryPolicy` instances that implement the `java.io.Closeable` interface,
- prevent events being delivered to configured `CacheEntryListeners` registered on the Cache,
- not return the name of the Cache when the `CacheManager.getCacheNames()` method is called.

Once closed any attempt to use an operational method on a Cache will throw an `IllegalStateException`.

Closing a Cache does not necessarily destroy the contents of a Cache. It simply signals to the owning CacheManager that the Cache is no longer required by the application and that future uses of a specific Cache instance should not be permitted. Depending on the implementation and Cache topology, eg: a storage-backed or distributed caches, the contents of a closed Cache may still be available and accessible by other applications or in fact via the Cache Manager that previously owned the Cache if an application calls `getCache` at some point in the future.

Destroying Caches

To destroy a Cache, release it from being managed and drop all of the cache entries, thus allowing a new cache, with the same name but possibly a different configuration to be created, the `CacheManager.destroyCache` method should be called.

```
/**
```

```

* Destroys a specifically named and managed {@link Cache}. Once destroyed
* a new {@link Cache} of the same name but with a different
* {@link Configuration} may be configured.
* <p/>
* This is equivalent to the following sequence of method calls:
* <ol>
*   <li>{@link javax.cache.Cache#clear()}</li>
*   <li>{@link javax.cache.Cache#close()}</li>
* </ol>
* followed by allowing the name of the {@link Cache} to be used for other
* {@link Cache} configurations.
* <p/>
* From the time this method is called, the specified {@link Cache} is not
* available for operational use. An attempt to call an operational method on
* the {@link Cache} will throw an {@link IllegalStateException}.
*
* @param cacheName the cache name
* @throws IllegalStateException if the {@link Cache} is {@link #isClosed()}
* @throws NullPointerException if cacheName is null
*/
void destroyCache(String cacheName);

```

Once destroyed:

- any attempt to use an operational method on instances of the said Cache will throw an `IllegalStateException`
- the destroyed Cache's name may be reused in a new cache by calling the `CacheManager.create` method, with the same or a different configuration.

Once destroyed a Cache is no longer available via a CacheManager. Destroying a Cache ensures that it is closed and all of the associated entries are no longer available by any application, both immediately and in the future, regardless of implementation or topology.

Closing CacheManagers

Closing a `CacheManager` via a call to the `CacheManager.close()` method or via the `CachingProvider.close(...)` methods has the effect of instructing a `CacheManager` to:

- close all of the Caches that it is currently managing, and
- release all resources that are currently being used to manage said Caches.

Once closed any attempt to use an operational method on a closed `CacheManager` or any of the Caches it was managing will throw an `IllegalStateException`.

After closing a CacheManager, another instance, possibly representing the previously managed Caches, may be acquired using the CachingProvider that originally produced the said CacheManager.

This is covered in the section on `CachingProviders`.

ClassLoadering

All Caches share the same `ClassLoader` that was configured for the `CacheManager` from which they were acquired when the said `CacheManager` was created.

To configure and acquire Caches that use different `ClassLoaders`, individual `CacheManagers` must be established to do so. For information on how to configure `CacheManagers`, consult the section on `CachingProviders`.

Caches

The primary artifact developers use to interact with a Cache is the `javax.cache.Cache` interface.

The `javax.cache.Cache` interface provides Map-like methods to enable access, update and remove access to Cache Entries.

The `javax.cache.Cache` interface is defined as follows:

```
package javax.cache;

import javax.cache.configuration.Configuration;
import javax.cache.integration.CompletionListener;
import java.io.Closeable;
import java.util.Iterator;
import java.util.Map;
import java.util.Set;

/**
 * A {@link Cache} is a Map-like data structure that provides temporary storage
 * of application data.
 * <p/>
 * Like {@link Map}s, {@link Cache}s
 * <ol>
 * <li>store key-value pairs, each referred to as an {@link Entry}</li>
 * <li>allow use Java Generics to improve application type-safety</li>
 * <li>are {@link Iterable}</li>
 * </ol>
 * <p/>
 * Unlike {@link Map}s, {@link Cache}s
 * <ol>
 * <li>do not allow null keys or values. Attempts to use null
 * will result in a {@link NullPointerException}</li>
 * <li>provide the ability to read values from a
 * {@link javax.cache.integration.CacheLoader} (read-through-caching)
 * when a value being requested is not in a cache</li>
 * <li>provide the ability to write values to a
 * {@link javax.cache.integration.CacheWriter} (write-through-caching)
 * when a value being created/updated/removed from a cache</li>
 * <li>provide the ability to observe cache entry changes</li>
 * <li>may capture and measure operational statistics</li>
 * <li>may be transactional</li>
 * </ol>
 * <p/>
 * A simple example of how to use a cache is:
 * <code>
 * String cacheName = "sampleCache";
 * CachingProvider provider = Caching.getCachingProvider();
 * CacheManager manager = provider.getCacheManager();
 * Cache<Integer, Date> cache = manager.getCache(cacheName,
 * Integer.class,
 * Date.class);
 * Date value1 = new Date();
 * Integer key = 1;
```

```

* cache.put(key, value1);
* Date value2 = cache.get(key);
* </code>
*
* @param <K> the type of key
* @param <V> the type of value
* @author Greg Luck
* @author Yannis Cosmadopoulos
* @author Brian Oliver
* @since 1.0
*/
public interface Cache<K, V> extends Iterable<Cache.Entry<K, V>>, Closeable {
/**
* Gets an entry from the cache.
* <p/>
* If the cache is configured read-through, and get would return null because
* the entry is missing from the cache, the Cache's
* {@link javax.cache.integration.CacheLoader} is called which will attempt
* to load the entry.
*
* @param key the key whose associated value is to be returned
* @return the element, or null, if it does not exist.
* @throws IllegalStateException if the cache is {@link #isClosed()}
* @throws NullPointerException if the key is null
* @throws CacheException if there is a problem fetching the value
* @throws ClassCastException if the implementation supports and is
* configured to perform runtime-type-checking,
* and the key type is incompatible with that
* which has been configured for the {@link Cache}
*/
V get(K key);

/**
* Gets a collection of entries from the {@link Cache}, returning them as
* {@link Map} of the values associated with the set of keys requested.
* <p/>
* If the cache is configured read-through, and a get would return null
* because an entry is missing from the cache, the Cache's
* {@link javax.cache.integration.CacheLoader} is called which will attempt
* to load the entry. This is done for each key in the set for which this is
* the case. If an entry cannot be loaded for a given key, the key will not be
* present in the returned Map.
* <p/>
*
* @param keys The keys whose associated values are to be returned.
* @return A map of entries that were found for the given keys. Keys not found
* in the cache are not in the returned map.
* @throws NullPointerException if keys is null or if keys contains a null
* @throws IllegalStateException if the cache is {@link #isClosed()}
* @throws CacheException if there is a problem fetching the values
* @throws ClassCastException if the implementation supports and is
* configured to perform runtime-type-checking,
* and any key type is incompatible with that
* which has been configured for the
* {@link Cache}
*/
}

```

```

Map<K, V> getAll(Set<? extends K> keys);

/**
 * Determines if the {@link Cache} contains an entry for the specified key.
 * <p/>
 * More formally, returns <tt>>true</tt> if and only if this cache contains a
 * mapping for a key <tt>k</tt> such that <tt>key.equals(k)</tt>.
 * (There can be at most one such mapping.)
 *
 * @param key key whose presence in this cache is to be tested.
 * @return <tt>>true</tt> if this map contains a mapping for the specified key
 * @throws NullPointerException if key is null
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException it there is a problem checking the mapping
 * @throws ClassCastException if the implementation supports and is
 * configured to perform runtime-type-checking,
 * and the key type is incompatible with that
 * which has been configured for the
 * {@link Cache}
 * @see java.util.Map#containsKey(Object)
 */
boolean containsKey(K key);

/**
 * Asynchronously loads the specified entries into the cache using the
 * configured {@link javax.cache.integration.CacheLoader} for the given keys.
 * <p/>
 * If an entry for a key already exists in the Cache, a value will be loaded
 * if and only if <code>replaceExistingValues</code> is true. If no loader
 * is configured for the cache, no objects will be loaded. If a problem is
 * encountered during the retrieving or loading of the objects,
 * an exception is provided to the {@link CompletionListener}. Once the
 * operation has completed, the specified CompletionListener is notified.
 * <p/>
 * Implementations may choose to load multiple keys from the provided
 * {@link Set} in parallel. Iteration however must not occur in parallel,
 * thus allow for non-thread-safe {@link Set}s to be used.
 * <p/>
 * The thread on which the completion listener is called is implementation
 * dependent. An implementation may also choose to serialize calls to
 * different CompletionListeners rather than use a thread per
 * CompletionListener.
 *
 * @param keys the keys to load
 * @param replaceExistingValues when true existing values in the Cache will
 * be replaced by those loaded from a CacheLoader
 * @param completionListener the CompletionListener (may be null)
 * @throws NullPointerException if keys is null or if keys contains a null.
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException thrown if there is a problem performing the
 * load. This may also be thrown on calling if
 * there are insufficient threads available to
 * perform the load.
 * @throws ClassCastException if the implementation supports and is
 * configured to perform runtime-type-checking,
 * and any key type is incompatible with that

```

```

*                                     which has been configured for the
*                                     {@link Cache}
*/
void loadAll(Set<? extends K> keys, boolean replaceExistingValues,
             CompletionListener completionListener);

/**
 * Associates the specified value with the specified key in the cache.
 * <p/>
 * If the {@link Cache} previously contained a mapping for the key, the old
 * value is replaced by the specified value. (A cache
 * <tt>c</tt> is said to contain a mapping for a key <tt>k</tt> if and only
 * if {@link #containsKey(Object) c.containsKey(k)} would return
 * <tt>>true</tt>.)
 *
 * @param key    key with which the specified value is to be associated
 * @param value  value to be associated with the specified key
 * @throws NullPointerException if key is null or if value is null
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      if there is a problem doing the put
 * @throws ClassCastException  if the implementation supports and is
 *                               configured to perform runtime-type-checking,
 *                               and the key or value types are incompatible
 *                               with those which have been configured for the
 *                               {@link Cache}
 * @see java.util.Map#put(Object, Object)
 * @see #getAndPut(Object, Object)
 * @see #getAndReplace(Object, Object)
 */
void put(K key, V value);

/**
 * Associates the specified value with the specified key in this cache,
 * returning an existing value if one existed.
 * <p/>
 * If the cache previously contained a mapping for
 * the key, the old value is replaced by the specified value. (A cache
 * <tt>c</tt> is said to contain a mapping for a key <tt>k</tt> if and only
 * if {@link #containsKey(Object) c.containsKey(k)} would return
 * <tt>>true</tt>.)
 * <p/>
 * The the previous value is returned, or null if there was no value associated
 * with the key previously.
 *
 * @param key    key with which the specified value is to be associated
 * @param value  value to be associated with the specified key
 * @return the value associated with the key at the start of the operation or
 *         null if none was associated
 * @throws NullPointerException if key is null or if value is null
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      if there is a problem doing the put
 * @throws ClassCastException  if the implementation supports and is
 *                               configured to perform runtime-type-checking,
 *                               and the key or value types are incompatible
 *                               with those which have been configured for the
 *                               {@link Cache}
 */

```

```

* @see #put(Object, Object)
* @see #getAndReplace(Object, Object)
*/
V getAndPut(K key, V value);

/**
 * Copies all of the entries from the specified map to the {@link Cache}.
 * <p/>
 * The effect of this call is equivalent to that of calling
 * {@link #put(Object, Object) put(k, v)} on this cache once for each mapping
 * from key <tt>k</tt> to value <tt>v</tt> in the specified map.
 * <p/>
 * The order in which the individual puts occur is undefined.
 * <p/>
 * The behavior of this operation is undefined if entries in the cache
 * corresponding to entries in the map are modified or removed while this
 * operation is in progress. or if map is modified while the operation is in
 * progress.
 * <p/>
 * In Default Consistency mode, individual puts occur atomically but not
 * the entire putAll. Listeners may observe individual updates.
 *
 * @param map mappings to be stored in this cache
 * @throws NullPointerException if map is null or if map contains null keys
 * or values.
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException if there is a problem doing the put.
 * @throws ClassCastException if the implementation supports and is
 * configured to perform runtime-type-checking,
 * and any of the key or value types are
 * incompatible with those that have been
 * configured for the {@link Cache}
 */
void putAll(java.util.Map<? extends K, ? extends V> map);

/**
 * Atomically associates the specified key with the given value if it is
 * not already associated with a value.
 * <p/>
 * This is equivalent to:
 * <code>
 * if (!cache.containsKey(key)) {}
 *     cache.put(key, value);
 *     return true;
 * } else {
 *     return false;
 * }
 * </code>
 * except that the action is performed atomically.
 *
 * @param key key with which the specified value is to be associated
 * @param value value to be associated with the specified key
 * @return true if a value was set.
 * @throws NullPointerException if key is null or value is null
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException if there is a problem doing the put

```



```

* @throws ClassCastException    if the implementation supports and is
*                               configured to perform runtime-type-checking,
*                               and the key or value types are incompatible
*                               with those that have been configured for the
*                               {@link Cache}
*/
boolean putIfAbsent(K key, V value);

/**
* Removes the mapping for a key from this cache if it is present.
* <p/>
* More formally, if this cache contains a mapping from key <tt>k</tt> to
* value <tt>v</tt> such that
* <code>(key==null ? k==null : key.equals(k))</code>, that mapping
* is removed. (The cache can contain at most one such mapping.)
* <p/>
* <p>Returns <tt>>true</tt> if this cache previously associated the key,
* or <tt>>false</tt> if the cache contained no mapping for the key.
* <p/>
* The cache will not contain a mapping for the specified key once the
* call returns.
*
* @param key key whose mapping is to be removed from the cache
* @return returns false if there was no matching key
* @throws NullPointerException if key is null
* @throws IllegalStateException if the cache is {@link #isClosed()}
* @throws CacheException      if there is a problem doing the put
* @throws ClassCastException   if the implementation supports and is
*                               configured to perform runtime-type-checking,
*                               and the key type is incompatible with that
*                               which has been configured for the
*                               {@link Cache}
*/
boolean remove(K key);

/**
* Atomically removes the mapping for a key only if currently mapped to the
* given value.
* <p/>
* This is equivalent to:
* <code>
*   if (cache.containsKey(key) && cache.get(key).equals(oldValue)) {
*       cache.remove(key);
*       return true;
*   } else {
*       return false;
*   }
* </code>
* except that the action is performed atomically.
*
* @param key      key whose mapping is to be removed from the cache
* @param oldValue value expected to be associated with the specified key
* @return returns false if there was no matching key
* @throws NullPointerException if key is null
* @throws IllegalStateException if the cache is {@link #isClosed()}
* @throws CacheException      if there is a problem doing the put

```

```

* @throws ClassCastException    if the implementation supports and is
*                               configured to perform runtime-type-checking,
*                               and the key or value types are incompatible
*                               with those that have been configured for the
*                               {@link Cache}
*/
boolean remove(K key, V oldValue);

/**
 * Atomically removes the entry for a key only if currently mapped to some value.
 * <p/>
 * This is equivalent to:
 * <code>
 *     if (cache.containsKey(key)) {
 *         V oldValue = cache.get(key);
 *         cache.remove(key);
 *         return oldValue;
 *     } else {
 *         return null;
 *     }
 * </code>
 * except that the action is performed atomically.
 *
 * @param key key with which the specified value is associated
 * @return the value if one existed or null if no mapping existed for this key
 * @throws NullPointerException if the specified key or value is null.
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      if there is a problem during the remove
 * @throws ClassCastException   if the implementation supports and is
 *                               configured to perform runtime-type-checking,
 *                               and the key type is incompatible with that
 *                               which has been configured for the {@link Cache}
 */
V getAndRemove(K key);

/**
 * Atomically replaces the entry for a key only if currently mapped to a
 * given value.
 * <p/>
 * This is equivalent to:
 * <code>
 *     if (cache.containsKey(key) && cache.get(key).equals(oldValue)) {
 *         cache.put(key, newValue);
 *         return true;
 *     } else {
 *         return false;
 *     }
 * </code>
 * except that the action is performed atomically.
 *
 * @param key      key with which the specified value is associated
 * @param oldValue value expected to be associated with the specified key
 * @param newValue value to be associated with the specified key
 * @return <tt>true</tt> if the value was replaced
 * @throws NullPointerException if key is null or if the values are null
 * @throws IllegalStateException if the cache is {@link #isClosed()}

```

```

* @throws CacheException          if there is a problem during the replace
* @throws ClassCastException      if the implementation supports and is
*                                configured to perform runtime-type-checking,
*                                and the key or value types are incompatible
*                                with those which have been configured for the
*                                {@link Cache}
*/
boolean replace(K key, V oldValue, V newValue);

/**
 * Atomically replaces the entry for a key only if currently mapped to some value.
 * <p/>
 * This is equivalent to
 * <code>
 *   if (cache.containsKey(key)) {
 *       cache.put(key, value);
 *       return true;
 *   } else {
 *       return false;
 *   }</code>
 * except that the action is performed atomically.
 *
 * @param key    key with which the specified value is associated
 * @param value  value to be associated with the specified key
 * @return <tt>true</tt> if the value was replaced
 * @throws NullPointerException if key is null or if value is null
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      if there is a problem during the replace
 * @throws ClassCastException  if the implementation supports and is
 *                                configured to perform runtime-type-checking,
 *                                and the key or value types are incompatible
 *                                with those that have been configured for the
 *                                {@link Cache}
 * @see #getAndReplace(Object, Object)
 */
boolean replace(K key, V value);

/**
 * Atomically replaces the value for a given key if and only if there is a
 * value currently mapped by the key.
 * <p/>
 * This is equivalent to
 * <code>
 *   if (cache.containsKey(key)) {
 *       V oldValue = cache.get(key);
 *       cache.put(key, value);
 *       return oldValue;
 *   } else {
 *       return null;
 *   }
 * </code>
 * except that the action is performed atomically.
 *
 * @param key    key with which the specified value is associated
 * @param value  value to be associated with the specified key
 * @return the previous value associated with the specified key, or

```

```

*         <tt>null</tt> if there was no mapping for the key.
* @throws NullPointerException if key is null or if value is null
* @throws IllegalStateException if the cache is {@link #isClosed()}
* @throws CacheException      if there is a problem during the replace
* @throws ClassCastException   if the implementation supports and is
*                               configured to perform runtime-type-checking,
*                               and the key or value types are incompatible
*                               with those that have been configured for the
*                               {@link Cache}
* @see java.util.concurrent.ConcurrentMap#replace(Object, Object)
*/
V getAndReplace(K key, V value);

/**
 * Removes entries for the specified keys.
 * <p/>
 * The order in which the individual removes will occur is undefined.
 *
 * @param keys the keys to remove
 * @throws NullPointerException if keys is null or if it contains a null key
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      if there is a problem during the remove
 * @throws ClassCastException   if the implementation supports and is
 *                               configured to perform runtime-type-checking,
 *                               and any key type is incompatible with that
 *                               which has been configured for the {@link Cache}
 */
void removeAll(Set<? extends K> keys);

/**
 * Removes all of the mappings from this cache.
 * <p/>
 * The order in which the individual removes will occur is undefined.
 * <p/>
 * This is potentially an expensive operation as listeners are invoked.
 * Use {@link #clear()} to avoid this.
 *
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      if there is a problem during the remove
 * @see #clear()
 */
void removeAll();

/**
 * Clears the contents of the cache, without notifying listeners or
 * {@link javax.cache.integration.CacheWriter}s.
 *
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      if there is a problem during the remove
 */
void clear();

/**
 * Obtains an immutable representation of the {@link Configuration} that
 * was used to configure the {@link Cache}.
 *

```

```

* @return the {@link javax.cache.configuration.Configuration}
*/
Configuration<K, V> getConfiguration();

/**
 * Invokes an {@link EntryProcessor} against the {@link Entry} specified by
 * the provided key. If an {@link Entry} does not exist for the specified key,
 * an
 * attempt is made to load it (if a loader is configured) or a surrogate {@link
 * Entry}, consisting of the key with a null value is used instead.
 * <p/>
 *
 * @param key          the key to the entry
 * @param entryProcessor the {@link EntryProcessor} to invoke
 * @param arguments    additional arguments to pass to the
 *                    {@link EntryProcessor}
 * @return the result of the processing, if any, defined by the
 *         {@link EntryProcessor} implementation
 * @throws NullPointerException if key or {@link EntryProcessor} are null
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      if an exception occurred while executing
 *                             the {@link EntryProcessor} (the causing
 *                             exception will be wrapped by the
 *                             CacheException)
 * @throws ClassCastException  if the implementation supports and is
 *                             configured to perform runtime-type-checking,
 *                             and the key or value types are incompatible
 *                             with those that have been configured for the
 *                             {@link Cache}
 * @see EntryProcessor
 */
<T> T invoke(K key,
            EntryProcessor<K, V, T> entryProcessor,
            Object... arguments);

/**
 * Invokes an {@link EntryProcessor} against the set of {@link Entry}s
 * specified by the set of keys. If an {@link Entry} does not exist for the
 * specified key, an attempt is made to load it (if a loader is configured) or a
 * surrogate {@link Entry}, consisting of the key with a null value is used
 * instead.
 * <p/>
 * The order in which the entries for the keys are processed is undefined.
 * Implementations may choose to process the entries in any order, including
 * concurrently. Furthermore there is no guarantee implementations will
 * use the same {@link EntryProcessor} instance to process each entry, as
 * the case may be in a non-local cache topology.
 *
 * @param keys          the set of keys for entries to process
 * @param entryProcessor the {@link EntryProcessor} to invoke
 * @param arguments    additional arguments to pass to the
 *                    {@link EntryProcessor}
 * @return the map of results of the processing per key, if any, defined by the
 *         {@link EntryProcessor} implementation. No mappings will be
 *         returned for {@link EntryProcessor}s that return a <code>null</code>
 *         value for a key

```

```

* @throws NullPointerException if keys or {@link EntryProcessor} are null
* @throws IllegalStateException if the cache is {@link #isClosed()}
* @throws CacheException      if an exception occurred while executing
*                               the {@link EntryProcessor} (the causing
*                               exception will be wrapped by the
*                               CacheException)
* @throws ClassCastException   if the implementation supports and is
*                               configured to perform runtime-type-checking,
*                               and the key or value types are incompatible
*                               with those that have been configured for the
*                               {@link Cache}
* @see EntryProcessor
*/
<T> Map<K, T> invokeAll(Set<? extends K> keys,
                      EntryProcessor<K, V, T> entryProcessor,
                      Object... arguments);

/**
 * Return the name of the cache.
 *
 * @return the name of the cache.
 */
String getName();

/**
 * Gets the {@link CacheManager} that owns and manages the {@link Cache}.
 *
 * @return the manager or null if the {@link Cache} is not
 *         managed
 */
CacheManager getCacheManager();

/**
 * Closing a {@link Cache} signals to the {@link CacheManager} that produced or
 * owns the said {@link Cache} that it should no longer be managed. At this
 * point
 * in time the {@link CacheManager}:
 * <ul>
 * <li>must close and release all resources being coordinated on behalf of the
 * Cache by the {@link CacheManager}. This includes calling the close
 * method on configured {@link javax.cache.integration.CacheLoader},
 * {@link javax.cache.integration.CacheWriter}, registered
 * {@link javax.cache.event.CacheEntryListener}s and {@link
 * javax.cache.expiry.ExpiryPolicy} instances that implement the
 * java.io.Closeable interface.
 * <li>prevent events being delivered to configured
 * {@link javax.cache.event.CacheEntryListener}s registered on the {@link Cache}
 * </li>
 * <li>not return the name of the Cache when the CacheManager getCacheNames()
 * method is called</li>
 * </ul>
 * Once closed any attempt to use an operational method on a Cache will throw an
 * {@link IllegalStateException}.
 */
void close();

```

```

/**
 * Determines whether this Cache instance has been closed. A Cache is
 * considered closed if;
 * <ol>
 * <li>the {@link #close()} method has been called</li>
 * <li>the associated {@link #getCacheManager()} has been closed, or</li>
 * <li>the Cache has been removed from the associated
 *     {@link #getCacheManager()}</li>
 * </ol>
 * <p/>
 * This method generally cannot be called to determine whether a Cache instance
 * is valid or invalid. A typical client can determine that a Cache is invalid
 * by catching any exceptions that might be thrown when an operation is attempted.
 *
 * @return true if this Cache instance is closed; false if it is still open
 */
boolean isClosed();

/**
 * Provides a standard way to access the underlying concrete caching
 * implementation to provide access to further, proprietary features.
 * <p/>
 * If the provider's implementation does not support the specified class,
 * the {@link IllegalArgumentException} is thrown.
 *
 * @param clazz the proprietary class or interface of the underlying concrete
 *     cache. It is this type which is returned.
 * @return an instance of the underlying concrete cache
 * @throws IllegalArgumentException if the caching provider doesn't support
 *     the specified class.
 */
<T> T unwrap(java.lang.Class<T> clazz);

/**
 * Registers a {@link javax.cache.event.CacheEntryListener}. The supplied
 * {@link CacheEntryListenerConfiguration} is used to instantiate a listener
 * and apply it to those events specified in the configuration.
 * @param cacheEntryListenerConfiguration a factory and related configuration
 *     for creating the listener
 * @see javax.cache.event.CacheEntryListener
 * @throws IllegalArgumentException is the same CacheEntryListenerConfiguration
 *     is used more than once
 */
void registerCacheEntryListener(
    CacheEntryListenerConfiguration<K, V> cacheEntryListenerConfiguration);

/**
 * Deregisters a listener, using as its unique identifier the
 * {@link CacheEntryListenerConfiguration} which was used to register it.
 * <p/>
 * Both listeners registered at configuration time,
 * and those created at runtime with {@link #registerCacheEntryListener} can
 * be deregistered.
 *
 * @param cacheEntryListenerConfiguration the factory and related configuration

```

```

*                                     that was used to create the
*                                     listener
*/
void deregisterCacheEntryListener(CacheEntryListenerConfiguration<K, V>
                                   cacheEntryListenerConfiguration);

/**
 * {@inheritDoc}
 * <p/>
 * The ordering of iteration over entries is undefined.
 * <p/>
 * During iteration, any entries that are a). read will have their appropriate
 * CacheEntryReadListeners notified and b). removed will have their appropriate
 * CacheEntryRemoveListeners notified.
 * <p/>
 * {@link java.util.Iterator#next()} may return null if the entry is no
 * longer present, has expired or has been evicted.
 */
Iterator<Cache.Entry<K, V>> iterator();

/**
 * A cache entry (key-value pair).
 */
interface Entry<K, V> {

    /**
     * Returns the key corresponding to this entry.
     *
     * @return the key corresponding to this entry
     */
    K getKey();

    /**
     * Returns the value stored in the cache when this entry was created.
     *
     * @return the value corresponding to this entry
     */
    V getValue();

    /**
     * Provides a standard way to access the underlying concrete cache entry
     * implementation in order to provide access to further, proprietary features.
     * <p/>
     * If the provider's implementation does not support the specified class,
     * the {@link IllegalArgumentException} is thrown.
     *
     * @param clazz the proprietary class or interface of the underlying
     *              concrete cache. It is this type which is returned.
     * @return an instance of the underlying concrete cache
     * @throws IllegalArgumentException if the caching provider doesn't support
     *         the specified class.
     */
    <T> T unwrap(Class<T> clazz);
}

/**

```



```

* A mutable representation of a {@link Cache} {@link Entry}.
*
* @param <K> the type of key
* @param <V> the type of value
*/
public interface MutableEntry<K, V> extends Entry<K, V> {

    /**
     * Checks for the existence of the entry in the cache
     *
     * @return true if the entry exists
     */
    boolean exists();

    /**
     * Removes the entry from the Cache
     */
    void remove();

    /**
     * Sets or replaces the value associated with the key
     * If {@link #exists} is false and setValue is called
     * then a mapping is added to the cache visible once the EntryProcessor
     * completes. Moreover a second invocation of {@link #exists()}
     * will return true.
     * <p/>
     *
     * @param value the value to update the entry with
     * @throws ClassCastException if the implementation supports and is
     *                             configured to perform runtime-type-checking,
     *                             and value type is incompatible with that
     *                             which has been configured for the
     *                             {@link Cache}
     */
    void setValue(V value);
}

/**
 * An invocable function that allows applications to perform compound
 * operations on a {@link Cache.Entry} atomically, according the defined
 * consistency of a {@link Cache}.
 * <p/>
 * Any {@link Cache.Entry} mutations will not take effect until after the
 * {@link EntryProcessor#process(javax.cache.Cache.MutableEntry, Object...)}
 * method has completed execution.
 * <p/>
 * If an exception is thrown by an {@link EntryProcessor}, the exception will
 * be returned wrapped in an {@link CacheException}. No changes will be made
 * to the {@link Cache.Entry}.
 * <p/>
 * Implementations may execute {@link EntryProcessor}s in situ, thus avoiding
 * locking, round-trips and expensive network transfers.
 * <p/>
 * {@link Cache.Entry} access, via a call to
 * {@link javax.cache.Cache.MutableEntry#getValue()}, will behave as if
 * {@link Cache#get(Object)} was called for the key. This includes updating

```

```

* necessary statistics, consulting the configured
* {@link javax.cache.expiry.ExpiryPolicy} and loading from a configured
* {@link javax.cache.integration.CacheLoader}.
* <p/>
* {@link Cache.Entry} mutation, via a call to
* {@link javax.cache.Cache.MutableEntry#setValue(Object)}, will behave
* as if {@link Cache#put(Object, Object)} was called for the key. This
* includes updating necessary statistics, consulting the configured
* {@link javax.cache.expiry.ExpiryPolicy}, notifying
* {@link javax.cache.event.CacheEntryListener}s and a writing to a configured
* {@link javax.cache.integration.CacheWriter}.
* <p/>
* {@link Cache.Entry} removal, via a call to
* {@link javax.cache.Cache.MutableEntry#remove()}, will behave
* as if {@link Cache#remove(Object)} was called for the key. This
* includes updating necessary statistics, notifying
* {@link javax.cache.event.CacheEntryListener}s and causing a delete on a
* configured {@link javax.cache.integration.CacheWriter}.
* <p/>
* As implementations may choose to execute {@link EntryProcessor}s remotely,
* {@link EntryProcessor}s, together with specified parameters and return
* values, may be required to implement {@link java.io.Serializable}.
*
* @param <K> the type of keys maintained by this cache
* @param <V> the type of cached values
*/
public interface EntryProcessor<K, V, T> {

    /**
     * Process an entry.
     *
     * @param entry the entry
     * @param arguments a number of arguments to the process.
     * @return the result of the processing, if any, which is user defined.
     */
    T process(Cache.MutableEntry<K, V> entry, Object... arguments);
}
}

```

Cache Type-Safety

The Java Caching API makes extensive use of Java Generics, as defined by JSR-14, to enable the development of compile-time type-safe applications when adopting Caching.

While available, compile-time type-safety does not guarantee runtime-type correctness for applications adopting caching. For some cache topologies, specifically those that store or communicate entries across Java process boundaries, Java runtime-type-information erasure and the inability to acquire and transfer generic type information may mean application types are unavailable to ensure type-safety of Cache operations in such environments. Care should always be taken to ensure Caches are configured using the appropriate key and value types so that implementations may perform type checking as necessary or required.

Compile-time Type-Safety

Compile-time type-safety is provided by declaring a `Cache` with the required generic types.

Example 1

In the following example a `Cache` is declared to have a key of type `String` and a value of type `Integer`. Compile time errors will be generated when incompatible values are specified when interacting with this cache.

```
Configuration config = new MutableConfiguration();

//create the cache
cacheManager.createCache(cacheName, config);

//get the cache
Cache<String, Integer> cache = cacheManager.getCache(cacheName);

//use the cache
String key = "key";
Integer value1 = 1;
cache.put("key", value1);
Integer value2 = cache.get(key);

//the following will not compile - incorrect types specified
//cache.put(2, "some value");
```

While it is possible to circumvent compile-time type-safety checking by declaring a `Cache` using raw types (not specifying generic type parameters), it is not a recommended practise as it permits simple programming errors to occur.

Example 2

In the following example a `Cache` is declared as a raw type. In this situation no compile-time type-checking can be performed (although type warnings may be generated).

```
Configuration config = new MutableConfiguration();
cacheManager.createCache(cacheName, config);
Cache cache = cacheManager.getCache(cacheName);
String key = "key";
Integer value1 = 1;
cache.put("key", value1);

cache.put(value1, "key1"); //not intended but will still compile and execute!
Integer value2 = (Integer) cache.get(key);
assertEquals(value1, value2);
```

Runtime Type-Safety

In addition to compile-time type-safety, developers may enable runtime type-safety through configuring a `Cache` with specific key and value types. For example, the `MutableConfiguration` class provides the following method to define the required key and value types for a `Cache`.

```

/**
 * Sets the expected type of keys and values for a {@link javax.cache.Cache}
 * configured with this {@link Configuration}. Setting both to null
 * means type-safety checks are not required.
 *
 * @param keyType the expected key type
 * @param valueType the expected value type
 * @return the {@link MutableConfiguration} to permit fluent-style method calls
 */
public MutableConfiguration<K, V> setTypes(Class<K> keyType, Class<V> valueType)

```

When a `Configuration` defines key and value types, a `Cache` returned by `CacheManager.getCache` must enforce that the requested key and value types are the same as those configured. To request a `Cache` with specific key and value types, the following `CacheManager` method must be used.

```

<K, V> Cache<K, V> getCache(String cacheName,
                           Class<K> keyType,
                           Class<V> valueType);

```

Implementations are not required to guarantee or perform key and value type checking at runtime for Cache operations. They are however required to ensure that Caches configured with specific types are returned with those specific types when requested using the `getCache` method.

When a `Configuration` does not define key and value types, or they are `null`, an implementation will not perform runtime type-checking when requesting a configured `Cache`. To request a `Cache` that does not define key and value types, the following `CacheManager` method must be used.

```

<K, V> Cache<K, V> getCache(String cacheName);

```

Attempting to use `getCache` without providing type parameters when a `Cache` has been configured with specific types or using `getCache` with specific type parameters when a `Cache` has been configured without specific types cause an `IllegalArgumentException` to be thrown.

Example 2

In this example a cache is configured to have a `String` key type and an `Integer` value type. The implementation then ensures that the declared types match the configured cache or an `IllegalArgumentException` is thrown.

```

CachingProvider cachingProvider = Caching.getCachingProvider();
CacheManager cacheManager = cachingProvider.getCacheManager();

MutableConfiguration<String, Integer> config = new
    MutableConfiguration<String, Integer>();
config.setTypes(String.class, Integer.class);
cacheManager.createCache("simpleCache", config);
Cache<String, Integer> simpleCache = cacheManager.getCache("simpleCache",
    String.class, Integer.class);

```

```
simpleCache.put("key1", 3);  
Integer value2 = simpleCache.get("key1");
```

While the Java Caching API provides mechanisms for both compile and runtime-type safety, type checking is only applicable to reifiable types of keys and values, including all generic collection types. For example a value of type `List<MyClass>` is not reifiable at runtime and thus may only be compared with the type `List.class`.

Expiry Policies

If an entry is expired, it is not returned from a cache. If no expiry policy has been configured for a cache, it defaults to the `Eternal` expiry policy, where cache entries do not expire.

Expiry policies may be set at configuration time by providing an `ExpiryPolicy` implementation, See below for the definition.

```
/**
 * Defines functions to determine when cache entries will expire based on
 * creation, access and modification operations.
 * <p/>
 * Each of the functions return a new {@link Duration} that of which specifies the
 * amount of time that must pass before a cache entry is considered expired.
 *
 * @param <K> the type of keys
 * @param <V> the type of values
 *
 * @author Brian Oliver
 * @author Greg Luck
 */
public interface ExpiryPolicy<K, V> {

    /**
     * Gets the duration before the newly Cache.Entry is considered expired.
     * <p/>
     * This method is called by the caching implementation after a Cache.Entry is
     * created, but before the said entry is added to a cache, to determine the
     * {@link Duration} before the said entry expires. If a {@link Duration#ZERO}
     * is returned the Cache.Entry is considered to be already expired and will
     * not be added to the Cache.
     * <p/>
     * Should an exception occur while determining the Duration, an implementation
     * specific default Duration will be used.
     *
     * @param entry the cache entry that was created
     * @return the new duration until the entry expires
     */
    Duration getExpiryForCreatedEntry(Entry<? extends K, ? extends V> entry);

    /**
     * Gets the duration before the accessed Cache.Entry is considered expired.
     * <p/>
     * This method is called by the caching implementation after a Cache.Entry is
     * accessed to determine the {@link Duration} before the said entry expires in
     * the future. If a {@link Duration#ZERO} is returned the Cache.Entry will be
     * considered expired for future access. Returning <code>null</code> will

```

```

* result in no change to the previously understood expiry {@link Duration}.
* <p/>
* Should an exception occur while determining the Duration, an implementation
* specific default Duration will be used.
*
* @param entry the cache entry that was accessed
* @return the new duration until the entry expires
*/
Duration getExpiryForAccessedEntry(Entry<? extends K, ? extends V> entry);

/**
* Gets the duration before the modified Cache.Entry is considered expired.
* <p/>
* This method is called by the caching implementation after a Cache.Entry is
* modified to determine the {@link Duration} before the updated entry expires.
* If a {@link Duration#ZERO} is returned the Cache.Entry is considered already
* expired. Returning <code>null</code> will result in no change to the
* previously understood expiry {@link Duration}.
* <p/>
* Should an exception occur while determining the Duration, an implementation
* specific default Duration will be used.
*
* @param entry the cache entry that was modified
* @return the new duration until the entry expires
*/
Duration getExpiryForModifiedEntry(Entry<? extends K, ? extends V> entry);
}

```

Entry expiry occurs a set time after certain cache operations are performed, the time defined using the `Duration` class. `Duration` is a pair made up of a `java.util.concurrent.TimeUnit` and a long `durationAmount`. The minimum allowed `TimeUnit` is `TimeUnit.MILLISECONDS`.

The expiry durations depend on the configured expiry policy and the cache operation performed. The following `ExpiryPolicy` methods are defined to determine suitable durations based on cache operations:

- `getExpiryForCreatedEntry()` - the duration for an entry when it is created
- `getExpiryForAccessedEntry()` - the new duration for an entry when it is accessed
- `getExpiryForModifiedEntry()` - the new duration for an entry when it is modified

When these methods are called the `ExpiryPolicy` will return one of the following:

- A duration equal to the cache's configured expiry duration
- `Duration.ZERO` which indicates the entry is now considered expired

In addition `getExpiryForModifiedEntry()` and `getExpiryForAccessedEntry()` may also return `null`, indicating the caching implementation should leave the entry's expiry time unchanged.

Duration#ZERO is very useful in a custom expiry policy for expiring a selected entry under control of business logic. See Example 1. There are also factory methods for creating other useful durations: 1 day, 1 hour, 30 minutes, 20 minutes, 10 minutes, 5 minutes and 1 minute.

Example 1

This example uses `getExpiryForModifiedEntry()` to set the expiry duration to `Duration.ZERO` when the a cart is closed and put back in the cache which causes them to expire.

```
public class ShoppingCartExpiryPolicy implements ExpiryPolicy<String,
    ShoppingCart> {

    @Override
    public Duration getExpiryForCreatedEntry(Entry<? extends String, ? extends
        ShoppingCart> entry) {
        return TWENTY_MINUTES;
    }

    @Override
    public Duration getExpiryForAccessedEntry(Entry<? extends String, ? extends
        ShoppingCart> entry) {
        //we don't change the expiry time for accesses
        return null;
    }

    @Override
    public Duration getExpiryForModifiedEntry(Entry<? extends String, ? extends
        ShoppingCart> entry) {
        ShoppingCart c = entry.getValue();

        //when a shopping cart is closed, we can expire it immediately,
        //otherwise we give shopping cart another 20 minutes. To cause this
        //to take immediate effect the cart must be put back in the cache.
        return c.isClosed() ? ZERO : TWENTY_MINUTES;
    }
}
```

The following table details how each of the cache methods interact with a configured `ExpiryPolicy`.

Method	<code>ExpiryPolicy.getExpiryForCreatedEntry</code> called?	<code>ExpiryPolicy.getExpiryForAccessedEntry</code> called?	<code>ExpiryPolicy.getExpiryForModifiedEntry</code> called?
<code>boolean containsKey(K key)</code>	No	No	No
<code>V get(K key)</code>	No, unless read-through caused a load	Yes	No

<code>Map<K,V> getAll(Collection<? extends K> keys)</code>	No, unless read-through caused a load	Yes	No
<code>V getAndPut(K key, V value)</code>	Yes (when the key is not associated with an existing value)	No	Yes (when the key is associated with an existing value)
<code>V getAndRemove(K key)</code>	No	No	No
<code>V getAndReplace(K key, V value)</code>	No	No	Yes (when the key is associated with an existing value)
<code>CacheManager getCacheManager()</code>	No	No	No
<code>CacheConfiguration getConfiguration()</code>	No	No	No
<code>String getName()</code>	No	No	No
<code>Iterator<Cache.Entry<K, V>> iterator()</code>	No	Yes (when an entry is visited by an iterator)	No
<code>void loadAll(Set<? extends K> keys, boolean replaceExistingValues, CompletionListener listener)</code>	Yes (when a key is not associated with a loaded value)	No	Yes (when a key is associated with a loaded value and the value should be replaced)
<code>void put(K key, V value)</code>	Yes (when the key is not associated with an existing value)	No	Yes (when the key is associated with an existing value)
<code>void putAll(Map<? extends K, ? extends V> map)</code>	Yes (when the key is not associated with an existing value)	No	Yes (when the key is associated with an existing value)

<code>boolean putIfAbsent(K key, V value)</code>	Yes (when the key is not associated with an existing value)	No	No
<code>boolean remove(K key)</code>	No	No	No
<code>boolean remove(K key, V oldValue)</code>	No	No	No
<code>void removeAll()</code>	No	No	No
<code>void removeAll(Set<? extends K> keys)</code>	No	No	No
<code><T> T invoke(K key, EntryProcessor<K, V, T> entryProcessor, Object... arguments)entryProcessor);</code>	Yes if read-through, <code>getValue</code> called and the <code>CacheLoader</code> is called.	Yes, if <code>getValue()</code> called	Yes, if <code>setValue()</code> called.
<code><T> Map<K, T> invokeAll(Set<? extends K> keys, EntryProcessor<K, V, T> entryProcessor, Object... arguments);</code>	Yes if read-through, <code>getValue</code> called and the <code>CacheLoader</code> is called.	Yes, if <code>getValue()</code> called	Yes, if <code>setValue()</code> called.
<code>boolean replace(K key, V value)</code>	No	No	Yes (when the key is associated with an existing value)
<code>boolean replace(K key, V oldValue, V newValue)</code>	No	Yes (when value is not replaced)	Yes (when value is replaced)
<code><T> T unwrap(Class<T> cls)</code>	No	No	No

Five expiry policy implementations are defined and included with the specification:

1. `Created` - expire a set time after creation.
2. `Modified` - expire a set time after creation. Refresh expiry when an entry is updated.

3. `Accessed` - expire a set time after creation. Refresh expiry when an entry is accessed (a read operation of some kind)
4. `Touched` - expire a set time after creation. Refresh expiry when an entry is updated or accessed
5. `Eternal` - never expire. This is the default.

Integration

Convenience methods have been created to ease integration with external resources. These are in the `javax.cache.integration` package.

Two interfaces `CacheLoader` and `CacheWriter` which are defined as follows:

```
/**
 * Used when a cache is read-through or when loading data into a cache via the
 * {@link javax.cache.Cache#loadAll(java.util.Set, boolean,
 * CompletionListener)} method.
 * <p/>
 * See {@link CacheWriter} which is the corollary for write-through caching.
 *
 * @param <K> the type of keys handled by this loader
 * @param <V> the type of values generated by this loader
 * @author Greg Luck
 * @author Yannis Cosmadopoulos
 * @since 1.0
 * @see CacheWriter
 */
public interface CacheLoader<K, V> {

    /**
     * Loads an object. Application developers should implement this
     * method to customize the loading of cache object. This method is called
     * by the caching service when the requested object is not in the cache. If
     * the object can't be loaded null should be returned.
     *
     * @param key the key identifying the object being loaded
     * @return The entry for the object that is to be stored in the cache or
     *         null if the object can't be loaded
     */
    Cache.Entry<K, V> load(K key);

    /**
     * Loads multiple objects. Application developers should implement this
     * method to customize the loading of cache object. This method is called
     * by the caching service when the requested object is not in the cache. If
     * an object can't be loaded, it is not returned in the resulting map.
     * <p/>
     *
     * @param keys keys identifying the values to be loaded
     * @return A map of key, values to be stored in the cache.
     */
    Map<K, V> loadAll(Set<? extends K> keys);
}
```

```

/**
 * A CacheWriter is used for write-through to an external resource.
 * <p/>
 * The semantics of write-through when an <code>Exception</code> is thrown are
 * implementation specific.
 * <p/>
 * Under Default Consistency, the non-batch writer methods are atomic with respect
 * to the corresponding cache operation.
 * <p/>
 * For batch methods under Default Consistency, the entire cache operation
 * is not required to be atomic in {@link Cache} and is therefore not required to
 * be atomic in the writer. As individual writer operations can fail, cache
 * operations are not required to occur until after the writer batch method has
 * returned or, in the case of partial success, thrown an exception. In the case of
 * partial success, the collection of entries must contain only those entries which
 * failed.
 * <p/>
 * The behaviour of the caching implementation in the case of partial success is
 * undefined.
 * <p/>
 * The semantics of Transactional Consistency are implementation specific.
 * <p/>
 * The entry passed into {@link #write(javax.cache.Cache.Entry)} is independent
 * of the cache mapping for that key, meaning that if the value changes in the
 * cache or is removed it does not change the said entry.
 *
 * @param <K> the type of keys maintained by this map
 * @param <V> the type of mapped values
 * @author Greg Luck
 * @since 1.0
 * @see CacheLoader
 */
public interface CacheWriter<K, V> {

    /**
     * Write the specified value under the specified key to the external resource.
     * <p/>
     * This method is intended to support both key/value creation and value update
     * for a specific key.
     *
     * @param entry the entry to be written
     * @throws javax.cache.CacheException if the write fails. If thrown the
     *         cache mutation will not occur.
     */
    void write(Cache.Entry<? extends K, ? extends V> entry);

    /**
     * Write the specified entries to the external resource. This method is intended
     * to support both insert and update.
     * <p/>

```

```

* The order in which individual writes occur is undefined, as
* {@link Cache#putAll(java.util.Map)} also has undefined ordering.
* <p/>
* If this operation fails (by throwing an exception) after a partial success,
* the writer must remove any successfully written entries from the entries
* collection so that the caching implementation knows what succeeded and can
* mutate the cache.
*
* @param entries a mutable collection to write. Upon invocation, it contains
* the entries to write for write-through. Upon return the
* collection must only contain entries that were not
* successfully written. (see partial success above)
* @throws javax.cache.CacheException if one or more of the writes fail. If
* thrown cache mutations will occur for
* entries which succeeded.
*/
void writeAll(Collection<Cache.Entry<? extends K, ? extends V>> entries);

```

```

/**
* Delete the cache entry from the external resource.
* <p/>
* Expiry of a cache entry is not a delete hence will not cause this method to
* be invoked.
*
* @param key the key that is used for the delete operation
* @throws javax.cache.CacheException if delete fails. If thrown the
* cache delete will not occur.
*/
void delete(Object key);

```

```

/**
* Remove data and keys from the external resource for the given collection of
* keys, if present.
* <p/>
* The order in which individual deletes occur is undefined, as
* {@link Cache#removeAll(java.util.Set)} also has undefined ordering.
* <p/>
* If this operation fails (by throwing an exception) after a partial success,
* the writer must remove any successfully written entries from the entries
* collection so that the caching implementation knows what succeeded and can
* mutate the cache.
* <p/>
* Expiry of a cache entry is not a delete hence will not cause this method to
* be invoked.
*
* @param keys a mutable collection of keys for entries to delete. Upon
* invocation, it contains the keys to delete for write-through.
* Upon return the collection must only contain the keys that were
* not successfully deleted. (see partial success above)

```

```

    * @throws javax.cache.CacheException if one or more deletes fail. If thrown
    *                                     cache deletes will occur for entries which succeeded.
    */
    void deleteAll(Collection<?> keys);
}

```

These interfaces are used as described below.

Cache Loading

The Cache loadAll method is used to load values from an external resource and is defined as follows:

```

/**
 * This method provides a means to "pre-load" objects into the cache. It will,
 * asynchronously, load the specified objects into the cache using the associated
 * cache loader for the given keys.
 * <p/>
 * If an entry for a key already exists in the Cache, a value will be loaded
 * if and only if replaceExistingValues is true. If no loader is
 * configured for the cache, no objects will be loaded. If a problem is
 * encountered during the retrieving or loading of the objects,
 * an exception is provided to the {@link CompletionListener}. Once the
 * operation has completed, the specified CompletionListener is notified.
 * <p/>
 * Implementations may choose to load multiple keys from the provided
 * iterable in parallel. Iteration must not occur in parallel, thus
 * allow for non-thread-safe Iterables, but loading may.
 * <p/>
 * The thread on which the completion listener is called is implementation
 * dependent. An implementation may also choose to serialize calls to
 * different CompletionListeners rather than use a thread per
 * CompletionListener.
 *
 * @param keys                the keys to load
 * @param replaceExistingValues when true existing values in the Cache will
 *                               be replaced by those loaded from a CacheLoader
 * @param completionListener  the CompletionListener (may be null)
 * @throws NullPointerException if keys is null or if keys contains a null.
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      thrown if there is a problem performing the
 *                               load. This may also be thrown on calling if
 *                               there are insufficient threads available to
 *                               perform the load.
 * @throws ClassCastException  if the implementation supports and is
 *                               configured to perform runtime-type-checking,
 *                               and any key type is incompatible with that
 *                               which has been configured for the {@link Cache}
 */
void loadAll(Set<? extends K> keys, boolean replaceExistingValues,
             CompletionListener completionListener);

```

For this method to be used a `CacheLoader` must have been set in `Configuration` when the cache was created. A cache is not required to be set to read-through caching mode to use this method.

Loading may take a significant amount of time. For this reason a `CompletionListener` can be passed in which is notified on completion or on exception. It is defined as follows:

```
/**
 * A CompletionListener is implemented by an application when it needs to be
 * notified of the completion of some Cache operation.
 * <p/>
 * When the operation is complete, the Cache provider notifies the application
 * by calling the {@link #onCompletion()} method of the {@link
 * CompletionListener}.
 * <p/>
 * If the operation fails for any reason, the Cache provider calls the
 * {@link #onException(Exception)} method of the {@link CompletionListener}.
 * <p/>
 * To support a Java Future-based approach to synchronously wait for a Cache
 * operation to complete, use a {@link CompletionListenerFuture}.
 * <p/>
 * A CompletionListener will use an implementation specific thread for the call.
 *
 * @author Brian Oliver
 * @see CompletionListenerFuture
 */
public interface CompletionListener {

    /**
     * Notifies the application that the operation completed successfully.
     */
    void onCompletion();

    /**
     * Notifies the application that the operation failed.
     *
     * @param e the Exception that occurred
     */
    void onException(Exception e);
}
```

There is also a blocking implementation of `CompletionListener`, `CompletionListenerFuture`. It implements both `CompletionListener` and `Future`. If the `onException(Exception e)` method of `CompletionListener` is called, the exception is wrapped in `ExecutionException` and rethrown by the `Future` `get()` and `get(long timeout, TimeUnit unit)` methods.

Example 1

Using CompletionListenerFuture.

```
HashSet<String> keys = new HashSet<>();
keys.add("234321kj");
keys.add("4fsldkj");

//create a completion future to use to wait for loadAll
CompletionListenerFuture future = new CompletionListenerFuture();

//load the values for the set of keys, replacing those that may already
//exist in the cache
cache.loadAll(keys, true, future);

//wait for the cache to load the keys
try {
    future.get();
} catch (InterruptedException e) {
    //future interrupted
    e.printStackTrace();
} catch (ExecutionException e) {
    //throwable was what was sent to onException(Exception e)
    Throwable throwable = e.getCause();
}
```

The `loadAll` method is useful for pre-loading a cache with data from an external resource. It may be required because the application logic assumes the data is there. Another usage is cache warming. Here it will not cause an application error if the data is absent from the cache, but it will affect performance or scalability.

Read-Through Caching

A read-through cache behaves exactly the same way as a non-read-through cache except that certain accessor methods will invoke the `CacheLoader` if the entry or entries are missing from the `Cache`.

Read-Through caching is set at configuration time by calling `setReadThrough(boolean isReadThrough)` on `MutableConfiguration`. A `CacheLoader Factory` must also have been defined. The `CacheLoader` is used to load entries from an external resource.

The effect on each method invocation when a cache is in read-through mode is described in the following table:

Method	Invoke Read-Through
<code>boolean containsKey(K key)</code>	No

<code>V get(K key)</code>	Yes
<code>Map<K,V> getAll(Collection<? extends K> keys)</code>	Yes. Invokes <code>loadAll()</code>
<code>V getAndPut(K key, V value)</code>	No
<code>V getAndRemove(K key)</code>	No
<code>V getAndReplace(K key, V value)</code>	No
<code><T> T invoke(K key, EntryProcessor<K, V, T> entryProcessor, Object... arguments)entryProcessor);</code>	Yes, if <code>getValue()</code> called.
<code><T> Map<K, T> invokeAll(Set<? extends K> keys, EntryProcessor<K, V, T> entryProcessor, Object... arguments);</code>	Yes, if <code>getValue</code> called.
<code>Iterator<Cache.Entry<K, V>> iterator()</code>	No
<code>void loadAll(Set<? extends K> keys, boolean replaceExistingValues, CompletionListener completionListener)</code>	Yes. Uses the <code>CacheLoader.loadAll()</code> method. Even when the cache is not read-through
<code>void put(K key, V value)</code>	No
<code>void putAll(Map<? extends K,? extends V> map)</code>	No
<code>boolean putIfAbsent(K key, V value)</code>	No
<code>boolean remove(K key)</code>	No
<code>boolean remove(K key, V oldValue)</code>	No
<code>void removeAll()</code>	No
<code>void removeAll(Set<? extends K> keys)</code>	No
<code>boolean replace(K key, V value)</code>	No
<code>boolean replace(K key, V oldValue, V newValue)</code>	No

Read-through caching is a useful idiom for lazily loading a cache. It is also useful in shielding the Cache user from the details of how an external resource is loaded into the cache.

When it is important for some or all cached content to be pre-loaded, use the `loadAll` method.

Write-Through Caching

A write-through cache behaves exactly the same way as a non-write-through cache except that certain mutative methods will invoke the CacheWriter.

Write-Through caching is set at configuration time by calling `setWriteThrough(boolean isWriteThrough)` on `MutableConfiguration`. A `CacheWriter` Factory must also have been defined. The `CacheWriter` is used to write and remove entries from an external resource.

The effect on each method invocation when a cache is in write-through mode is described in the following table:

Method	Invoke Write-Through
<code>boolean containsKey(K key)</code>	No
<code>V get(K key)</code>	No
<code>Map<K,V> getAll(Collection<? extends K> keys)</code>	No
<code>V getAndPut(K key, V value)</code>	Yes
<code>V getAndRemove(K key)</code>	Yes
<code>V getAndReplace(K key, V value)</code>	Yes
<code><T> T invoke(K key, EntryProcessor<K, V, T> entryProcessor, Object... arguments)</code>	Yes, if <code>setValue()</code> is called.
<code><T> Map<K, T> invokeAll(Set<? extends K> keys, EntryProcessor<K, V, T> entryProcessor, Object... arguments);</code>	Yes, if <code>setValue()</code> is called.
<code>Iterator<Cache.Entry<K, V>> iterator()</code>	No
<code>void loadAll(Set<? extends K> keys, boolean replaceExistingValues, CompletionListener completionListener)</code>	No
<code>void put(K key, V value)</code>	Yes
<code>void putAll(Map<? extends K,? extends V> map)</code>	Yes, <code>writeAll</code> will be called
<code>boolean putIfAbsent(K key, V value)</code>	Yes, if this method returns true
<code>boolean remove(K key)</code>	Yes
<code>boolean remove(K key, V oldValue)</code>	Yes, if this method returns true

<code>void removeAll()</code>	Yes
<code>void removeAll(Set<? extends K> keys)</code>	Yes
<code>boolean replace(K key, V value)</code>	Yes
<code>boolean replace(K key, V oldValue, V newValue)</code>	Yes, if this method returns true

Write-through caching is a useful idiom for keeping an external resource updated with cache changes. It shields the Cache user from the details of how an external resource is written to.

Cache Entry Listeners

The `javax.cache.event` package contains classes and interfaces for event processing.

Events and Event Types

A `CacheEntryEvent` is defined as follows:

```
/**
 * A Cache entry event base class.
 *
 * @param <K> the type of key
 * @param <V> the type of value
 * @author Greg Luck
 * @since 1.0
 */
public abstract class CacheEntryEvent<K, V> extends EventObject
    implements Cache.Entry<K, V> {

    private EventType eventType;

    /**
     * Constructs a cache entry event from a given cache as source
     *
     * @param source the cache that originated the event
     */
    public CacheEntryEvent(Cache source, EventType eventType) {
        super(source);
        this.eventType = eventType;
    }

    /**
     * {@inheritDoc}
     */
    @Override
    public final Cache getSource() {
        return (Cache) super.getSource();
    }

    /**
     * Returns the previous value, that of which existed prior to the
     * modification of the Entry value.
     *
     * @return the previous value or null if there was no
     *         previous value
     */
    public abstract V getOldValue();

}
```

```

    * Whether the old value is available.
    *
    * @return true if the old value is populated
    */
public abstract boolean isOldValueAvailable();

/**
 * Gets the event type of this event
 *
 * @return the event type.
 */
public final EventType getEventType() {
    return eventType;
}
}

```

There are four types of event, as enumerated by the EventType enum, defined as follows:

```

/**
 * The type of event received by the listener.
 *
 * @author Greg Luck
 */
public enum EventType {

    /**
     * An event type indicating that the cache entry was created.
     */
    CREATED,

    /**
     * An event type indicating that the cache entry was updated. i.e. a previous
     * mapping existed
     */
    UPDATED,

    /**
     * An event type indicating that the cache entry was removed.
     */
    REMOVED,

    /**
     * An event type indicating that the cache entry has expired.
     */
    EXPIRED
}

```

CacheEntryListeners

These events are propagated to `CacheEntryListeners` which may be registered with a `Cache`. The tagging interface for these, which defines the behaviour which must be honoured by all sub-interfaces is defined below:

```
/**
 * A tagging interface for cache entry listeners.
 * <p/>
 * Sub-interfaces exist for the various cache events allowing a listener to be
 * created which implements only those listeners it is interested in.
 * <p/>
 * Listeners should be implemented with care. In particular it is important to
 * consider their impact on performance and latency.
 * <p/>
 * Listeners:
 * <ul>
 * <li>are fired after the entry is mutated in the cache</li>
 * <li>if synchronous are fired, for a given key, in the order in which events
 * occur</li>
 * <li>block the calling thread until the listener returns,
 * where the listener was registered as synchronous</li>
 * <li>which are asynchronous iterate through multiple events with an undefined
 * ordering, except that events on the same key are in the the order in which the
 * events occur.</li>
 * </ul>
 * Listeners follow the observer pattern. An exception thrown by a
 * listener does not cause the cache operation to fail.
 * <p/>
 * <p/>
 * A listener which mutates a cache on the CacheManager may cause a deadlock.
 * Detection and response to deadlocks is implementation specific.
 * <p/>
 * A listener on a transactional cache is executed orthogonally to the transaction.
 * If synchronous it is executed after the mutation and not after the transaction
 * commits, and if asynchronous the timing is undefined. A listener which throws
 * an exception will not affect the transaction. A transaction which is rolled back
 * will not unfire a listener.
 * <p/>
 * A listener will be notified of events for the time it is registered. Listeners
 * are not required to be durable.
 *
 * @param <K> the type of key
 * @param <V> the type of value
 * @author Yannis Cosmadopoulos
 * @author Greg Luck
 * @see CacheEntryCreatedListener
 * @see CacheEntryUpdatedListener
```

```

* @see CacheEntryRemovedListener
* @see CacheEntryExpiredListener
* @see EventType
* @since 1.0
*/
public interface CacheEntryListener<K, V> extends EventListener {

}

```

There are four sub-interfaces, corresponding to each of the `EventTypes`, defined as follows:

```

/**
 * Invoked after a cache entry is created, or if a batch call is made, after the
 * entries are created.
 * <p/>
 * If an entry for the key existed prior to the operation it is not invoked,
 * instead {@link CacheEntryUpdatedListener} is invoked.
 *
 * @param <K> the type of key
 * @param <V> the type of value
 * @author Yannis Cosmadopoulos
 * @author Greg Luck
 * @see CacheEntryUpdatedListener
 * @since 1.0
 */
public interface CacheEntryCreatedListener<K, V> extends CacheEntryListener<K, V> {

    /**
     * Called after one or more entries have been created.
     *
     * @param events The entries just created.
     * @throws CacheEntryListenerException if there is problem executing the listener
     */
    void onCreated(Iterable<CacheEntryEvent<? extends K, ? extends V>> events)
        throws CacheEntryListenerException;

}

/**
 * Invoked if an existing cache entry is updated, or if a batch call is made,
 * after the entries are updated.
 * <p/>
 *
 * @param <K> the type of key
 * @param <V> the type of value
 * @author Yannis Cosmadopoulos
 * @author Greg Luck
 * @see CacheEntryCreatedListener
 * @since 1.0
 */

```



```

public interface CacheEntryUpdatedListener<K, V> extends CacheEntryListener<K, V> {

    /**
     * Called after one or more entries have been updated.
     *
     * @param events The entries just updated.
     * @throws CacheEntryListenerException if there is problem executing the listener
     */
    void onUpdated(Iterable<CacheEntryEvent<? extends K, ? extends V>> events)
        throws CacheEntryListenerException;
}

/**
 * Invoked if a cache entry is removed, or if a batch call is made, after the
 * entries are removed.
 *
 * @param <K> the type of key
 * @param <V> the type of value
 * @author Yannis Cosmadopoulos
 * @author Greg Luck
 * @since 1.0
 */
public interface CacheEntryRemovedListener<K, V> extends CacheEntryListener<K, V> {

    /**
     * Called after one or more entries have been removed. If no entry existed for
     * a key an event is not raised for it.
     *
     * @param events The entries just removed.
     * @throws CacheEntryListenerException if there is problem executing the listener
     */
    void onRemoved(Iterable<CacheEntryEvent<? extends K, ? extends V>> events)
        throws CacheEntryListenerException;
}

/**
 * Invoked if a cache entry or entries are evicted due to expiration.
 *
 * @param <K> the type of key
 * @param <V> the type of value
 * @author Greg Luck
 * @since 1.0
 */
public interface CacheEntryExpiredListener<K, V> extends CacheEntryListener<K, V> {

    /**
     * Called after one or more entries have been expired by the cache. This is not
     * necessarily when an entry is expired, but when the cache detects the expiry.

```

```

*
* @param events The entries just removed.
* @throws CacheEntryListenerException if there is problem executing the listener
*/
void onExpired(Iterable<CacheEntryEvent<? extends K, ? extends V>> events)
    throws CacheEntryListenerException;
}

```

The motivation for this design is to allow efficient implementation of distributed listeners.

Registration of Listeners

Listeners are not assumed to be in-process with a cache. To avoid registration of instances which may not support `Serialization`, instead `CacheEntryListenerConfigurations` are used. These may be added to `MutableConfiguration` using `MutableConfiguration.addCacheEntryListenerConfiguration` at configuration time or to a `Cache` using `Cache.registerCacheEntryListener` at runtime.

Listeners may be deregistered at runtime using `Cache.deregisterCacheEntryListener`.

These are defined as shown below:

```

/**
 * Defines the configuration requirements for a
 * {@link javax.cache.event.CacheEntryListener} and a {@link Factory} for its
 * creation.
 *
 * @param <K> the type of keys
 * @param <V> the type of values
 * @author Brian Oliver
 * @author Greg Luck
 */
public interface CacheEntryListenerConfiguration<K, V> {
    /**
     * Obtains the {@link Factory} for the
     * {@link javax.cache.event.CacheEntryListener}.
     *
     * @return the {@link Factory} for the
     *         {@link javax.cache.event.CacheEntryListener}
     */
    Factory<CacheEntryListener<? super K, ? super V>> getCacheEntryListenerFactory();

    /**
     * Determines if the old value should be provided to the
     * {@link CacheEntryListener}.
     *
     * @return true if the old value is required by the
     *         {@link CacheEntryListener}
     */
}

```

```

    */
    boolean isOldValueRequired();

    /**
     * Obtains the {@link Factory} for the
     * {@link javax.cache.event.CacheEntryEventFilter} that
     * should be applied prior to notifying the {@link CacheEntryListener}.
     * When null no filtering is applied and all appropriate events
     * are notified.
     *
     * @return the {@link Factory} for the
     *         {@link javax.cache.event.CacheEntryEventFilter} or null
     *         if no filtering is required
     */
    Factory<CacheEntryEventFilter<? super K, ? super V>>
    getCacheEntryEventFilterFactory();

    /**
     * Determines if the thread that caused an event to be created should be
     * blocked (not return from the operation causing the event) until the
     * {@link CacheEntryListener} has been notified.
     *
     * @return true if the thread that created the event should block
     */
    boolean isSynchronous();
}

```

For convenience, a `MutableCacheEntryListenerConfiguration` is provided in the package, which may be used.

Multiple `CacheEntryListenerConfigurations` can be added to a `Configuration`. When the cache is initiated the `CacheEntryListeners` are created using the `Factory` are registered. A cache may have any number of listeners for the same or different `EventTypes`. There is no ordering guarantee between listeners for their creation or dispatch of events.

Invocation of Listeners

Cache Listeners:

- are fired after the entry is mutated in the cache
- if synchronous, are fired, for a given key, in the order in which events occur, blocking the calling thread until the listener returns
- if asynchronous, iterate through multiple events with an undefined ordering, except that events on the same key must be processed in the order in which the events occur.

Listeners follow the observer pattern. An exception thrown by a listener does not cause the cache operation to fail.

A listener which mutates a cache on the CacheManager may cause a deadlock. Detection and response to deadlocks is implementation specific.

A listener on a transactional cache is executed orthogonally to the transaction. If synchronous it is executed after the mutation and not after the transaction commits, and if asynchronous the timing is undefined. A listener which throws an exception will not affect the transaction. A transaction which is rolled back will not unfire a listener.

A registered listener will be invoked at most once by a caching implementation for each event.

i.e. listeners have once across a cluster semantics, not broadcast and execute in each node semantics.

A listener is not required to be in-process with the originating event.

In a distributed implementation, the listener may be implemented anywhere.

A listener may have a CacheEntryEventFilter, as part of its CacheEntryListenerConfiguration. These are defined as shown below:

```
/**
 * A function which may be used to check {@link CacheEntryEvent}s prior to being
 * dispatched to {@link CacheEntryListener}s.
 * <p/>
 * A filter must not create side effects.
 *
 * @param <K> the type of key
 * @param <V> the type of value
 * @author Greg Luck
 * @author Brian Oliver
 * @since 1.0
 */
public interface CacheEntryEventFilter<K, V> {

    /**
     * Evaluates specified {@link CacheEntryEvent}.
     *
     * @param event the event that occurred
     * @return true if the evaluation passes, otherwise false.
     *         The effect of returning true is that listener will be invoked
     * @throws CacheEntryListenerException if there is problem executing the listener
     */
    boolean evaluate(CacheEntryEvent<? extends K, ? extends V> event)
        throws CacheEntryListenerException;
}
```

A CacheEntryEventFilter is not required to be in-process with the originating event.

In a distributed implementation, the filter may be implemented wherever it gives the best performance advantage.

The table below summarises which listeners are invoked by each cache operation. Conditions are on the state of the entry before the operation. Expiry is always “No”. The exact timing of expiry is caching implementation specific.

Method	Created	Expired	Removed	Update
<code>boolean containsKey(K key)</code>	No	No	No	No
<code>V get(K key)</code>	Yes, if created by read-through	No	No	No
<code>Map<K,V> getAll(Collection<? extends K> keys)</code>	Yes, if created by read-through	No	No	No
<code>V getAndPut(K key, V value)</code>	if missing	No	No	if there
<code>V getAndRemove(K key)</code>	No	No	if there	No
<code>V getAndReplace(K key, V value)</code>	No	No	No	if there
<code><T> T invoke(K key, EntryProcessor<K, V> entryProcessor);</code>	Yes, if <code>setValue()</code> created an entry, or <code>getValue()</code> created an entry by read-through	No	Yes, if <code>remove()</code> was called	Yes, if <code>setValue()</code> updated an entry
<code><T> Map<K, T> invokeAll(Set<? extends K> keys, EntryProcessor<K, V, T> entryProcessor, Object... arguments);</code>	Yes, if <code>setValue()</code> created an entry, or <code>getValue()</code> created an entry by read-through	No	Yes, if <code>remove()</code> was called	Yes, if <code>setValue()</code> updated an entry
<code>Iterator<Cache.Entry<K, V>> iterator()</code>	No	No	Yes, if <code>remove()</code>	No

			was called	
<code>void loadAll(Set<? extends K> keys, boolean replaceExistingValues, CompletionListener completionListener);</code>	if missing	No	No	if there
<code>void put(K key, V value)</code>	if missing	No	No	if there
<code>void putAll(Map<? extends K, ? extends V> map)</code>	if missing	No	No	if there
<code>boolean putIfAbsent(K key, V value)</code>	if missing	No	No	No
<code>boolean remove(K key)</code>	No	No	if there	No
<code>boolean remove(K key, V oldValue)</code>	No	No	if there && equal	No
<code>void removeAll()</code>	No	No	if there	No
<code>void removeAll(Set<? extends K> keys)</code>	No	No	if there	No
<code>boolean replace(K key, V value)</code>	No	No	No	if there
<code>boolean replace(K key, V oldValue, V newValue)</code>	No	No	No	if there && equal

Entry Processors

A `javax.cache.Cache.EntryProcessor` is an invocable function, much like a `java.util.concurrent.Callable`, that applications may use to efficiently perform compound Cache operations, including access, update and removal atomically on a Cache Entry, without requiring explicit locking or transactions.

For example, an application that may wish to inspect the value of a Cache entry, calculate a new value, update the entry and return some other value atomically, could do so using a custom `EntryProcessor` implementation.

The `javax.cache.Cache.EntryProcessor` interface is defined as follows:

```
/**
 * An invocable function that allows applications to perform compound
 * operations on a {@link Cache.Entry} atomically, according the defined
 * consistency of a {@link Cache}.
 * <p/>
 * Any {@link Cache.Entry} mutations will not take effect until after the
 * {@link EntryProcessor#process(javax.cache.Cache.MutableEntry, Object...)}
 * method has completed execution.
 * <p/>
 * If an exception is thrown by an {@link EntryProcessor}, the exception will
 * be returned wrapped in an {@link CacheException}. No changes will be made
 * to the {@link Cache.Entry}.
 * <p/>
 * Implementations may execute {@link EntryProcessor}s in situ, thus avoiding
 * locking, round-trips and expensive network transfers.
 * <p/>
 * {@link Cache.Entry} access, via a call to
 * {@link javax.cache.Cache.MutableEntry#getValue()}, will behave as if
 * {@link Cache#get(Object)} was called for the key. This includes updating
 * necessary statistics, consulting the configured
 * {@link javax.cache.expiry.ExpiryPolicy} and loading from a configured
 * {@link javax.cache.integration.CacheLoader}.
 * <p/>
 * {@link Cache.Entry} mutation, via a call to
 * {@link javax.cache.Cache.MutableEntry#setValue(Object)}, will behave
 * as if {@link Cache#put(Object, Object)} was called for the key. This
 * includes updating necessary statistics, consulting the configured
 * {@link javax.cache.expiry.ExpiryPolicy}, notifying
 * {@link javax.cache.event.CacheEntryListener}s and a writing to a configured
 * {@link javax.cache.integration.CacheWriter}.
 * <p/>
 * {@link Cache.Entry} removal, via a call to
 * {@link javax.cache.Cache.MutableEntry#remove()}, will behave
 * as if {@link Cache#remove(Object)} was called for the key. This
 * includes updating necessary statistics, notifying
```

```

* {@link javax.cache.event.CacheEntryListener}s and causing a delete on a
* configured {@link javax.cache.integration.CacheWriter}.
* <p/>
* As implementations may choose to execute {@link EntryProcessor}s remotely,
* {@link EntryProcessor}s, together with specified parameters and return
* values, may be required to implement {@link java.io.Serializable}.
*
* @param <K> the type of keys maintained by this cache
* @param <V> the type of cached values
*/
public interface EntryProcessor<K, V, T> {

    /**
     * Process an entry.
     *
     * @param entry      the entry
     * @param arguments a number of arguments to the process.
     * @return the result of the processing, if any, which is user defined.
     */
    T process(Cache.MutableEntry<K, V> entry, Object... arguments);
}

```

To invoke an `EntryProcessor` on a Cache Entry, applications must use the `Cache.invoke` method.

```

/**
 * Invokes an {@link EntryProcessor} against the {@link Entry} specified by
 * the provided key. If an {@link Entry} does not exist for the specified
 * key, one is loaded (if a loader is configured) or an empty {@link Entry}
 * is created.
 *
 * @param key          the key to the entry
 * @param entryProcessor the {@link EntryProcessor} to invoke
 * @param arguments    additional arguments to pass to the
 *                    {@link EntryProcessor}
 * @return the result of the processing, if any, defined by the
 *         {@link EntryProcessor} implementation
 *
 * @throws NullPointerException if key or {@link EntryProcessor} are null
 * @throws IllegalStateException if the cache is {@link #isClosed()}
 * @throws CacheException      if an exception occurred while executing
 *                             the {@link EntryProcessor} (the causing
 *                             exception will be wrapped by the
 *                             CacheException)
 * @throws ClassCastException   if the implementation supports and is
 *                             configured to perform runtime-type-checking,
 *                             and the key or value types are incompatible
 *                             with those that have been configured for the
 *                             {@link Cache}
 */

```



```

    * @see EntryProcessor
    */
    <T> T invoke(K key,
                EntryProcessor<K, V, T> entryProcessor,
                Object... arguments);

```

The following example demonstrates atomically incrementing the value of an entry with an `EntryProcessor`.

```

CachingProvider provider = Caching.getCachingProvider();
CacheManager manager = provider.getCacheManager();

MutableConfiguration<String, Integer> configuration =
    new MutableConfiguration<String, Integer>()
        .setTypes(String.class, Integer.class);

manager.createCache("example", configuration);
Cache<String, Integer> cache = manager.getCache(
    "example", configuration);

String key = "counter";
cache.put(key, 1);

int previous = cache.invoke(key,
                             new IncrementProcessor<String>());

assert previous == 1;
assert cache.get(key) == 2;

```

With the `IncrementProcessor` `Entry Processor` being defined as follows:

```

/**
 * An {@link Cache.EntryProcessor} that increments an {@link Integer}.
 *
 * @param <K> the type of keys
 */
public static class IncrementProcessor<K>
    implements Cache.EntryProcessor<K, Integer, Integer> {

    @Override
    public Integer process(Cache.MutableEntry<K, Integer> entry,
                          Object... arguments) {
        if (entry.exists()) {
            Integer current = entry.getValue();
            entry.setValue(current + 1);
            return current;
        } else {
            entry.setValue(0);
            return -1;
        }
    }
}

```

```
    }  
  }  
}
```

Implementations that support remote or distributed Caching Topologies may choose to execute Entry Processors in a remote process. In such circumstances implementations may require EntryProcessors, the invocation parameters and return types to implement java.lang.Serializable or be serializable in some manner. Alternatively implementations may choose to simply serialize the EntryProcessor class name, together with the specified invocation parameters and execute the request remotely by instantiating the EntryProcessor class and calling it with the deserialized invocation parameters.

As the outcome of an EntryProcessor is atomic, so are the interactions with Cache Loaders, Cache Writers, Cache Entry Listeners and Expiry Policies.

A application will never observe the intermediate events or side-effect for individual calls to MutableEntry getValue, setValue, remove etc while an EntryProcessor is being invoked. Rather applications will only observe the "net" result of an operation performed by an EntryProcessor on a Cache Entry.

For Example: An Entry Processor that has the following calls:

```
V v1 = entry.getValue();  
entry.setValue(v2);  
entry.remove();  
entry.setValue(v3);  
v4 = entry.getValue();
```

Will produce a single Cache Entry Listener event; an update from v1 to v3.

Like other `javax.cache.Cache` methods that support operations against multiple Cache Entries, `EntryProcessors` may also be invoked across multiple entries.

```
/**  
 * Invokes an {@link EntryProcessor} against the set of {@link Entry}s  
 * specified by the set of keys. If an {@link Entry} does not exist for a  
 * specified key, an attempt is made to load it (if a loader is configured)  
 * or an empty {@link Entry} is created and used instead.  
 * <p/>  
 * The order in which the entries for the keys are processed is undefined.  
 * Implementations may choose to process the entries in any order, including  
 * concurrently. Furthermore there is no guarantee implementations will
```

```

* use the same {@link EntryProcessor} instance to process each entry, as
* the case may be in a non-local cache topology.
*
* @param keys          the set of keys for entries to process
* @param entryProcessor the {@link EntryProcessor} to invoke
* @param arguments    additional arguments to pass to the
*                    {@link EntryProcessor}
* @return the map of results of the processing per key, if any, defined by the
*         {@link EntryProcessor} implementation. No mappings will be
*         returned for {@link EntryProcessor}s that return a <code>null</code>
*         value for a key
*
* @throws NullPointerException if keys or {@link EntryProcessor} are null
* @throws IllegalStateException if the cache is {@link #isClosed()}
* @throws CacheException      if an exception occurred while executing
*                             the {@link EntryProcessor} (the causing
*                             exception will be wrapped by the
*                             CacheException)
* @throws ClassCastException  if the implementation supports and is
*                             configured to perform runtime-type-checking,
*                             and the key or value types are incompatible
*                             with those that have been configured for the
*                             {@link Cache}
* @see EntryProcessor
*/
<T> Map<K, T> invokeAll(Set<? extends K> keys,
                      EntryProcessor<K, V, T> entryProcessor,
                      Object... arguments);

```

Caching Providers

Caching Providers are a core concept of the Java Caching API. It is through a `CachingProvider` that developers acquire `CacheManagers`, from which they interact with `Caches`.

A `CachingProvider` provides a means of:

- acquiring a default `CacheManager` instance
- establishing `CacheManager` instances, uniquely identified by an implementation specific URIS

e.g. an implementation might request a `CacheManager` configured declaratively with an implementation specific configuration file on the classpath, which the implementation allows to be loaded with a resource path.

```
cachingProvider.getCacheManager("/sample/ConfigurationFile.xml");
```

- scoping and managing `CacheManager` instances by URI and `ClassLoader`
- closing and releasing specific or collections of related `CacheManagers`
- querying the capabilities of a `CachingProvider` implementation, including support for optional features.

CacheManager Identity and Configuration

`CacheManagers` are logically identified by the URI that was used to establish them within the context of a `CachingProvider`. While applications often use the default URI as defined by a `CachingProvider` as a means of acquiring a `CacheManager`, applications may additionally use implementation specific URIs for advanced configuration of `CacheManagers`.

For example an implementation may permit a URI to be used as the location of a configuration file, say for pre-configured caches.

```
cachingProvider.getCacheManager("/sample/ConfigurationFile.xml");
```

Two or more `CacheManagers` defined with the same URI within an application deployment are said to be logically equivalent and depending on the implementation may manage the same caches.

For example two applications using the same URI with an implementation supporting distributed caching topologies may logically share the same cache names and content. In such situations, changes to cache entries in one application may be visible to the other application.

Alternatively two applications using the same URI with an implementation that only supports local

caching topologies may use the same cache names, but will not share cache content. In such situations changes to cache entries in one application may not be visible to the other application.

The following table outlines how a CacheManager URI may effect visibility of caches of the same name for application deployments using implementations that support only local (non-shared) v's distributed (or shared) cache topologies.

CacheManager URI	Local (non-shared) Cache Topology	Distributed / Shared Cache Topology
Same	Caches will have the same configurations . Caches may have different entries .	Caches will have the same configurations . Caches will have the same entries .
Different	Caches may have different configurations . Caches may have different entries .	Caches may have different configurations . Caches may have different entries .

The CachingProvider interface is defined as follows:

```
package javax.cache.spi;

import javax.cache.CacheManager;
import javax.cache.configuration.OptionalFeature;
import java.net.URI;
import java.util.Properties;

/**
 * Provides mechanisms to create, request and later manage the life-cycle of
 * configured {@link CacheManager}s, identified by {@link URI}s and scoped by
 * {@link ClassLoader}s.
 * <p/>
 * The meaning and semantics of the {@link URI} used to identify a
 * {@link CacheManager} is implementation dependent. For applications to remain
 * implementation independent, they should avoid attempting to create {@link URI}s
 * and instead use that which is returned by {@link #getDefaultURI()}.
 *
 * @author Brian Oliver
 */
public interface CachingProvider extends Closeable {

    /**
     * Requests a {@link CacheManager} configured according to the implementation
     * specific {@link URI} be made available that uses the provided
     * {@link ClassLoader} for loading underlying classes.
     * <p/>

```

```

* Multiple calls to this method with the same {@link URI} and
* {@link ClassLoader} must return the same {@link CacheManager} instance,
* except if a previously returned {@link CacheManager} has been closed.
* <p/>
* Properties are used in construction of a {@link CacheManager} and do not form
* part of the identity of the CacheManager. i.e. if a second call is made to
* with the same {@link URI} and {@link ClassLoader} but different properties,
* the {@link CacheManager} created in the first call is returned.
*
* @param uri          an implementation specific URI for the
*                    {@link CacheManager} (null means use
*                    {@link #getDefaultURI()})
* @param classLoader the {@link ClassLoader} to use for the
*                    {@link CacheManager} (null means use
*                    {@link #getDefaultClassLoader()})
* @param properties  the {@link Properties} for the {@link CachingProvider}
*                    to create the {@link CacheManager} (null means no
*                    implementation specific Properties are required)
* @throws javax.cache.CacheException when a {@link CacheManager} for the
*                                     specified arguments could not be produced
*/

/**
 * Obtains the default {@link ClassLoader} that will be used by the
 * {@link CachingProvider}.
 *
 * @return the default {@link ClassLoader} used by the {@link CachingProvider}
 */
ClassLoader getDefaultClassLoader();

/**
 * Obtains the default {@link URI} for the {@link CachingProvider}.
 * <p/>
 * Use this method to obtain a suitable {@link URI} for the
 * {@link CachingProvider}.
 *
 * @return the default {@link URI} for the {@link CachingProvider}
 */
URI getDefaultURI();

/**
 * Obtains the default {@link Properties} for the {@link CachingProvider}.
 * <p/>
 * Use this method to obtain suitable {@link Properties} for the
 * {@link CachingProvider}.
 *
 * @return the default {@link Properties} for the {@link CachingProvider}
 */
Properties getDefaultProperties();

```

```

/**
 * Requests a {@link CacheManager} configured according to the implementation
 * specific {@link URI} that uses the provided {@link ClassLoader} for loading
 * underlying classes.
 * <p/>
 * Multiple calls to this method with the same {@link URI} and
 * {@link ClassLoader} must return the same {@link CacheManager} instance,
 * accept if a previously returned {@link CacheManager} has been closed.
 *
 * @param uri          an implementation specific {@link URI} for the
 *                    {@link CacheManager} (null means
 *                    use {@link #getDefaultURI()})
 * @param classLoader the {@link ClassLoader} to use for the
 *                    {@link CacheManager} (null means
 *                    use {@link #getDefaultClassLoader()})
 * @throws javax.cache.CacheException when a {@link CacheManager} for the
 *                                     specified arguments could not be produced
 */
CacheManager getCacheManager(URI uri, ClassLoader classLoader);

/**
 * Requests a {@link CacheManager} configured according to the
 * {@link #getDefaultURI()} and {@link #getDefaultProperties()} be made
 * available that using the {@link #getDefaultClassLoader()} for loading
 * underlying classes.
 * <p/>
 * Multiple calls to this method must return the same {@link CacheManager}
 * instance, accept if a previously returned {@link CacheManager} has been
 * closed.
 */
CacheManager getCacheManager();

/**
 * Closes all of the {@link CacheManager} instances and associated resources
 * created and maintained by the {@link CachingProvider} across all
 * {@link ClassLoader}s.
 * <p/>
 * After closing the {@link CachingProvider} will still be operational. It
 * may still be used for acquiring {@link CacheManager} instances, though
 * those will now be new.
 */
void close();

/**
 * Closes all {@link CacheManager} instances and associated resources created
 * by the {@link CachingProvider} using the specified {@link ClassLoader}.
 * <p/>
 * After closing the {@link CachingProvider} will still be operational. It
 * may still be used for acquiring {@link CacheManager} instances, though

```

```

* those will now be new for the specified {@link ClassLoader} .
*
* @param classLoader the {@link ClassLoader} to release
*/
void close(ClassLoader classLoader);

/**
* Closes all {@link CacheManager} instances and associated resources created
* by the {@link CachingProvider} for the specified {@link URI} and
* {@link ClassLoader} .
*
* @param uri the {@link URI} to release
* @param classLoader the {@link ClassLoader} to release
*/
void close(URI uri, ClassLoader classLoader);

/**
* Determines whether an optional feature is supported by the
* {@link CachingProvider}.
*
* @param optionalFeature the feature to check for
* @return true if the feature is supported
*/
boolean isSupported(OptionalFeature optionalFeature);
}

```

Although optional, in Java SE environments the primary means of acquiring a `CachingProvider` instance is to use the `Caching` bootstrap class.

The `Caching` bootstrap provides three mechanisms for locating and instantiating one or more available `CachingProvider` implementations by:

- assuming implementations are defined as a `Service` and resolving them through the use of a `java.util.ServiceLoader`
- allowing a developer to specify the default implementation by using the `javax.cache.CachingProvider` Java System Property to define the fully qualified class name of the desired `CachingProvider`
- allowing applications to explicitly request a specific implementation using the fully qualified class name of the desired `CachingProvider`

While developers may alternatively use implementation dependent techniques for acquiring `CachingProviders` doing so may reduce the portability of their applications between `CachingProvider` implementations.

In Java EE environments, `CachingProviders` are used internally to resolve `CacheManagers`, those of

which are used to inject the required Caches into applications.

For a `CachingProvider` implementation to be automatically located by the Caching bootstrap class `java.util.ServiceLoader`, the fully qualified class name(s) of the `CachingProvider` implementation(s) an application will use must be defined in a `META-INF/services/javax.cache.spi.CachingProvider` configuration file as described in the JAR File Specification.

The `javax.cache.spi.CachingProvider` configuration file serves to define the specific `CachingProvider` implementation class(es) to the Caching bootstrap class, thus allowing it to automatically locate, load and provide appropriate instances to applications on request.

The content of a `javax.cache.spi.CachingProvider` configuration file is simply one or more fully qualified class names, each on a separate line, each specifying the name of an available `CachingProvider` implementation.

For example:

A Java Caching API implementor, ACME Caching Products, ships a JAR called `acme.jar`, that of which contains a `CachingProvider` implementation. The contents of the JAR includes both the `CachingProvider` implementation and the `javax.cache.spi.CachingProvider` configuration file.

```
META-INF/services/javax.cache.spi.CachingProvider
com/acme/cache/ACMECachingProvider.class
...
```

The contents of the `META-INF/services/javax.cache.spi.CachingProvider` file is nothing more than the name of the implementation class:

```
com.acme.cache.ACMECachingProvider
```

Applications may use multiple `CachingProvider` implementations simply by correctly configuring the `META-INF/services/javax.cache.spi.CachingProvider` file. When multiple `CachingProviders` are available, a request to return the default `CachingProvider` from the Caching bootstrap class will result in an exception.

The methods defined by the Caching bootstrap class are defined as follows:

```
package javax.cache;

import javax.cache.spi.CachingProvider;
import java.security.AccessController;
import java.security.PrivilegedAction;
import java.util.Iterator;
import java.util.LinkedHashMap;
import java.util.ServiceLoader;
import java.util.WeakHashMap;
```

```

/**
 * The {@link Caching} class provides a convenient means for an application to
 * acquire an appropriate {@link CachingProvider} implementation.
 * <p/>
 * While defined as part of the specification, its use is not mandatory.
 * Applications and/or containers may instead choose to directly instantiate a
 * {@link CachingProvider} implementation based on implementation specific
 * instructions.
 * <p/>
 * When using the {@link Caching} class, {@link CachingProvider} implementations
 * are automatically discovered when they follow the conventions outlined by the
 * Java Development Kit {@link ServiceLoader} class.
 * <p/>
 * For a {@link CachingProvider} to be automatically discoverable by the
 * {@link Caching} class, the fully qualified class name of the
 * {@link CachingProvider} implementation must be declared in the following
 * file:
 * <pre>
 *     META-INF/services/javax.cache.spi.CachingProvider
 * </pre>
 * that of which is resolvable via the class path.
 * <p/>
 * For example, in the reference implementation the contents of this file are:
 * <code>org.jsr107.ri.RICachingProvider</code>
 * <p/>
 * Alternatively when the fully qualified class name of a
 * {@link CachingProvider} implementation is specified using the system property
 * <code>javax.cache.CachingProvider</code>, that implementation will be used
 * as the default {@link CachingProvider}.
 * <p/>
 * All {@link CachingProvider}s that are automatically detected or explicitly
 * declared and loaded by the {@link Caching} class are maintained in an
 * internal registry. Consequently when a previously loaded
 * {@link CachingProvider} is requested, it will be simply returned from the
 * internal registry, without reloading and/or instantiating the said
 * implementation again.
 * <p/>
 * As required by some applications and containers, multiple co-existing
 * {@link CachingProvider}s implementations, from the same or different
 * implementors are permitted at runtime.
 * <p/>
 * To iterate through those that are currently registered a developer may use
 * the following methods:
 * <ol>
 * <li>{@link #getCachingProviders()}</li>
 * <li>{@link #getCachingProviders(ClassLoader)}</li>
 * </ol>
 * To request a specific {@link CachingProvider} implementation, a developer
 * should use either the {@link #getCachingProvider(String)} or
 * {@link #getCachingProvider(String, ClassLoader)} method.
 * <p/>
 * Where multiple {@link CachingProvider}s are present, the
 * {@link CachingProvider} returned by getters {@link #getCachingProvider()} and
 * {@link #getCachingProvider(ClassLoader)} is undefined and as a result a
 * {@link CacheException} will be thrown when attempted.

```

```

*
* @author Brian Oliver
* @author Greg Luck
* @see java.util.ServiceLoader
* @see javax.cache.spi.CachingProvider
*/
public final class Caching {

    /**
     * Obtains the {@link ClassLoader} to use for API methods that don't
     * explicitly require a {@link ClassLoader} but internally require one.
     * <p/>
     * By default this is the {@link Thread#getContextClassLoader()}.
     *
     * @return the default {@link ClassLoader}
     */
    public static ClassLoader getDefaultClassLoader()

    /**
     * Set the {@link ClassLoader} to use for API methods that don't explicitly
     * require a {@link ClassLoader}, but internally use one.
     *
     * @param classLoader the {@link ClassLoader} or <code>null</code> if the
     *     calling {@link Thread#getContextClassLoader()} should
     *     be used
     */
    public void setDefaultClassLoader(ClassLoader classLoader)

    /**
     * Obtains the single {@link CachingProvider} visible to the default
     * {@link ClassLoader}, which is {@link Thread#getContextClassLoader()}.
     *
     * @return the {@link CachingProvider}
     * @throws CacheException should zero, or more than one
     *     {@link CachingProvider} be available on the
     *     classpath, or it could not be loaded
     */
    public static CachingProvider getCachingProvider()

    /**
     * Obtains the single {@link CachingProvider} visible to the specified
     * {@link ClassLoader}.
     *
     * @param classLoader the {@link ClassLoader} to use for loading the
     *     {@link CachingProvider}
     * @return the {@link CachingProvider}
     * @throws CacheException should zero, or more than one
     *     {@link CachingProvider} be available on the
     *     classpath, or it could not be loaded
     * @see #getCachingProviders(ClassLoader)
     */
    public static CachingProvider getCachingProvider(ClassLoader classLoader)

    /**
     * Obtains the {@link CachingProvider}s that are available via the
     * {@link #getDefaultClassLoader()}.

```

```

* <p/>
* If a <code>javax.cache.cachingprovider</code> system property is defined,
* only that {@link CachingProvider} specified by that property is returned.
* Otherwise all {@link CachingProvider}s that are available via a
* {@link ServiceLoader} for {@link CachingProvider}s using the default
* {@link ClassLoader} (including those previously requested via
* {@link #getCachingProvider(String)}) are returned.
*
* @return an {@link Iterable} of {@link CachingProvider}s loaded by the
*         specified {@link ClassLoader}
*/
public static Iterable<CachingProvider> getCachingProviders()

/**
* Obtains the {@link CachingProvider}s that are available via the specified
* {@link ClassLoader}.
* <p/>
* If a <code>javax.cache.cachingprovider</code> system property is defined,
* only that {@link CachingProvider} specified by that property is returned.
* Otherwise all {@link CachingProvider}s that are available via a
* {@link ServiceLoader} for {@link CachingProvider}s using the specified
* {@link ClassLoader} (including those previously requested via
* {@link #getCachingProvider(String, ClassLoader)}) are returned.
*
* @param classLoader the {@link ClassLoader} of the returned
*                   {@link CachingProvider}s
* @return an {@link Iterable} of {@link CachingProvider}s loaded by the
*         specified {@link ClassLoader}
*/
public static Iterable<CachingProvider> getCachingProviders(
    ClassLoader classLoader)

/**
* Obtain the {@link CachingProvider} that is implemented by the specified
* fully qualified class name using the {@link #getDefaultClassLoader()}.
* Should this {@link CachingProvider} already be loaded it is simply returned,
* otherwise an attempt will be made to load and instantiate the specified
* class (using a no-args constructor).
*
* @param fullyQualifiedClassName the fully qualified class name of the
*                               {@link CachingProvider}
* @return the {@link CachingProvider}
* @throws CacheException if the {@link CachingProvider} cannot be created
*/
public static CachingProvider getCachingProvider(String fullyQualifiedClassName)

/**
* Obtain the {@link CachingProvider} that is implemented by the specified
* fully qualified class name using the provided {@link ClassLoader}.
* Should this {@link CachingProvider} already be loaded it is returned,
* otherwise an attempt will be made to load and instantiate the specified
* class (using a no-args constructor).
*
* @param fullyQualifiedClassName the fully qualified class name of the
*                               {@link CachingProvider}
* @param classLoader             the {@link ClassLoader} to load the

```

```
*                                     {@link CachingProvider}
* @return the {@link CachingProvider}
* @throws CacheException if the {@link CachingProvider} cannot be created
*/
public static CachingProvider getCachingProvider(
    String fullyQualifiedClassName,
    ClassLoader classLoader)
}
```

Transactions

Transactions are an optional requirement of this specification. Transactions have the meaning provided by the JTA specification^[10]. If implemented, transactions must work as specified here.

Two transaction modes are supported:

- Global Transactions, also known as XA Transactions.
- Local Transactions, where the transaction boundary is the CacheManager.

The motivation for providing transactions is that it is often very important for caches to stay strongly consistent with databases, message queues and other XA Resources. Without transaction support cache entries will not be guaranteed to be consistent with these.

A given `Cache` can support one of Local or XA transactions but not both at the same time.

All or nothing

If a transactions are enabled on a cache, all operations on it must happen within a transaction context otherwise a `javax.cache.transaction.TransactionException` will be thrown.

XA Transactions

XA Transactions for caches will work as per the JTA specification and the isolation level chosen.

XA transactions require the presence of a Transaction Manager.

An attempt to operate on an XA cache outside of an XA transaction context will throw a `CacheException`.

Example 1

An example using programmatic transaction control is:

```
//Get a global transaction assuming in a Java EE app server
UserTransaction utxg = jndiContext.lookup("java:comp/UserTransaction");

// start the transactions
utxg.begin();

// do work
cache1.put("key1", "value");
cache2.remove("key3");
cache3.put("key5", "value4");

// commit the transactions
utxg.commit();
```

Enlistment

Enlistment is the process by which the Transaction Manager is told that an XA Resource is going to participate in a transaction.

The `javax.transaction.xa` package defines the contract between the transaction manager and the resource manager, which allows the transaction manager to enlist and delist resource objects (supplied by the resource manager driver) in JTA transactions.

Enlistments is done using `TransactionManager.getTransaction().enlistResource(XAResource xaRes)`. The way in which a reference is obtained to a `TransactionManager` is not defined by the JTA specification. Java EE application servers typically use a well-known JNDI path to obtain that reference which is vendor specific.

On the first XA Resource operation `start()` is called by the `TransactionManager`.

XA is connection oriented. Caches are connection agnostic which creates an impedance mismatch.

Because we do not close connections, `XAResource.end()` is never called.

We expect the Transaction Managers to call `end()` for us before the two-phase commit protocol is started (even though this is not specified in the JTA specification). This is the behaviour of most existing Transaction Managers.

The JTA spec requires “Interleaving multiple transaction contexts using the same resource may be done by the transaction manager as long as `XAResource.start` and `XAResource.end` are invoked properly for each transaction context switch. Each time the resource is used with a different transaction, the method `XAResource.end` must be invoked for the previous transaction that was associated with the resource, and `XAResource.start` must be invoked for the current transaction context.”

Because we do not call `end` if the `XAResource` was at the `CacheManager` level this would imply only a single transaction could be done at a time across the `CacheManager`. This is a possible implementation.

It is suggested that a new `XAResource` is created for each transaction so that a single cache manager will have as many `XAResources` open as there are transactions. In this way concurrent transactions are supported. The interleaving issue is also avoided.

Another possible implementation is to create an `XAResource` per cache operation. This is highly concurrent but requires more calls to the expensive `enlistResource()` method.

Recovery

Caches must implement recovery protocols as defined by JTA. In particular

`XAResource.recover()` must be supported.

In the case of a local in-process cache where there is no durable store, and the process has restarted thus coming up with an empty cache, it is acceptable for `recover()` to return an empty array of `javax.transaction.xa.Xid[]`. The Transaction Manager may report a Heuristic Exception in this case. *This will not prevent pending transactions being correctly recovered for other XAResources. Further because the local cache is empty any attempt to read an affected value will be a cache miss so the user logic will go elsewhere for the value.*

Local Transactions

A cache supporting transactions will support Local Transactions with the four isolation levels.

Local transactions do not require the presence of a Transaction Manager.

Local Transactions allow single phase commit across multiple cache operations against one or more caches in the same CacheManager, whether distributed or local. This lets you apply multiple changes to a CacheManager all within a single transaction. If you also want to apply changes to other resources such as a database then you need to open a transaction to them and manually handle commit and rollback to ensure consistency. (Or use XA Transactions).

For example, we have two puts to Cache A and one remove to Cache B, and 4 puts to Cache C. These can all be accommodated in a single local transaction.

The JTA API is used to control local transactions. The `javax.transaction.UserTransaction` interface provides the application the ability to control transaction boundaries programmatically.

```
//Get a transaction
UserTransaction utx = cacheManager.getUserTransaction();
// start a transaction
utx.begin();
// do work
cache1.put("key1", "value");
cache2.remove("key3");

// commit the work
utx.commit();
```

The best practice as with all local transactions is to place these steps in a try-catch block and call `rollback()` if an exception is thrown. See the JTA spec for details^[10].

An attempt to operate on a cache outside of a local transaction context will cause a `javax.cache.transaction.TransactionException`.

It is possible for a single thread to have begun a XA Transaction and a local transaction. Cache operations will then be accepted for both XA and local transaction caches because both transaction

contexts exist. However the transactions are separate.

So another programmatic (bean managed in EJB language) example showing this where cache1 and cache2 are configured for Local Transactions and cache3 is configured for XA Transactions would be:

```
//Get a local transaction
UserTransaction utxl = cacheManager.getUserTransaction();

//Get a global transaction assuming in a Java EE app server
UserTransaction utxg = jndiContext.lookup("java:comp/UserTransaction");

// start the transactions
utxl.begin();
utxg.begin();

// do work
cache1.put("key1", "value");
cache2.remove("key3");
cache3.put("key5", "value4");

// commit the transactions
utxl.commit();
utxg.commit();
```

Though this works, it is not particularly useful as one transaction can succeed on commit and the other can fail.

Local Transactions have their own exceptions that can be thrown, which are all subclasses of CacheException. They are:

- `TransactionException` - a general exception
- `TransactionInterruptedException` - if `Thread.interrupt()` got called while the cache was processing a transaction
- `TransactionTimeoutException` - if a cache operation or commit is called after the transaction timeout has elapsed

Recovery

This is relevant to XA Transactions. For local transactions if you crash before `commit()` then the changes will depend on your isolation level.

Isolation Levels

The isolation level for a cache must be set at creation time and remains immutable for the lifetime of the cache.

The isolation levels `READ_COMMITTED`, `READ_UNCOMMITTED`, `REPEATABLE_READ` and `SERIALIZABLE` are required.

READ_COMMITTED

Mutating changes are not visible to other transactions in a local cache or a distributed cache until `COMMIT` has been called.

Until then Implementations are free to either:

- *return the old copy*
- *block until commit or rollback is called*

Both approaches satisfy the `READ_COMMITTED` isolation level.

READ_UNCOMMITTED

Cache mutations may be visible to other transactions in a local cache or a distributed cache, subject to any propagation delay of the implementation, as if transactions were not being used.

On `commit()`, no value changes will be observed.

On `rollback()`, the values will be reverted to their previous values, which will of course be a visible change.

On timeout, the JTA specification states that rollback is called. So on timeout the old values will be reverted too. Exactly when the rollback occurs will be implementation dependent.

SERIALIZABLE

Mutating changes are not visible to other transactions in a local cache or a distributed cache until `commit()` has been called.

Further no changes to the cache made by other transactions are visible to this transaction until it completes.

The `SERIALIZABLE` isolation level offers one further protection over `REPEATABLE_READ`, protection from phantom reads.

An alternative is to exclusively write lock a collection of keys of interest before starting your

transaction. We could use lockAll(Collection keys). This would create a read-write lock. Other transactions would block until this transaction.

This behaviour could be achieved pessimistically with a ReadWrite lock over the entire cache or also achieved optimistically by triggering a RollbackException if any changes made to the keys used (for reads or writes) in this transaction have been made.

REPEATABLE_READ

Mutating changes are not visible to other transactions in a local cache or a distributed cache until `commit()` has been called.

No changes to the cache made by other transactions are visible to this transaction once they have been read or written by this transaction.

This behaviour could be achieved pessimistically with a read-write lock acquired over the keys as they are used or also achieved optimistically by triggering a RollbackException if any changes made to the keys used (for writes) in this transaction have been made.

Caching Annotations

Caching annotations provide method interceptors for user supplied classes which interact with caches. The annotations and support classes are in the `javax.cache.annotation` package. The following annotations are defined:

- `@CacheDefaults`
- `@CacheResult`
- `@CacheRemoveEntry`
- `@CacheRemoveAll`
- `@CachePut`

Each annotation defines the underlying cache operations that are to be performed using the Java API. The same result must be achieved whether by annotation or by using the defined Java API operations.

Annotations therefore provide an additional API for interacting with caches. Annotations are provided only for the most commonly used cache methods.

In order to use annotations in an application, a library or framework which processes these annotations and intercepts calls to annotated application objects is required. In an application, the method and configuration of processing caching annotations on classes is left to the implementation.

This would commonly be provided by a dependency injection framework such as defined by CDI in Java EE. The RI includes example implementations for use with CDI, Spring and Guice.

Annotations

`@CacheDefaults`

This is a class level annotation used to define default property values for all caching related annotations used in a class. The `cacheName`, `cacheResolverFactory`, and `cacheKeyGenerator` properties may be specified though all are optional.

If `@CacheDefaults` is specified on a class but no method level caching annotations exist then the `@CacheDefaults` annotation is ignored.

The following example specifies a cache named “cities” as the default cache name for annotations in the class. The `@CacheResult` annotation on the `getCity` method will use this cache name at runtime.

Example 1

This example shows how to use the `@CacheDefaults` Annotation.

```
@CacheDefaults(cacheName="cities")
public class CitySource {
    @CacheResult
```

```

    public City getCity(int lat, int lon) {
        //...
    }
}

```

@CacheResult

This is a method level annotation used to mark methods whose returned value is cached, using a key generated from the method parameters, and returned from the cache on later calls with the same parameters.

It is defined as follows:

```

/**
 * When a method annotated with {@link CacheResult} is invoked a
 * {@link GeneratedCacheKey} will be generated and
 * {@link javax.cache.Cache#get(Object)} is called before the annotated method
 * actually executes. If a value is found in the cache it is returned and the
 * annotated method is never actually executed. If no value is found the
 * annotated method is invoked and the returned value is stored in the cache
 * with the generated key.
 * <p/>
 * null return values are cached by default, this behavior can be disabled by
 * setting the {@link #cacheNull()} property to false.
 * <p/>
 * Exceptions are not cached by default. Caching of exceptions can be enabled by
 * specifying an {@link #exceptionCacheName()}. If an exception cache is specified
 * it is checked before invoking the annotated method and if a cached exception is
 * found it is re-thrown.
 * <p/>
 * The {@link #cachedExceptions()} and {@link #nonCachedExceptions()} properties
 * can be used to control which exceptions are cached and which are not.
 * <p/>
 * To always invoke the annotated method and still cache the result set
 * {@link #skipGet()} to true. This will disable the pre-invocation
 * {@link javax.cache.Cache#get(Object)} call. If {@link #exceptionCacheName()} is
 * specified the pre-invocation exception check is also disabled. This feature is
 * useful for methods that create or update objects to be cached.
 * <p/>
 * Example of caching the Domain object with a key generated from the
 * <code>String</code> and <code>int</code> parameters.
 * <p/>
 * With no {@link #cacheName()} specified a cache name of
 * "my.app.DomainDao.getDomain(java.lang.String,int)" will be generated.
 * <p><blockquote><pre>
 * package my.app;
 * <p/>
 * public class DomainDao {
 *     &#64;CacheResult
 *     public Domain getDomain(String domainId, int index) {
 *         ...
 *     }
 * }
 * </pre></blockquote></p>

```

```

* Example using the {@link GeneratedCacheKey} annotation so that only the domainId
* parameter is used in key generation:
* <p><blockquote><pre>
* package my.app;
* </pre>
* public class DomainDao {
*     &#64;CacheResult
*     public Domain getDomain(@CacheKey String domainId, Monitor mon) {
*         ...
*     }
* }
* </pre></blockquote></p>
* If exception caching is enabled via specification of
* {@link #exceptionCacheName()} the following rules are used to determine if a
* thrown exception is cached:
* <ol>
* <li>If {@link #cachedExceptions()} and {@link #nonCachedExceptions()} are both
* empty then all exceptions are cached</li>
* <li>If {@link #cachedExceptions()} is specified and
* {@link #nonCachedExceptions()} is not specified then only exceptions
* which pass an instanceof check against the cachedExceptions list are cached</li>
* <li>If {@link #nonCachedExceptions()} is specified and
* {@link #cachedExceptions()} is not specified then all exceptions
* which do not pass an instanceof check against the nonCachedExceptions list are
* cached</li>
* <li>If {@link #cachedExceptions()} and {@link #nonCachedExceptions()} are both
* specified then exceptions which pass an instanceof check against the
* cachedExceptions list but do not pass an instanceof check against the
* nonCachedExceptions list are cached</li>
* </ol>
*
* @author Eric Dalquist
* @author Rick Hightower
* @see CacheKey
* @since 1.0
*/
@Target({ElementType.METHOD, ElementType.TYPE})
@Retention(RetentionPolicy.RUNTIME)
public @interface CacheResult {

    /**
     * The name of the cache.
     * <p/>
     * If not specified defaults first to {@link CacheDefaults#cacheName()} an if
     * that is not set it defaults to:
     * package.name.ClassName.methodName (package.ParameterType,package.ParameterType)
     */
    @Nonbinding String cacheName() default "";

    /**
     * If set to true the pre-invocation {@link javax.cache.Cache#get(Object)} is
     * skipped and the annotated method is always executed with the returned value
     * being cached as normal. This is useful for create or update methods which
     * should always be executed and have their returned value placed in the cache.
     * <p/>
     * If true and an {@link #exceptionCacheName()} is specified the pre-invocation

```

```

* check for a thrown exception is also skipped. If an exception is thrown during
* invocation it will be cached following the standard exception caching rules.
* <p/>
* Defaults to false.
*
* @see CachePut
*/
@Nonbinding boolean skipGet() default false;

/**
* If set to false null return values will not be cached. If true (the default)
* null return values will be cached.
* <p/>
* Defaults to true.
*/
@Nonbinding boolean cacheNull() default true;

/**
* The {@link CacheResolverFactory} used to find the {@link CacheResolver} to
* use at runtime.
* <p/>
* The default resolver pair will resolve the cache by name from the default
* {@link javax.cache.CacheManager}
*/
@Nonbinding Class<? extends CacheResolverFactory> cacheResolverFactory()
    default CacheResolverFactory.class;

/**
* The {@link CacheKeyGenerator} to use to generate the {@link GeneratedCacheKey}
* for interacting with the specified Cache.
* <p/>
* Defaults to a key generator that uses
* {@link java.util.Arrays#deepHashCode(Object[])} and
* {@link java.util.Arrays#deepEquals(Object[], Object[])} with the array
* returned by {@link CacheKeyInvocationContext#getKeyParameters()}
*
* @see CacheKey
*/
@Nonbinding Class<? extends CacheKeyGenerator> cacheKeyGenerator()
    default CacheKeyGenerator.class;

/**
* The name of the cache to cache exceptions.
* <p/>
* If not specified no exception caching is done.
*/
@Nonbinding String exceptionCacheName() default "";

/**
* Defines zero (0) or more exception {@link Class classes}, which must be a
* subclass of {@link Throwable}, indicating which exception types which
* <b>must</b> be cached. Only consulted if {@link #exceptionCacheName()} is
* specified.
*/
@Nonbinding Class<? extends Throwable>[] cachedExceptions() default {};

```

```

/**
 * Defines zero (0) or more exception {@link Class Classes}, which must be a
 * subclass of {@link Throwable}, indicating which exception types
 * <b>must not</b> be cached. Only consulted if {@link #exceptionCacheName()}
 * is specified.
 */
@Nonbinding Class<? extends Throwable>[] nonCachedExceptions() default {};
}

```

The `@CacheKey` annotation can be used to select a subset of the parameters for key generation.

Options

1. Toggle caching of null return values via the `cacheNull` property.
2. Optional caching and re-throwing of exceptions with their own named cache, includes the ability to only cache specific exceptions.
3. `skipGet`. Optional skipping of the pre-execution `Cache.get` call, useful when the annotated method should always be executed and the returned value placed in the cache.

`@CacheResult` will be ignored if placed on static methods.

@CachePut

This is a method level annotation used to mark methods where one of the method parameters should be stored in the cache. One parameter must be annotated with `@CacheValue` marking it as the parameter to be cached. If no `@CacheValue` annotation is specified a `CacheAnnotationConfigurationException` must be thrown either at application initialization time or on method invocation.

It is defined as follows:

```

/**
 * When a method annotated with {@link CachePut} is invoked a {@link
 * GeneratedCacheKey} will be generated and {@link javax.cache.Cache#put(Object,
 * Object)} will be invoked on the specified cache storing the value marked with
 * {@link CacheValue}. Null values are cached by default but this behavior can be
 * disabled via the {@link #cacheNull()} property.
 * <p/>
 * The default behavior is to call {@link javax.cache.Cache#put(Object, Object)}
 * after the annotated method is invoked, this behavior can be changed by setting
 * {@link #afterInvocation()} to false in which case
 * {@link javax.cache.Cache#put(Object, Object)} will be called before the
 * annotated method is invoked.
 * <p/>
 * Example of caching the Domain object with a key generated from the String and
 * int parameters. The {@link CacheValue} annotation is used to designate which
 * parameter should be stored in the "domainDao" cache.
 * <p><blockquote><pre>
 * package my.app;

```



```

* <p/>
* public class DomainDao {
*     &#64;CachePut(cacheName="domainCache")
*     public void updateDomain(String domainId, int index, &#64;CacheValue Domain
* domain) {
*         ...
*     }
* }
* </pre></blockquote></p>
* Exception Handling, only used if {@link #afterInvocation()} is true.
* <ol>
* <li>If {@link #cacheFor()} and {@link #noCacheFor()} are both empty then all
* exceptions prevent the put</li>
* <li>If {@link #cacheFor()} is specified and {@link #noCacheFor()} is not
* specified then only exceptions which pass an instanceof check against the
* cacheFor list result in a put</li>
* <li>If {@link #noCacheFor()} is specified and {@link #cacheFor()} is not
* specified then all exceptions which do not pass an instanceof check against the
* noCacheFor result in a put</li>
* <li>If {@link #cacheFor()} and {@link #noCacheFor()} are both specified then
* exceptions which pass an instanceof check against the cacheFor list but do not
* pass an instanceof check against the noCacheFor list result in a put</li>
* </ol>
*
* @author Eric Dalquist
* @author Rick Hightower
* @see CacheValue
* @see CacheKey
* @since 1.0
*/
@Target({ElementType.METHOD, ElementType.TYPE})
@Retention(RetentionPolicy.RUNTIME)
public @interface CachePut {

    /**
     * The name of the cache.
     * <p/>
     * If not specified defaults first to {@link CacheDefaults#cacheName()} an if
     * that is not set it
     * defaults to:
     * package.name.ClassName.methodName(package.ParameterType,package.ParameterType)
     */
    @Nonbinding String cacheName() default "";

    /**
     * When {@link javax.cache.Cache#put(Object, Object)} should be called. If true
     * it is called after the annotated method
     * invocation completes successfully. If false it is called before the annotated
     * method is invoked.
     * <p/>

```

```

* Defaults to true.
* <p/>
* If true and the annotated method throws an exception the rules governing
* {@link #cacheFor()} and {@link #noCacheFor()}
* will be followed.
*/
@Nonbinding boolean afterInvocation() default true;

/**
* If set to false null {@link CacheValue} values will not be cached. If true
* (the default) null {@link CacheValue}
* values will be cached.
* <p/>
* Defaults to true
*/
@Nonbinding boolean cacheNull() default true;

/**
* The {@link CacheResolverFactory} used to find the {@link CacheResolver} to
* use at runtime.
* <p/>
* The default resolver pair will resolve the cache by name from the default
* {@link javax.cache.CacheManager}
*/
@Nonbinding Class<? extends CacheResolverFactory> cacheResolverFactory()
    default CacheResolverFactory.class;

/**
* The {@link CacheKeyGenerator} to use to generate the {@link
* GeneratedCacheKey} for interacting with the specified Cache.
* <p/>
* Defaults to a key generator that uses
* {@link java.util.Arrays#deepHashCode(Object[])}
* and
* {@link java.util.Arrays#deepEquals(Object[], Object[])} with the array
* returned by
* {@link CacheKeyInvocationContext#getKeyParameters()}
*
* @see CacheKey
*/
@Nonbinding Class<? extends CacheKeyGenerator> cacheKeyGenerator()
    default CacheKeyGenerator.class;

/**
* Defines zero (0) or more exception {@link Class} classes, which must be a
* subclass of {@link Throwable}, indicating which exception types <b>must</b>
* cause
* the parameter to be cached. Only used if {@link #afterInvocation()} is true.
*/
@Nonbinding Class<? extends Throwable>[] cacheFor() default {};

```

```

/**
 * Defines zero (0) or more exception {@link Class Classes}, which must be a
 * subclass of {@link Throwable}, indicating which exception types <b>must
 * not</b>
 * cause the parameter to be cached. Only used if {@link #afterInvocation()} is
 * true.
 */
@Nonbinding Class<? extends Throwable>[] noCacheFor() default {};
}

```

The `@CacheKey` annotation can be used to select a subset of the parameters for key generation. The `@CacheValue` annotated parameter is never included in key generation.

Options

1. Toggle caching of null parameter values via the `cacheNull` property.
2. Specify if the `Cache.put` call will happen before or after method execution.
3. If caching happens after invocation then an exception thrown by the annotated method can cancel the `Cache.put` call.

`@CachePut` will be ignored if placed on static methods.

`@CacheRemoveEntry`

This is a method level annotation used to mark methods where the invocation results in an entry being removed from the specified `Cache`.

It is defined as follows:

```

/**
 * When a method annotated with {@link CacheRemoveEntry} is invoked a {@link
 * GeneratedCacheKey} will be generated and
 * {@link javax.cache.Cache#remove(Object)} will be invoked on the specified
 * cache.
 * <p/>
 * The default behavior is to call {@link javax.cache.Cache#remove(Object)} after
 * the annotated method is invoked, this behavior can be changed by setting
 * {@link #afterInvocation()} to false in which case
 * {@link javax.cache.Cache#remove(Object)} will be called before the annotated
 * method is invoked.
 * <p/>
 * Example of removing a specific Domain object from the "domainCache". A {@link
 * GeneratedCacheKey} will be generated from the String and int parameters and
 * used to call {@link javax.cache.Cache#remove(Object)} after the deleteDomain
 * method completes successfully.
 * <p><blockquote><pre>
 * package my.app;
 * <p/>
 * public class DomainDao {
 *     &#64;CacheRemoveEntry(cacheName="domainCache")

```

```

*   public void deleteDomain(String domainId, int index) {
*       ...
*   }
* }
* </pre></blockquote></p>
* Exception Handling, only used if {@link #afterInvocation()} is true.
* <ol>
* <li>If {@link #evictFor()} and {@link #noEvictFor()} are both empty then all
* exceptions prevent the remove</li>
* <li>If {@link #evictFor()} is specified and {@link #noEvictFor()} is not
* specified then only exceptions which pass an instanceof check against the
* evictFor list result in a
* remove</li>
* <li>If {@link #noEvictFor()} is specified and {@link #evictFor()} is not
* specified then all exceptions which do not pass an instanceof check against the
* noEvictFor result in a
* remove</li>
* <li>If {@link #evictFor()} and {@link #noEvictFor()} are both specified then
* exceptions which pass an instanceof check against the evictFor list but do not
* pass an instanceof check against the noEvictFor list result in a remove</li>
* </ol>
*
* @author Eric Dalquist
* @author Rick Hightower
* @see CacheKey
* @since 1.0
*/
@Target({ElementType.METHOD, ElementType.TYPE})
@Retention(RetentionPolicy.RUNTIME)
public @interface CacheRemoveEntry {

    /**
     * The name of the cache.
     * <p/>
     * If not specified defaults first to {@link CacheDefaults#cacheName()},
     * and if that is not set then to:
     * package.name.ClassName.methodName (package.ParameterType,package.ParameterType)
     */
    @Nonbinding String cacheName() default "";

    /**
     * When {@link javax.cache.Cache#remove(Object)} should be called. If true it
     * is called after the annotated method invocation completes successfully. If
     * false it is called before the annotated method is invoked.
     * <p/>
     * Defaults to true.
     * <p/>
     * If true and the annotated method throws an exception the put will not be
     * executed.
     */
    @Nonbinding boolean afterInvocation() default true;

    /**
     * The {@link CacheResolverFactory} used to find the {@link CacheResolver} to
     * use at runtime.
     * <p/>

```

```

* The default resolver pair will resolve the cache by name from the default
* {@link javax.cache.CacheManager}
*/
@Nonbinding Class<? extends CacheResolverFactory> cacheResolverFactory()
    default CacheResolverFactory.class;

/**
* The {@link CacheKeyGenerator} to use to generate the {@link
* GeneratedCacheKey} for interacting with the specified Cache.
* <p/>
* Defaults to a key generator that uses
* {@link java.util.Arrays#deepHashCode(Object[])}
* and {@link java.util.Arrays#deepEquals(Object[], Object[])} with the array
* returned by {@link CacheKeyInvocationContext#getKeyParameters()}
*
* @see CacheKey
*/
@Nonbinding Class<? extends CacheKeyGenerator> cacheKeyGenerator()
    default CacheKeyGenerator.class;

/**
* Defines zero (0) or more exception {@link Class classes}, which must be a
* subclass of {@link Throwable}, indicating which exception types must cause
* a cache evict. Only used if {@link #afterInvocation()} is true.
*/
@Nonbinding Class<? extends Throwable>[] evictFor() default {};

/**
* Defines zero (0) or more exception {@link Class Classes}, which must be a
* subclass of {@link Throwable}, indicating which exception types must
* <b>not</b> cause a cache evict. Only used if {@link #afterInvocation()} is
* true.
*/
@Nonbinding Class<? extends Throwable>[] noEvictFor() default {};
}

```

The `@CacheKey` annotation can be used to select a subset of the parameters for key generation.

Options

1. Specify if the `Cache.remove` call will happen before or after method execution
2. If removal happens after invocation then an exception thrown by the annotated method can cancel the `Cache.remove` call.

`@CacheRemoveEntry` will be ignored if placed on static methods.

`@CacheRemoveAll`

This is a method level annotation used to mark methods where the invocation results in all entries being removed from the specified Cache.

It is defined as follows:

```

/**
 * When a method annotated with {@link CacheRemoveAll} is invoked all elements in
 * the specified cache will be removed via the
 * {@link javax.cache.Cache#removeAll()} method
 * <p/>
 * The default behavior is to call {@link javax.cache.Cache#removeAll()} after the
 * annotated method is invoked, this behavior can be changed by setting {@link
 * #afterInvocation()} to false in which case {@link javax.cache.Cache#removeAll()}
 * will be called before the annotated method is invoked.
 * <p/>
 * Example of removing all Domain objects from the "domainCache". {@link
 * javax.cache.Cache#removeAll()} will be called after deleteAllDomains() returns
 * successfully.
 * <p><blockquote><pre>
 * package my.app;
 * <p/>
 * public class DomainDao {
 *     &#64;CacheRemoveAll(cacheName="domainCache")
 *     public void deleteAllDomains() {
 *         ...
 *     }
 * }
 * </pre></blockquote></p>
 * Exception Handling, only used if {@link #afterInvocation()} is true.
 * <ol>
 * <li>If {@link #evictFor()} and {@link #noEvictFor()} are both empty then all
 * exceptions prevent the removeAll</li>
 * <li>If {@link #evictFor()} is specified and {@link #noEvictFor()} is not
 * specified then only exceptions which pass an instanceof check against the
 * evictFor list result in a removeAll</li>
 * <li>If {@link #noEvictFor()} is specified and {@link #evictFor()} is not
 * specified then all exceptions which do not pass an instanceof check against the
 * noEvictFor result in a removeAll</li>
 * <li>If {@link #evictFor()} and {@link #noEvictFor()} are both specified then
 * exceptions which pass an instanceof check against the evictFor list but do not
 * pass an instanceof check against the noEvictFor list result in a removeAll</li>
 * </ol>
 *
 * @author Eric Dalquist
 * @author Rick Hightower
 * @since 1.0
 */
@Target({ElementType.METHOD, ElementType.TYPE})
@Retention(RetentionPolicy.RUNTIME)
public @interface CacheRemoveAll {

    /**
     * /**
     * The name of the cache.

```

```

* <p/>
* If not specified defaults first to {@link CacheDefaults#cacheName()} an if
* that is not set it defaults to:
* package.name.ClassName.methodName(package.ParameterType,package.ParameterType)
*/
@Nonbinding String cacheName() default "";

/**
* When {@link javax.cache.Cache#removeAll()} should be called. If true it is
* called after the annotated method
* invocation completes successfully. If false it is called before the annotated
* method is invoked.
* <p/>
* Defaults to true.
* <p/>
* If true and the annotated method throws an exception the put will not be
* executed.
*/
@Nonbinding boolean afterInvocation() default true;

/**
* The {@link CacheResolverFactory} used to find the {@link CacheResolver} to
* use
* at runtime.
* <p/>
* The default resolver pair will resolve the cache by name from the default
* {@link javax.cache.CacheManager}
*/
@Nonbinding Class<? extends CacheResolverFactory> cacheResolverFactory()
    default CacheResolverFactory.class;

/**
* Defines zero (0) or more exception {@link Class classes}, which must be a
* subclass of {@link Throwable}, indicating which exception types must cause
* a cache removeAll. Only used if {@link #afterInvocation()} is true.
*/
@Nonbinding Class<? extends Throwable>[] evictFor() default {};

/**
* Defines zero (0) or more exception {@link Class Classes}, which must be a
* subclass of {@link Throwable}, indicating which exception types must
* <b>not</b>
* cause a cache removeAll. Only used if {@link #afterInvocation()} is true.
*/
@Nonbinding Class<? extends Throwable>[] noEvictFor() default {};
}

```

Options

1. Specify if the `Cache.removeAll` call will happen before or after method execution.
2. If removal happens after invocation then an exception thrown by the annotated method can cancel the `Cache.removeAll` call.

`@CacheRemoveAll` will be ignored if placed on static methods.

@CacheKey

This is a parameter level annotation used to mark parameters that are used to generate the `GeneratedCacheKey` via the `CacheKeyGenerator`. At execution time the values of the parameters annotated with `@CacheKey` are placed in the `CacheKeyInvocationContext.getKeyParameters()` array.

Usable with:

- `@CacheResult`,
- `@CachePut`, and
- `@CacheRemoveEntry`

@CacheValue

This is a parameter level annotation used to mark the parameter to be cached for a method annotated with `@CachePut`. A parameter annotated with `@CachePut` will never be included in the `CacheKeyInvocationContext.getKeyParameters()` array.

Usable with:

- `@CachePut`

Example 2

This example shows usage of many of the above annotation.

```
/**
 * An implementation of BlogManager that uses a variety of annotations
 * @author Rick Hightower
 */
@CacheDefaults(cacheName = "blgMngr")
public class ClassLevelCacheConfigBlogManagerImpl implements BlogManager {

    private static Map<String, Blog> map = new HashMap<String, Blog>();

    @CacheResult
    public Blog getEntryCached(String title) {
        return map.get(title);
    }
}
```



```

public Blog getEntryRaw(String title) {
    return map.get(title);
}

/**
 * @see manager.BlogManager#clearEntryFromCache(java.lang.String)
 */
@CacheRemoveEntry
public void clearEntryFromCache(String title) {
}

public void clearEntry(String title) {
    map.put(title, null);
}

@CacheRemoveAll
public void clearCache() {
}

public void createEntry(Blog blog) {
    map.put(blog.getTitle(), blog);
}

@CacheResult
public Blog getEntryCached(String randomArg, @CacheKey String title,
                           String randomArg2) {
    return map.get(title);
}

```

Cache Resolution

All of the method level annotations allow for specification of a `CacheResolverFactory` and cache name which are used to determine the Cache to interact with at runtime.

It is defined below:

```

/**
 * Determines the {@link CacheResolver} to use for an annotated method. The
 * {@link CacheResolver} will be retrieved once per annotated method.
 * <p/>
 * Implementations MUST be thread-safe.
 *
 * @author Eric Dalquist
 * @since 1.0
 */
public interface CacheResolverFactory {

```

```

/**
 * Get the {@link CacheResolver} used at runtime for resolution of the
 * {@link javax.cache.Cache} for the {@link CacheResult}, {@link CachePut},
 * {@link CacheRemoveEntry}, or {@link CacheRemoveAll} annotation.
 *
 * @param cacheMethodDetails The details of the annotated method to get the
 *                             {@link CacheResolver} for. @return The {@link
 *                             CacheResolver} instance to be
 *                             used by the interceptor.
 */
CacheResolver getCacheResolver(CacheMethodDetails<? extends Annotation>
                               cacheMethodDetails);

/**
 * Get the {@link CacheResolver} used at runtime for resolution of the {@link
 * javax.cache.Cache} for the {@link CacheResult} annotation to cache exceptions.
 * <p/>
 * Will only be called if {@link CacheResult#exceptionCacheName()} is not empty.
 *
 * @param cacheMethodDetails The details of the annotated method to get the
 *                             {@link CacheResolver} for.
 * @return The {@link CacheResolver} instance to be used by the interceptor.
 */
CacheResolver getExceptionCacheResolver(CacheMethodDetails<CacheResult>
                                         cacheMethodDetails);
}

```

Cache Name

If no cache name is specified either on the method level annotation or at the class level with the `@CacheDefaults` annotation then the name is generated as the following:

```
package.name.ClassName.methodName(package.ParameterType,package.Parameter
Type)
```

The `@CacheResult` annotation has an additional `exceptionCacheName` property. If this property is not specified there is no default exception cache name and no exception cache is used.

CacheResolverFactory

The specified `CacheResolverFactory` must be called exactly once per annotated method to determine the `CacheResolver` to use for each execution of the annotated method. When an annotated method is executed the previously retrieved `CacheResolver` is used to determine the `Cache` to use based on the `CacheInvocationContext`.

If `javax.cache.annotation.CacheResolverFactory` is specified on the annotation and the `@CacheDefaults` then the default `CacheResolverFactory` logic must be used.

The default `CacheResolverFactory` does the following, in order:

1. Get the `CacheManager` to use via:

```

CachingProvider provider = Caching.getCachingProvider();
CacheManager cacheManager =
provider.getCacheManager(provider.getDefaultURI(),
provider.getDefaultClassLoader());

```

2. Call `CacheManager.getCache(String cacheName)` with the cache name
3. If a `Cache` is not returned, a default cache is created using:

```

Cache cache = cacheManager.createCache(cacheName, new
MutableConfiguration<Object, Object>());

```

4. Create a `CacheResolver` that wraps the found/created `Cache` and always returns the `Cache`.

If the `CacheResolverFactory` throws an exception the exception must be propagated up to the application code that triggered the execution of the `CacheResolverFactory`.

CacheResolver

The `CacheResolver` is returned by the `CacheResolver` factory and is meant to be called on every invocation of the annotated method it was returned for, returning the `Cache` to use for that invocation.

It is defined as follows:

```

/**
 * Determines the {@link Cache} to use for an intercepted method invocation.
 * <p/>
 * Implementations MUST be thread-safe.
 *
 * @author Eric Dalquist
 * @see CacheResolverFactory
 * @since 1.0
 */
public interface CacheResolver {

    /**
     * Resolve the {@link Cache} to use for the {@link CacheInvocationContext}.
     *
     * @param cacheInvocationContext The context data for the intercepted method
     *                               invocation
     * @return The {@link Cache} instance to be used by the interceptor
     */
    <K, V> Cache<K, V> resolveCache(CacheInvocationContext<? extends Annotation>
                                cacheInvocationContext);
}

```

If the `CacheResolver` throws an exception the exception must be propagated up to the application

code that triggered the execution of the `CacheResolverFactory`.

Key Generation

The `@CacheResult`, `@CachePut`, and `@CacheRemoveEntry` annotations all require a cache key to be generated and all of these annotations allow for specification of a `CacheKeyGenerator` implementation.

The specified `CacheKeyGenerator` will be called once for every annotated method invocation. Information about the annotated method and the current invocation is provided by the `CacheKeyInvocationContext`. The method parameters the developer specified to be used in the key are contained in the `CacheInvocationParameter` array returned by the `getKeyParameters()` method. A custom `CacheKeyGenerator` can use whatever information at its disposal to generate the `GeneratedCacheKey`.

If `javax.cache.annotation.CacheKeyGenerator` is specified on the annotation and the `@CacheDefaults` then the default `CacheKeyGenerator` logic must be used.

Default `CacheKeyGenerator` Rules:

1. Create an `Object[]` using `CacheInvocationParameter.getValue()` from the array returned by `CacheKeyInvocationContext.getKeyParameters()`
2. Create a `CacheKey` instance that wraps the `Object[]` and uses `Arrays.deepHashCode` to calculate its `hashCode` and `Arrays.deepEquals` for comparison to other keys.

If the `CacheKeyGenerator` throws an exception the exception must be propagated up to the application code that triggered the execution of the `CacheKeyGenerator`.

Annotation Support Classes

`CacheMethodDetails`

Static information about a method with a caching annotation. Used by the `CacheResolverFactory` to determine the `CacheResolver` to use at runtime.

It is defined as follows:

```
/**
 * Static information about a method annotated with one of:
 * {@link CacheResult}, {@link CachePut}, {@link CacheRemoveEntry}, or {@link
 * CacheRemoveAll}
 * <p/>
 * Used with {@link CacheResolverFactory#getCacheResolver(CacheMethodDetails)} to
 * determine the {@link CacheResolver} to use with the method.
 *
 * @param <A> The type of annotation this context information is for. One of
 *             {@link
 *             javax.cache.annotation.CacheResult},
 *             {@link javax.cache.annotation.CachePut}, {@link
 *             javax.cache.annotation.CacheRemoveEntry}, or
```

```

*           {@link javax.cache.annotation.CacheRemoveAll}.
* @author Eric Dalquist
* @version $Revision$
* @see CacheResolverFactory
*/
public interface CacheMethodDetails<A extends Annotation> {
    /**
     * The annotated method
     *
     * @return The annotated method
     */
    Method getMethod();

    /**
     * An immutable Set of all Annotations on this method
     *
     * @return An immutable Set of all Annotations on this method
     */
    Set<Annotation> getAnnotations();

    /**
     * The caching related annotation on the method.
     * One of: {@link CacheResult}, {@link CachePut}, {@link CacheRemoveEntry}, or
     * {@link CacheRemoveAll}
     *
     * @return The caching related annotation on the method.
     */
    A getCacheAnnotation();

    /**
     * The cache name resolved by the implementation.
     * <p/>
     * The cache name is determined by first looking at the cacheName attribute of
     * the method level annotation. If that attribute is not set then the class
     * level {@link CacheDefaults} annotation is checked. If that annotation does
     * not exist or does not have its cacheName attribute set then the following
     * cache name generation rules are followed:
     * <p/>
     * "fully qualified class name"."method name"("fully qualified parameter class
     * names")
     * <p/>
     * For example:
     * <p><blockquote><pre>
     * package my.app;
     * <p/>
     * public class DomainDao {
     *     &#64;CacheResult
     *     public Domain getDomain(String domainId, int index) {
     *         ...
     *     }
     * }
     * </pre></blockquote></p>

```

```

* }
* </pre></blockquote></p>
* Results in the cache name: "my.app.DomainDao.getDomain(java.lang.String,int)"
*
* @return The fully resolved cache name
*/
String getCacheName();
}

```

CacheInvocationContext

Runtime information about the execution of a method with a caching annotation. Used by the `CacheResolver` to determine the Cache to use. Extends `CacheMethodDetails` so all static information is also available.

It is defined as follows:

```

/**
 * Runtime information about an intercepted method invocation for a method
 * annotated with {@link CacheResult}, {@link CachePut}, {@link CacheRemoveEntry},
 * or {@link CacheRemoveAll}
 * <p/>
 * Used with {@link CacheResolver#resolveCache(CacheInvocationContext)} to
 * determine the {@link javax.cache.Cache} to use at runtime for the method
 * invocation.
 *
 * @param <A> The type of annotation this context information is for. One of
 *     {@link
 *     javax.cache.annotation.CacheResult},
 *     {@link javax.cache.annotation.CachePut}, {@link
 *     javax.cache.annotation.CacheRemoveEntry}, or
 *     {@link javax.cache.annotation.CacheRemoveAll}.
 * @author Eric Dalquist
 * @version $Revision$
 * @see CacheResolver
 */
public interface CacheInvocationContext<A extends Annotation>
    extends CacheMethodDetails<A> {

    /**
     * @return The object the intercepted method was invoked on.
     */
    Object getTarget();

    /**
     * Returns a clone of the array of all method parameters.
     *
     * @return An array of all parameters for the annotated method
     */
}

```

```

CacheInvocationParameter[] getAllParameters();

/**
 * Return an object of the specified type to allow access to the
 * provider-specific API. If the provider's
 * implementation does not support the specified class, the {@link
 * IllegalArgumentException} is thrown.
 *
 * @param cls the class of the object to be returned. This is normally either the
 *           underlying implementation class or an interface that it
 *           implements.
 * @return an instance of the specified class
 * @throws IllegalArgumentException if the provider doesn't support the
 *           specified
 *           class.
 */
<T> T unwrap(java.lang.Class<T> cls);
}

```

CacheKeyInvocationContext

Runtime information about the execution of a method where key generation will take place (annotated with one of `@CacheResult`, `@CachePut`, or `@CacheRemoveEntry`). Used by the `CacheKeyGenerator` to create the `GeneratedCacheKey` to use. Extends `CacheInvocationContext` so all standard runtime and static information is also available.

It is defined as follows:

```

/**
 * Runtime information about an intercepted method invocation for a method
 * annotated with {@link CacheResult}, {@link CachePut}, or
 * {@link CacheRemoveEntry}.
 * <p/>
 * Used with {@link CacheKeyGenerator#generateCacheKey(CacheKeyInvocationContext)}
 * to generate a {@link GeneratedCacheKey} for the invocation.
 *
 * @param <A> The type of annotation this context information is for. One of
 *           {@link javax.cache.annotation.CacheResult},
 *           {@link javax.cache.annotation.CachePut}, or
 *           {@link javax.cache.annotation.CacheRemoveEntry}
 * @author Eric Dalquist
 * @see CacheKeyGenerator
 */
public interface CacheKeyInvocationContext<A extends Annotation>
    extends CacheInvocationContext<A> {

/**
 * Returns a clone of the array of all method parameters to be used by the
 * {@link

```

```

* CacheKeyGenerator} in creating a {@link GeneratedCacheKey}. The returned array
* may be the same as or a subset of the array returned by
* {@link #getAllParameters()}
* <p/>
* Parameters in this array are selected by the following rules:
* <ul>
* <li>If no parameters are annotated with {@link CacheKey} or {@link
* CacheValue}
* then all parameters are included</li>
* <li>If a {@link CacheValue} annotation exists and no {@link CacheKey} then
* all
* parameters except the one annotated with {@link CacheValue} are included</li>
* <li>If one or more {@link CacheKey} annotations exist only those parameters
* with the {@link CacheKey} annotation are included</li>
* </ul>
*
* @return An array of all parameters to be used in cache key generation
*/
CacheInvocationParameter[] getKeyParameters();

/**
* When a method is annotated with {@link CachePut} this is the parameter
* annotated with {@link CacheValue}
*
* @return The parameter to cache, will never be null for methods annotated with
*         {@link CachePut}, will be null for methods not annotated with {@link
*         CachePut}
*/
CacheInvocationParameter getValueParameter();
}

```

CacheInvocationParameter

Runtime information about a parameter for a method execution. Includes parameter annotations, position, type and value. Provided by CacheInvocationContext and CacheKeyInvocationContext

It is defined as follows:

```

/**
* A parameter to an intercepted method invocation. Contains the parameter value
* as well static type and annotation information about the parameter.
*
* @author Eric Dalquist
* @version $Revision$
*/
public interface CacheInvocationParameter {

/**
* The parameter type as declared on the method.

```



```

    */
    Class<?> getRawType();

    /**
     * @return The parameter value
     */
    Object getValue();

    /**
     * @return An immutable Set of all Annotations on this method parameter, never
     * null.
     */
    Set<Annotation> getAnnotations();

    /**
     * The index of the parameter in the original parameter array as returned by
     * {@link CacheInvocationContext#getAllParameters()}
     *
     * @return The index of the parameter in the original parameter array.
     */
    int getParameterPosition();
}

```

GeneratedCacheKey

Created by the `CacheKeyGenerator` interface the `GeneratedCacheKey` is used as the key in any cache interacted with by the annotations. All `GeneratedCacheKeys` must be immutable and serializable.

It is defined as follows:

```

/**
 * A {@link Serializable}, immutable, thread-safe object that is used as a key,
 * that of which is automatically generated using a CacheKeyGenerator.
 * <p/>
 * The implementation MUST follow the Java contract for {@link Object#hashCode()}
 * and {@link Object#equals(Object)} to ensure correct behavior.
 * <p/>
 * It is recommended that implementations also override {@link Object#toString()}
 * and provide a human-readable string representation of the key.
 *
 * @author Eric Dalquist
 * @see CacheKeyGenerator
 * @since 1.0
 */
public interface GeneratedCacheKey extends Serializable {

    /**
     * The immutable hash code of the cache key.

```

```

*
* @return The hash code of the object
* @see Object#hashCode()
*/
@Override
int hashCode();

/**
* Compare this {@link GeneratedCacheKey} with another. If the two objects
* are equal their {@link #hashCode()} values MUST be equal as well.
*
* @param o The other object to compare to.
* @return true if the objects are equal
* @see Object#equals(Object)
*/
@Override
boolean equals(Object o);
}

```

Annotations Interactions

Annotation Inheritance and Ordering

This specification defers to section 2.1 of the Common Annotations for Java specification^[2] for annotation inheritance. Order of interceptor execution with regards to annotations outside of this specification is not defined and left to the annotation support implementation.

Multiple Annotations

Only one method level caching annotation can be specified on a method and only one parameter level caching annotation can be specified on a parameter. If more than one annotation is specified on a method or on a parameter then a `CacheAnnotationConfigurationException` must be thrown either at application initialization time or on method invocation.

Transactions

If a cache is transactional, then a transaction context must exist when a caching annotated method is executed. If a transaction does not exist when the method is executed the `Cache` will throw a `CacheException`.

Management

The `javax.cache.management` package contains `MXBean` interfaces for cache management and statistics.

Enabling and Disabling

By default, both management and statistics are disabled. To turn them on at configuration time, use the following methods on `MutableConfiguration`:

- `setManagementEnabled(boolean enabled)` to turn on management
- `setStatisticsEnabled(boolean enabled)` to turn on statistics

To enable or disable them at runtime, the following methods are provided on `CacheManager`:

```
/**
 * Controls whether management is enabled. If enabled the
 * {@link javax.cache.management.CacheMXBean} for each cache is registered in
 * the platform MBean server. The platform MBeanServer is obtained using
 * {@link java.lang.management.ManagementFactory#getPlatformMBeanServer()}
 * <p/>
 * Management information includes the name and configuration information for
 * the cache.
 * <p/>
 * Each cache's management object must be registered with an ObjectName that
 * is unique and has the following type and attributes:
 * <p/>
 * Type:
 * <code>javax.cache:type=Cache</code>
 * <p/>
 * Required Attributes:
 * <ul>
 * <li>CacheManager the name of the CacheManager
 * <li>Cache the name of the Cache
 * </ul>
 *
 * @param cacheName the name of the cache to register
 * @param enabled true to enable management, false to disable.
 */
void enableManagement(String cacheName, boolean enabled);

/**
 * Enables or disables statistics gathering for a managed {@link Cache} at
 * runtime.
 * <p/>
 * Each cache's statistics object must be registered with an ObjectName that
 * is unique and has the following type and attributes:
 * <p/>
```

```

* Type:
* <code>javax.cache:type=CacheStatistics</code>
* <p/>
* Required Attributes:
* <ul>
* <li>CacheManager the name of the CacheManager
* <li>Cache the name of the Cache
* </ul>
*
* @param cacheName the name of the cache to register
* @param enabled true to enable statistics, false to disable.
* @throws IllegalStateException if the cache is {@link #isClosed()}
* @throws NullPointerException if cacheName is null
*/
void enableStatistics(String cacheName, boolean enabled);

```

MXBean Definitions

The CacheMXBean provides details of cache configuration and is defined as follows:

```

/**
 * A management bean for cache. It provides configuration information. It does not
 * allow mutation of configuration or mutation of the cache.
 * <p/>
 * Each cache's management object must be registered with an ObjectName that is
 * unique and has the following type and attributes:
 * <p/>
 * Type:
 * <code>javax.cache:type=Cache</code>
 * <p/>
 * Required Attributes:
 * <ul>
 * <li>CacheManager the name of the CacheManager
 * <li>Cache the name of the Cache
 * </ul>
 * <p/>
 *
 * @author Greg Luck
 * @author Yannis Cosmadopoulos
 * @since 1.0
 */
@MXBean
public interface CacheMXBean {

    /**
     * Determines if a {@link javax.cache.Cache} should operate in read-through mode.
     * <p/>
     * When in read-through mode, cache misses that occur due to cache entries
     * not existing as a result of performing a "get" call via one of

```

```

* {@link javax.cache.Cache#get(Object)},
* {@link javax.cache.Cache#getAll(java.util.Set)},
* {@link javax.cache.Cache#getAndRemove(Object)} and/or
* {@link javax.cache.Cache#getAndReplace(Object, Object)} will appropriately
* cause the configured {@link javax.cache.integration.CacheLoader} to be
* invoked.
* <p/>
* The default value is <code>>false</code>.
*
* @return <code>>true</code> when a {@link javax.cache.Cache} is in
* "read-through" mode.
* @see javax.cache.integration.CacheLoader
*/
boolean isReadThrough();

/**
* Determines if a {@link javax.cache.Cache} should operate in "write-through"
* mode.
* <p/>
* When in "write-through" mode, cache updates that occur as a result of
* performing "put" operations called via one of
* {@link javax.cache.Cache#put(Object, Object)},
* {@link javax.cache.Cache#getAndRemove(Object)},
* {@link javax.cache.Cache#removeAll()},
* {@link javax.cache.Cache#getAndPut(Object, Object)}
* {@link javax.cache.Cache#getAndRemove(Object)},
* {@link javax.cache.Cache#getAndReplace(Object, Object)},
* {@link javax.cache.Cache#invoke(Object,
* javax.cache.Cache.EntryProcessor, Object...)}
* will appropriately cause the configured
* {@link javax.cache.integration.CacheWriter} to be invoked.
* <p/>
* The default value is <code>>false</code>.
*
* @return <code>>true</code> when a {@link javax.cache.Cache} is in
* "write-through" mode.
* @see javax.cache.integration.CacheWriter
*/
boolean isWriteThrough();

/**
* Whether storeByValue (true) or storeByReference (false).
* When true, both keys and values are stored by value.
* <p/>
* When false, both keys and values are stored by reference.
* Caches stored by reference are capable of mutation by any threads holding
* the reference. The effects are:
* <ul>
* <li>if the key is mutated, then the key may not be retrievable or
* removable</li>

```

```

* <li>if the value is mutated, then all threads in the JVM can potentially
* observe those mutations,
* subject to the normal Java Memory Model rules.</li>
* </ul>
* Storage by reference only applies to the local heap. If an entry is moved off
* heap it will
* need to be transformed into a representation. Any mutations that occur after
* transformation
* may not be reflected in the cache.
* <p/>
* When a cache is storeByValue, any mutation to the key or value does not affect
* the key of value stored in the cache.
* <p/>
* The default value is <code>>true</code>.
*
* @return true if the cache is store by value
*/
boolean isStoreByValue();

/**
* Checks whether statistics collection is enabled in this cache.
* <p/>
* The default value is <code>>false</code>.
*
* @return true if statistics collection is enabled
*/
boolean isStatisticsEnabled();

/**
* Checks whether management is enabled on this cache.
* <p/>
* The default value is <code>>false</code>.
*
* @return true if management is enabled
*/
boolean isManagementEnabled();

/**
* Checks whether transactions are enabled for this cache.
* <p/>
* Note that in a transactional cache, entries being mutated within a
* transaction cannot be expired by the cache.
* <p/>
* The default value is <code>>false</code>.
*
* @return true if transaction are enabled
*/
boolean isTransactionsEnabled();

/**

```

```

    * Gets the transaction isolation level.
    * <p/>
    * The default value is {@link javax.cache.transaction.IsolationLevel#NONE}.
    *
    * @return the isolation level.
    */
    IsolationLevel getTransactionIsolationLevel();

/**
 * Gets the transaction mode.
 * <p/>
 * The default value is {@link javax.cache.transaction.Mode#NONE}.
 *
 * @return the the mode of the cache.
 */
    Mode getTransactionMode();
}

```

The `CacheStatisticsMXBean` provides statistics for a cache, and is defined as follows:

```

/**
 * Cache statistics.
 * <p/>
 * Statistics are accumulated from the time a cache is created. They can be reset
 * to zero using {@link #clear}.
 * <p/>
 * There are no defined consistency semantics for statistics. Refer to the
 * implementation for precise semantics.
 * <p/>
 * Each cache's statistics object must be registered with an ObjectName that is
 * unique and has the following type and attributes:
 * <p/>
 * Type:
 * <code>javax.cache:type=CacheStatistics</code>
 * <p/>
 * Required Attributes:
 * <ul>
 * <li>CacheManager the URI of the CacheManager
 * <li>Cache the name of the Cache
 * </ul>
 *
 * @author Greg Luck
 * @since 1.0
 */
@MXBean
public interface CacheStatisticsMXBean {

/**

```

```

* Clears the statistics counters to 0 for the associated Cache.
*/
void clear();

/**
* The number of get requests that were satisfied by the cache.
* <p/>
* In a caches with multiple tiered storage, a hit may be implemented as a hit
* to the cache or to the first tier.
*
* @return the number of hits
*/
long getCacheHits();

/**
* This is a measure of cache efficiency.
* <p/>
* It is calculated as:
* {@link #getCacheHits} divided by {@link #getCacheGets ()} * 100.
*
* @return the percentage of successful hits, as a decimal e.g 75.
*/
float getCacheHitPercentage();

/**
* A miss is a get request which is not satisfied.
* <p/>
* In a simple cache a miss occurs when the cache does not satisfy the request.
* <p/>
* In a caches with multiple tiered storage, a miss may be implemented as a miss
* to the cache or to the first tier.
* <p/>
* In a read-through cache a miss is an absence of the key in the cache which
* will trigger a call to a CacheLoader. So it is still a miss even though the
* cache will load and return the value.
* <p/>
* Refer to the implementation for precise semantics.
*
* @return the number of misses
*/
long getCacheMisses();

/**
* Returns the percentage of cache accesses that did not find a requested entry
* in the cache.
* <p/>
* This is calculated as {@link #getCacheMisses()} divided by
* {@link #getCacheGets()} * 100.
*
* @return the percentage of accesses that failed to find anything

```



```

*/
float getCacheMissPercentage();

/**
 * The total number of requests to the cache. This will be equal to the sum of
 * the hits and misses.
 * <p/>
 * A "get" is an operation that returns the current or previous value. It does
 * not include checking for the existence of a key.
 * <p/>
 * In a caches with multiple tiered storage, a gets may be implemented as a get
 * to the cache or to the first tier.
 *
 * @return the number of gets
 */
long getCacheGets();

/**
 * The total number of puts to the cache.
 * <p/>
 * A put is counted even if it is immediately evicted.
 * <p/>
 * Replaces, where a put occurs which overrides an existing mapping is counted
 * as a put.
 *
 * @return the number of hits
 */
long getCachePuts();

/**
 * The total number of removals from the cache. This does not include evictions,
 * where the cache itself initiates the removal to make space.
 *
 * @return the number of hits
 */
long getCacheRemovals();

/**
 * The total number of evictions from the cache. An eviction is a removal
 * initiated by the cache itself to free up space. An eviction is not treated as
 * a removal and does not appear in the removal counts.
 *
 * @return the number of evictions from the cache
 */
long getCacheEvictions();

/**
 * The mean time to execute gets.
 * <p/>
 * In a read-through cache the time taken to load an entry on miss is not

```

```

    * included in get time.
    *
    * @return the time in μs
    */
float getAverageGetTime();

/**
 * The mean time to execute puts.
 *
 * @return the time in μs
 */
float getAveragePutTime();

/**
 * The mean time to execute removes.
 *
 * @return the time in μs
 */
float getAverageRemoveTime();
}

```

Accessing Management Information

There are no accessor methods provided for either management or statistics. Instead the JMX infrastructure is used. When enabled if enabled the `MBeans` are registered with the platform MBean server. It is obtained using the `getPlatformMBeanServer()` method on `java.lang.management.ManagementFactory`.

The beans can then be obtained from the `MBeanServer` in the usual way as defined by JMX.

The convention for JMX attribute names follows the JavaBeans^[15] convention for properties. So, the accessor `getCacheHitPercentage()` on `CacheStatisticsMXBean` corresponds to the JMX attribute `CacheHitPercentage`.

Example 1

This example shows how to read the `CacheHitPercentage` for the cache named “simpleCache”.

```

CachingProvider cachingProvider = Caching.getCachingProvider();
CacheManager cacheManager = cachingProvider.getCacheManager();

MutableConfiguration<String, Integer> config =
    new MutableConfiguration<String, Integer>();
config.setStoreByValue(false)
    .setTypes(String.class, Integer.class)
    .setExpiryPolicyFactory(AccessedExpiryPolicy.factoryOf(ONE_HOUR))
    .setStatisticsEnabled(true);

```

```

cacheManager.createCache("simpleCache", config);
Cache<String, Integer> cache = cacheManager.getCache("simpleCache",
    String.class, Integer.class);

Set<ObjectName> registeredObjectNames = null;
MBeanServer mBeanServer = ManagementFactory.getPlatformMBeanServer();

ObjectName objectName = new ObjectName("javax.cache:type=CacheStatistics"
    + ",CacheManager=" + (cache.getCacheManager().getURI().toString())
    + ",Cache=" + cache.getName());
System.out.println(mBeanServer.getAttribute(objectName,
    "CacheHitPercentage"));
}

```

Statistics Effects of Cache Operations

The following table shows which cache operations affect the statistics counters. In the table some if a hit will occur if a mapping exists, and a miss if one does not the table will have yes in each column. If a cache is set to read-through mode, the lack of a mapping will cause a miss, even if a CacheLoader loads an entry and the cache operation returns it from the call.

Method	Puts	Removals	Hits	Misses
<code>boolean containsKey(K key)</code>	No	No	Yes	Yes
<code>V get(K key)</code>	No	No	Yes	Yes
<code>Map<K,V> getAll(Collection<? extends K> keys)</code>	No	No	Yes	Yes
<code>V getAndPut(K key, V value)</code>	Yes	No	Yes	Yes
<code>V getAndRemove(K key)</code>	No	Yes	Yes	Yes
<code>V getAndReplace(K key, V value)</code>	Yes	No	Yes	Yes
<code><T> T invoke(K key, EntryProcessor<K, V, T> entryProcessor, Object... arguments)</code>	Yes, if <code>setValue(V value)</code> was called.	Yes, if <code>remove()</code> was called.	Yes	Yes
<code><T> Map<K, T> invokeAll(Set<? extends K> keys, EntryProcessor<K, V, T> entryProcessor, Object... arguments);</code>	Yes, if <code>setValue(V value)</code> was called.	Yes, if <code>remove()</code> was called.	Yes	Yes
<code>Iterator<Cache.Entry<K, V>> iterator() ?</code>	No	Yes, if <code>remove()</code> was	Yes	No

		called.		
<code>void loadAll(Set<? extends K> keys, boolean replaceExistingValues, CompletionListener completionListener)</code>	No	No	No	No
<code>void put(K key, V value)</code>	Yes	No	No	No
<code>void putAll(Map<? extends K, ? extends V> map)</code>	Yes	No	No	No
<code>boolean putIfAbsent(K key, V value)</code>	Yes	No	Yes	Yes
<code>boolean remove(K key)</code>	No	Yes	No	No
<code>boolean remove(K key, V oldValue)</code>	No	Yes, if the method returns true	Yes	Yes
<code>void removeAll()</code>	No	Yes	No	No
<code>void removeAll(Set<? extends K> keys)</code>	No	Yes	No	No
<code>boolean replace(K key, V value)</code>	Yes	No	Yes	Yes
<code>boolean replace(K key, V oldValue, V newValue)</code>	Yes	No	Yes	Yes

Portability Recommendations

The following recommendations should be followed to improve application portability between implementations of the Java Cache API:

1. Custom Key classes must correctly implement `hashCode` and `equals`.
2. To support the default store-by-value `Cache` semantics, Custom Key and Value classes should be serializable.
3. Caches should not use forward slashes (/) or colons (:) as part of their names.
4. Applications should use default URIs and Properties when requesting `CacheManagers`.
5. For maximum portability, applications should avoid using optional features of the specification, or use the `CachingProvider.isSupported` method to exploit optional features when they are present.

For example, store-by-reference in-process implementations will have much higher performance than store-by-value because keys and values may be referenced directly.

6. Keep proprietary configuration in proprietary declarative configuration files rather than using proprietary programmatic `Cache` construction.
7. Avoid use of `Cache.unwrap` and `Cache.Entry.unwrap`.

These are used to gain access the proprietary backing `Cache` and `Cache.Entry` respectively. Using proprietary APIs reduces portability.

8. Do not make assumptions about `Cache` topology.

For example, assuming a listener will be executed locally, and creating a dependence on local application class instances, which it may do in one cache implementation but not in another.

Glossary

Application	A Java application that uses the Java Caching API.
Cache	A named and configured collection of Entries.
Cache Manager	A container for caches, which holds references to them.
Cache Operation	An invocation of a method on <code>Cache</code> .
Caching Provider	An implementation of this specification. See Caching Implementation
CacheLoader	A user-defined Class which is used to load key/value pairs into a <code>Cache</code> on demand.
CacheWriter	A user-defined Class which is used to write key/value pairs into a cache after a put operation.
Cache Store	A place where cache data is kept. Caches may have multiple stores.
CacheEventListener	A user-defined Class which listens to <code>Cache</code> events.
Compile Time	The time when source code is compiled into Java byte code
Configuration Time	The time when a new cache is being configured and before it is available for use
Developer	An application developer using the Java Caching API.
Entry	A cache entry, consisting of a unique key and a value
Eviction	The process of removing entries from a <code>Cache</code> when the <code>Cache</code> has exceeded a resource limit.
Expiry	The process of ensuring entries are no longer available to an application because they are no longer considered valid.
Expiry Policy	A policy that defines when a <code>Cache.Entry</code> is considered expired, and therefore should not be available to an application.
External Resource	A resource, external to the cache, loaded from or written to by a <code>CacheLoader</code> or <code>CacheWriter</code> respectively.
Implementer	The supplier of a caching implementation.
Implementation	An implementation of this specification.
Key	A way of unambiguously identifying a unique item in a <code>Cache</code> .
Operational Method	A method that when executed causes an operation to be performed eg: <code>Cache.put(...)</code> , <code>Cache.get(...)</code> . Non-Operational

	methods however are those that typically perform simple “status” requests and do not rely on underlying resources being available eg: <code>Cache.isClosed()</code> , <code>Cache.getName()</code> .
Read-Through	If a mapping is missing from the cache, a loader will be invoked to read data in from an external resource.
Runtime	The time when methods on a cache are invokable.
Store By Reference	The cache stores entries using their key and value references when using an on-heap store.
Store By Value	The cache stores entries not using their key and value references when using an on-heap store.
Value	The value stored in a <code>Cache</code> . Any <code>Java Object</code> can be a value.
Write-Through	On a cache mutation, a writer will be invoked to write data to an external resource.

Bibliography

- [1] JSR345: Enterprise JavaBeans, v. 3.2. EJB Core Contracts and Requirements. <http://jcp.org/en/jsr/proposalDetails?id=345>
- [2] JSR-250: Common Annotations for the Java™ Platform 1.1. <http://jcp.org/en/jsr/detail?id=250>.
- [3] JSR-175: A Metadata Facility for the Java™ Programming language. <http://jcp.org/en/jsr/detail?id=175>.
- [4] JSR-221: JDBC 4.1 Specification. <http://java.sun.com/products/jdbc>.
- [5] Enterprise JavaBeans, Simplified API, v 3.0. <http://java.sun.com/products/ejb>.
- [6] JAR File Specification, <http://java.sun.com/j2se/1.5.0/docs/guide/jar/jar.html>.
- [7] JSR-338: Java™ Persistence 2.1 <http://jcp.org/en/jsr/detail?id=338>
- [8] JSR-336: Java™ SE 7. <http://jcp.org/en/jsr/detail?id=336>
- [9] JSR342: Java™ Platform, Enterprise Edition 7 (Java EE 7) Specification, <http://www.jcp.org/en/jsr/detail?id=342>
- [10] JTA Specification 2.0, <http://jcp.org/en/jsr/detail?id=907>
- [11] CDI Specification 1.1, <http://jcp.org/en/jsr/detail?id=346>
- [12] Expression Language 3.0 <http://jcp.org/en/jsr/detail?id=341>
- [13] @CacheResult and Ehcache prior art <http://code.google.com/p/ehcache-spring-annotations/>
- [14] @CacheResult with Grails prior art <http://gpc.github.com/grails-springcache/docs/manual/guide/4.%20Content%20Caching.html>
- [15] Information on properly implementing equals and hashCode (<http://java.sun.com/developer/Books/effectivejava/Chapter3.pdf>)

Appendix A - Revision History

This appendix lists the significant changes that have been made during the development of JSR107.

Early Draft 1

Created initial draft. This document was incomplete.

Public Review Draft

Created first complete draft.

Second Public Review Draft

1. <https://github.com/jsr107/jsr107spec/issues/170> - Minor Typo.
2. <https://github.com/jsr107/jsr107spec/issues/169> - Minor Typo.
3. <https://github.com/jsr107/jsr107spec/issues/168> - Minor Typo.
4. <https://github.com/jsr107/jsr107spec/issues/167> - Minor Typo.
5. <https://github.com/jsr107/jsr107spec/issues/166> - Minor Typo.
6. <https://github.com/jsr107/jsr107spec/issues/162> - `DefaultCacheResolver` updated for new creational mechanism.
7. <https://github.com/jsr107/jsr107spec/pull/172> - Minor JavaDoc formatting.
8. <https://github.com/jsr107/jsr107spec/issues/184> Minor JavaDoc corrections on `CacheLoader`.
9. <https://github.com/jsr107/jsr107spec/issues/182> Fix omission in Section 8.4, invocations of listeners.
10. <https://github.com/jsr107/jsr107spec/issues/181> Clarify behaviour of when write-through is called by `invoke` and `invokeAll` in Section 7.3.
11. <https://github.com/jsr107/jsr107spec/issues/180> Section 7.2, Read-Through Caching. Add behaviour to `invoke` and `invokeAll` where `getValue()` is called.
12. <https://github.com/jsr107/jsr107spec/issues/178> Section 6. Expiry Policy. Changes to the method list table.
13. <https://github.com/jsr107/jsr107spec/issues/176> Minor Typos.
14. <https://github.com/jsr107/jsr107spec/issues/173> Contradiction around the read-through behaviour of `getAndRemove` and `getAndReplace`

15. <https://github.com/jsr107/jsr107spec/issues/179> Renamed `TouchedPolicy` to `TouchedExpiryPolicy`
16. <https://github.com/jsr107/jsr107spec/issues/191> Remove obsolete entries from the expert group listing.
17. <https://github.com/jsr107/jsr107spec/issues/190> Numerous spec doc typos, grammar and formatting edits after proof read.
18. <https://github.com/jsr107/jsr107spec/issues/188> Fix threading bug identified by FindBugs in `CompletionListenerFuture`
19. <https://github.com/jsr107/jsr107spec/issues/185> Change `MutableEntry.getValue()` to `Entry.getVlaue()` in JavaDoc which is more correct.
20. <https://github.com/jsr107/jsr107spec/issues/177> Changes to examples in Section 5.2, Type-Safety.
21. <https://github.com/jsr107/jsr107spec/issues/175> Removed `IllegalStateException` on `CacheManager.destroyCache(String cacheName)`
22. <https://github.com/jsr107/jsr107spec/issues/174> Changed `configureCache` to two methods: `getOrCreateCache`, which works the same, and `createCache`, for create only.
23. <https://github.com/jsr107/jsr107spec/issues/187> Added dynamic methods for registering and deregistering `CacheEntryListeners` on `Cache`.
24. <https://github.com/jsr107/jsr107spec/issues/210> Add a very simple `Caching.getCache(String cacheName, Long.class, String.class);`
25. <https://github.com/jsr107/jsr107spec/issues/194> Remove `CacheManager.isSupported` as it serves no purpose above having it on `CachingProvider`.
26. <https://github.com/jsr107/jsr107spec/issues/207> Remove `Cache.getOrCreate`. Separate methods are provided for creation and lookup.
27. <https://github.com/jsr107/jsr107spec/issues/212> Typo.
28. <https://github.com/jsr107/jsr107spec/issues/211> Minor clarifications to Section 2.2 and Section 15.
29. <https://github.com/jsr107/jsr107spec/issues/197> Clarify cache topologies.
30. <https://github.com/jsr107/jsr107spec/issues/200> Simplify and update simple example.

31. <https://github.com/jsr107/jsr107spec/issues/202> Clarified the scope and appropriate use of Cache Names.
32. <https://github.com/jsr107/jsr107spec/issues/199> Documented the purpose of CacheManager URIs.
33. <https://github.com/jsr107/jsr107spec/issues/213> Clarify that it's a non-objective for the Caching API to keep a Cache in sync with an external resource.
34. <https://github.com/jsr107/jsr107spec/issues/205> Clarified when Cache Configuration validation occurs.
35. <https://github.com/jsr107/jsr107spec/issues/186> Clarified EntryProcessor atomicity semantics.
36. <https://github.com/jsr107/jsr107spec/issues/206> Introduced Portability Recommendations section.
37. <https://github.com/jsr107/jsr107spec/issues/201> Specification Corrections
38. <https://github.com/jsr107/jsr107spec/issues/203> JavaDoc Corrections
39. <https://github.com/jsr107/jsr107spec/issues/222> Clarified CacheWriter deleteAll and writeAll method requirements for a mutable collection.