

ArjunaCore 4.11.0

Administration Guide

AC-AG-06/16/08



Legal Notices

The information contained in this documentation is subject to change without notice.

Arjuna Technologies Limited makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Arjuna Technologies Limited shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Java™ and J2EE is a U.S. trademark of Sun Microsystems, Inc. Microsoft® and Windows NT® are registered trademarks of Microsoft Corporation. Oracle® is a registered U.S. trademark and Oracle9™, Oracle9 Server™ Oracle9 Enterprise Edition™ are trademarks of Oracle Corporation. Unix is used here as a generic term covering all versions of the UNIX® operating system. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

Copyright

JBoss, Home of Professional Open Source Copyright 2006, JBoss Inc., and individual contributors as indicated by the @authors tag. All rights reserved.

See the copyright.txt in the distribution for a full listing of individual contributors. This copyrighted material is made available to anyone wishing to use, modify, copy, or redistribute it subject to the terms and conditions of the GNU General Public License, v. 2.0. This program is distributed in the hope that it will be useful, but WITHOUT A WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

See the GNU General Public License for more details. You should have received a copy of the GNU General Public License, v. 2.0 along with this distribution; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, * MA 02110-1301, USA.

Software Version

ArjunaCore 4.11.0

Restricted Rights Legend

Use, duplication, or disclosure is subject to restrictions as set forth in contract subdivision (c)(1)(ii) of the Rights in Technical Data and Computer Software clause 52.227-FAR14.

© Copyright 2010 JBoss Inc.

Contents

Table of Contents

About This Guide.....	4	Introduction.....	6
What This Guide Contains.....	4	Starting and Stopping the transaction manager.	6
Audience.....	4	ObjectStore management.....	7
Prerequisites.....	4	ArjunaCore runtime information.....	7
Organization.....	4	Failure recovery administration.....	8
Documentation Conventions.....	4	The Recovery Manager.....	8
Additional Documentation.....	5	Configuring the Recovery Manager.....	8
Contacting Us.....	5	Output.....	9
		Periodic Recovery.....	9
		Expired entry removal.....	10
		Errors and Exceptions.....	11
Administration of ArjunaCore.....	6		

About This Guide

What This Guide Contains

The Administration Guide contains information on how to use ArjunaCore 4.11.0.

Audience

This guide is most relevant to engineers who are responsible for administering ArjunaCore 4.11.0 installations.

Prerequisites

You should have installed ArjunaCore 4.11.0.

Organization

This guide contains the following chapters:

- **Chapter 1, Administration of ArjunaCore:** contains information on about how to administer ArjunaCore.

Documentation Conventions

The following conventions are used in this guide:

Convention	Description
<i>Italic</i>	In paragraph text, italic identifies the titles of documents that are being referenced. When used in conjunction with the Code text described below, italics identify a variable that should be replaced by the user with an actual value.
Bold	Emphasizes items of particular importance.
Code	Text that represents programming code.
Function Function	A path to a function or dialog box within an interface. For example, "Select File Open." indicates that you should select the Open function from the File menu.
() and	Parentheses enclose optional items in command syntax. The vertical bar separates syntax items in a list of choices. For example, any of the following three items can be entered in this syntax: persistPolicy (Never OnTimer OnUpdate NoMoreOftenThan)
Note:	A note highlights important supplemental information.
Caution:	A caution highlights procedures or information that is necessary to avoid damage to equipment, damage to software, loss of data, or invalid test results.

Table 1 Formatting Conventions

Additional Documentation

In addition to this guide, the following guides are available in the ArjunaCore 4.11.0 documentation set:

- **ArjunaCore 4.11.0 Release Notes:** Provides late-breaking information about ArjunaCore 4.11.0.
- **ArjunaCore 4.11.0 Installation Guide:** This guide provides instructions for installing ArjunaCore 4.11.0.
- **ArjunaCore 4.11.0 Users Guide:** Provides guidance for writing applications.

Contacting Us

Questions or comments about ArjunaCore 4.11.0 should be directed to our support team.

Administration of ArjunaCore

Introduction

Apart from ensuring that the run-time system is executing normally, there is little continuous administration needed for the *ArjunaCore* software. There are a few points however, that should be made:

- The present implementation of the *ArjunaCore* system provides no security or protection for data. The objects stored in the *ArjunaCore* object store are (typically) owned by the user who ran the application that created them. The Object Store and Object Manager facilities make no attempt to enforce even the limited form of protection that Unix/Windows provides. There is no checking of user or group IDs on access to objects for either reading or writing.
- Persistent objects created in the Object Store *never* go away unless the `StateManager.destroy` method is invoked on the object or some application program explicitly deletes them. This means that the Object Store gradually accumulates garbage (especially during application development and testing phases). At present we have no automated garbage collection facility. Further, we have not addressed the problem of dangling references. That is, a persistent object, A, may have stored a `Uid` for another persistent object, B, in its passive representation on disk. There is nothing to prevent an application from deleting B even though A still contains a reference to it. When A is next activated and attempts to access B, a run-time error will occur.
- There is presently no support for version control of objects or database reconfiguration in the event of class structure changes. This is a complex research area that we have not addressed. At present, if you change the definition of a class of persistent objects, you are entirely responsible for ensuring that existing instances of the object in the Object Store are converted to the new representation. The *ArjunaCore* software can neither detect nor correct references to old object state by new operation versions or vice versa.
- Object store management is critically important to the transaction service.

Starting and Stopping the transaction manager

By default the transaction manager starts up in an active state such that new transactions can be created immediately. If you wish to have more control over this it is possible to set the `CoordinatorEnvironmentBean.startDisabled` configuration option to YES and in

which case no transactions can be created until the transaction manager is enabled (via a call to `TxControl.enable`).

It is possible to stop the creation of new transactions at any time by calling `TxControl.disable`. Transactions that are currently executing will not be affected. By default recovery will be allowed to continue and the transaction system will still be available to manage recovery requests from other instances in a distributed environment. (See the Failure Recovery Guide for further details). However, if you wish to disable recovery as well as remove any resources it maintains, then you can pass `true` to `TxControl.disable`; the default is to use `false`.

If you wish to shut the system down completely then it may also be necessary to terminate the background transaction reaper (see the Programmers Guide for information about what the reaper does.) In order to do this you may want to first prevent the creation of new transactions (if you are not creating transactions with timeouts then this step is not necessary) using `TxControl.disable`. Then you should call `TransactionReaper.terminate`. This method takes a boolean parameter: if `true` then the method will wait for the normal timeout periods associated with any transactions to expire before terminating the transactions; if `false` then transactions will be forced to terminate (rollback or have their outcome set such that they can only ever rollback) immediately.

Note: if you intent to restart the recovery manager later after having terminated it then you **MUST** use the `terminate` method with asynchronous behavior set to `false`.

ObjectStore management

Within the transaction service installation, the object store is updated regularly whenever transactions are created, or when *Transactional Objects for Java* is used. In a failure free environment, the only object states which should reside within the object store are those representing objects created with the *Transactional Objects for Java* API. However, if failures occur, transaction logs may remain in the object store until crash recovery facilities have resolved the transactions they represent. As such it is very important that the contents of the object store are not deleted without due care and attention, as this will make it impossible to resolve in doubt transactions. In addition, if multiple users share the same object store it is important that they realise this and do not simply delete the contents of the object store assuming it is an exclusive resource.

ArjunaCore runtime information

Compile time configuration information is available via the `com.arjuna.common.util.ConfigurationInfo` class. Runtime configuration is embodied in the various `<name>EnvironmentBean` classes, see the configuration section of the user guide. These beans have corresponding MBean interfaces and may be linked to JMX for remote inspection of the configuration if desired.

Failure recovery administration

The failure recovery subsystem of *ArjunaCore* will ensure that results of a transaction are applied consistently to all resources affected by the transaction, even if any of the application processes or the machine hosting them crash or lose network connectivity. In the case of machine (system) crash or network failure, the recovery will not take place until the system or network are restored, but the original application does not need to be restarted – recovery responsibility is delegated to the Recovery Manager process (see below). Recovery after failure requires that information about the transaction and the resources involved survives the failure and is accessible afterward: this information is held in the ActionStore, which is part of the ObjectStore.

| Caution: If the ObjectStore is destroyed or modified, recovery may not be possible.

Until the recovery procedures are complete, resources affected by a transaction that was in progress at the time of the failure may be inaccessible. For database resources, this may be reported as tables or rows held by “in-doubt transactions”. For TransactionalObjects for Java resources, an attempt to activate the Transactional Object (as when trying to get a lock) will fail.

The Recovery Manager

The failure recovery subsystem of *ArjunaCore* requires that the stand-alone Recovery Manager process be running for each ObjectStore (typically one for each node on the network that is running *ArjunaCore* applications). The RecoveryManager file is located in the *arjunacore* jar file within the package `com.arjuna.ats.arjuna.recovery.RecoveryManager`. To start the Recovery Manager issue the following command:

```
java com.arjuna.ats.arjuna.recovery.RecoveryManager
```

If the `-test` flag is used with the Recovery Manager then it will display a “Ready” message when initialised, i.e.,

```
java com.arjuna.ats.arjuna.recovery.RecoveryManager -test
```

Configuring the Recovery Manager

The RecoveryManager reads the properties defined in the `arjuna.properties` file and then also reads the property file `RecoveryManager.properties`, from the same directory as it found the `arjuna` properties file. An entry for a property in the RecoveryManager properties file will override an entry for the same property in the main TransactionService properties file. Most of the entries are specific to the Recovery Manager.

A default version of `RecoveryManager.properties` is supplied with the distribution – this can be used without modification, except possibly the debug tracing fields (see below, Output). The rest of this section discusses the issues relevant in setting the properties to other values (in the order of their appearance in the default version of the file)

Output

It is likely that installations will want to have some form of output from the RecoveryManager, to provide a record of what recovery activity has taken place. RecoveryManager uses the logging tracing mechanism provided by the Arjuna Common Logging Framework (CLF), which provides a high level interface that hides differences that exist between existing logging APIs such as Jakarta log4j or JDK 1.4 logging API. CLF indirections all logging via the Apache Commons Logging framework and configuration is assumed to occur through that framework.

With the CLF applications make logging calls on logger objects. Loggers may use logging Levels to decide if they are interested in a particular log message. Each log message has an associated log Level, that gives the importance and urgency of a log message. The set of possible Log Levels are `DEBUG`, `INFO`, `WARN`, `ERROR` and `FATAL`. Defined Levels are ordered according to their integer values as follows: `DEBUG < INFO < WARN < ERROR < FATAL`.

Messages describing the start and the periodical behavior made by the RecoveryManager are output using the `INFO` level. If other debug tracing is wanted, the finer debugging level should be set appropriately.

Setting the normal recovery messages to the `INFO` level allows the RecoveryManager to produce a moderate level of reporting. If nothing is going on, it just reports the entry into each module for each periodic pass. To disable `INFO` messages produced by the RecoveryManager, the logging level could be set to the higher level: `ERROR`. Setting the level to `ERROR` means that the RecoveryManager will only produce error, warning or fatal messages.

Periodic Recovery

The RecoveryManager scans the ObjectStore and other locations of information, looking for transactions and resources that require, or may require recovery. The scans and recovery processing are performed by recovery modules, (instances of classes that implement the `com.arjuna.ats.arjuna.recovery.RecoveryModule` interface), each with responsibility for a particular category of transaction or resource. The set of recovery modules used are dynamically loaded, using properties found in the RecoveryManager property file.

The interface has two methods: `periodicWorkFirstPass` and `periodicWorkSecondPass`. At an interval (defined by property `com.arjuna.ats.arjuna.recovery.periodicRecoveryPeriod`), the RecoveryManager will call the first pass method on each property, then wait for a brief period (defined by property `com.arjuna.ats.arjuna.recovery.recoveryBackoffPeriod`), then call the second pass of each module. Typically, in the first pass, the module scans (e.g. the relevant part of the ObjectStore) to find transactions or resources that are in-doubt (i.e. are part way through the commitment process). On the second pass, if any of the same items are still in-doubt, it is possible the original application process has crashed and the item is a candidate for recovery.

An attempt, by the RecoveryManager, to recover a transaction that is still progressing in the original process(es) is likely to break the consistency. Accordingly, the recovery modules use

a mechanism (implemented in the `com.arjuna.ats.arjuna.recovery.TransactionStatusManager` package) to check to see if the original process is still alive, and if the transaction is still in progress. The `RecoveryManager` only proceeds with recovery if the original process has gone, or, if still alive, the transaction is completed. (If a server process or machine crashes, but the transaction-initiating process survives, the transaction will complete, usually generating a warning. Recovery of such a transaction is the `RecoveryManager`'s responsibility).

It is clearly important to set the interval periods appropriately. The total iteration time will be the sum of the `periodicRecoveryPeriod`, `recoveryBackoffPeriod` and the length of time it takes to scan the stores and to attempt recovery of any in-doubt transactions found, for all the recovery modules. The recovery attempt time may include connection timeouts while trying to communicate with processes or machines that have crashed or are inaccessible (which is why there are mechanisms in the recovery system to avoid trying to recover the same transaction for ever). The total iteration time will affect how long a resource will remain inaccessible after a failure – `periodicRecoveryPeriod` should be set accordingly (default is 120 seconds). The `recoveryBackoffPeriod` can be comparatively short (default is 10 seconds) – its purpose is mainly to reduce the number of transactions that are candidates for recovery and which thus require a “call to the original process to see if they are still in progress

Note: In previous versions of ArjunaCore there was no contact mechanism, and the backoff period had to be long enough to avoid catching transactions in flight at all. From 3.0, there is no such risk.

Two recovery modules (implementations of the `com.arjuna.ats.arjuna.recovery.RecoveryModule` interface) are supplied with *ArjunaCore*, supporting various aspects of transaction recovery including JDBC recovery. It is possible for advanced users to create their own recovery modules and register them with the `RecoveryManager`. The recovery modules are registered with the `RecoveryManager` using `RecoveryEnvironmentBean.recoveryExtensions`. These will be invoked on each pass of the periodic recovery in the sort-order of the property names – it is thus possible to predict the ordering (but note that a failure in an application process might occur while a periodic recovery pass is in progress). The default `RecoveryExtension` settings are:

```
<entry key="RecoveryEnvironmentBean.recoveryExtensions">
    com.arjuna.ats.internal.arjuna.recovery.AtomicActionRecoveryModule
    com.arjuna.ats.internal.txoj.recovery.TORecoveryModule
</entry>
```

Expired entry removal

The operation of the recovery subsystem will cause some entries to be made in the `ObjectStore` that will not be removed in normal progress. The `RecoveryManager` has a facility for scanning for these and removing items that are very old. Scans and removals are performed by implementations of the `com.arjuna.ats.arjuna.recovery.ExpiryScanner` interface. Implementations of this interface are loaded by giving the class names as the value of a property `RecoveryEnvironmentBean.expiryScanners`. The `RecoveryManager` calls the `scan()` method on each loaded `ExpiryScanner` implementation at an interval determined by the property `RecoveryEnvironmentBean.expiryScanInterval`. This value is given in *hours* – default is 12. An `expiryScanInterval` value of zero will suppress any expiry scanning. If the value as supplied is positive, the first scan is performed when `RecoveryManager` starts; if the

value is negative, the first scan is delayed until after the first interval (using the absolute value)

The kinds of item that are scanned for expiry are:

TransactionStatusManager items : one of these is created by every application process that uses ArjunaCore – they contain the information that allows the RecoveryManager to determine if the process that initiated the transaction is still alive, and what the transaction status is. The expiry time for these is set by the property `com.arjuna.ats.arjuna.recovery.transactionStatusManagerExpiryTime` (in hours – default is 12, zero means never expire). The expiry time should be greater than the lifetime of any single *ArjunaCore*-using process.

The Expiry Scanner properties for these are:

```
<entry key="RecoveryEnvironmentBean.expiryScanners">
  com.arjuna.ats.internal.arjuna.recovery.ExpiredTransactionStatusManager
  Scanner
</entry>
```

Errors and Exceptions

In this section we shall cover the types of errors and exceptions which may be thrown or reported during a transactional application and give probable indications of their causes.

- `NO_MEMORY`: the application has run out of memory (thrown an `OutOfMemoryError`) and *ArjunaCore* has attempted to do some cleanup (by running the garbage collector) before re-throwing the exception. This is probably a transient problem and retrying the invocation should succeed.
- `com.arjuna.ats.arjuna.exceptions.FatalError`: an error has occurred which means that the transaction system must shut down. Prior to this error being thrown the transaction service will have ensured that all running transactions have rolled back. If caught, the application should tidy up and exit. If further work is attempted, application consistency may be violated.
- `com.arjuna.ats.arjuna.exceptions.LicenceError`: an attempt has been made to use the transaction service in a manner inconsistent with the current licence. The transaction service will not allow further forward progress for existing or new transactions.
- `com.arjuna.ats.arjuna.exceptions.ObjectStoreError`: an error occurred while the transaction service attempted to use the object store. Further forward progress is not possible.
- Object store warnings about access problems on states may occur during the normal execution of crash recovery. This is the result of multiple concurrent attempts to perform recovery on the same transaction. It can be safely ignored.