

JBoss Transactions API 4.2

Programmers Guide

JBTA-PG-3/22/06



Legal Notices

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

JBoss Inc. 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. JBoss Inc. 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.

Arjuna is a trademark of Hewlett-Packard Company and used here under licence.

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

JBoss Transactions API 4.2

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 2006 JBoss Inc.

Contents

About This Guide	5	JDBC and transactions	23
What This Guide Contains.....	5	Using the transactional JDBC driver.....	23
Audience	5	Managing transactions	23
Prerequisites.....	5	Restrictions.....	23
Organization	5	Transactional drivers.....	23
Documentation Conventions.....	6	Loading drivers	24
Additional Documentation.....	6	Connections.....	24
Contacting Us	7	Making the connection.....	25
An introduction to the JTA	9	JDBC 2.0.....	25
The Java Transaction API.....	9	XADataSources.....	25
Transactions.....	11	Java Naming and Directory Interface (JNDI).....	25
The API.....	11	Dynamic class instantiation.....	26
UserTransaction.....	11	Using the connection.....	27
TransactionManager	11	Connection pooling	28
Suspending and resuming a transaction.....	12	Reusing connections	28
The Transaction interface	13	Terminating the transaction.....	28
Resource enlistment.....	13	AutoCommit.....	28
Transaction synchronization	14	Setting isolation levels	29
Transaction equality.....	14	Examples.....	30
The Resource Manager	15	JDBC example	30
The XAResource interface.....	15	Failure recovery example.....	32
Extended XAResource control	16	Configuring JBossJTA	37
Opening a Resource Manager.....	16	Configuration options	37
Closing a Resource Manager	17	Using JBossJTA in application servers.....	38
Threads of control.....	17	JBoss Application Server.....	38
Transaction association.....	17	Configuration	38
Externally controlled connections	18	The services.....	39
Resource sharing.....	18	Ensuring Transactional Context is Propagated to the Server.....	39
Local and global transactions	19	Index	40
Dynamic Registration	19		
Transaction recovery	20		
Failure recovery	20		
Recovering non-JBossJTA XAConnections.....	21		

About This Guide

What This Guide Contains

This document describes how to use the Arjuna Technologies JTA and JDBC implementations, referred to within this document as JBoss Transactions API 4.2 (JBossJTA).

Audience

This guide is most relevant to engineers who are responsible for administering JBoss Transactions API 4.2 installations.

Prerequisites

Knowledge of JTA and JDBC.

Organization

This guide contains the following chapters:

- **Chapter 1, An introduction to the JTA:** An introduction to JTA by describing the concept of the high level interface to manage transactions.
- **Chapter 2, Transactions:** contains an overview of the interfaces defined by JTA demarcate and manage transactions from the user view and from the application server view.
- **Chapter 3, The Resource Manager:** gives an overview of the way to manage Resource Managers compliant to the X/Open XA protocol.
- **Chapter 4, Transaction Recovery:** This chapter, gives an overview of the behavior applied to recovery XA resource managers, and describes information needed by JBossJTA to manage the recovery.
- **Chapter 5, JDBC and transaction:** This chapter, describes how JBossJTA supports transactional applications that access databases using JDBC.
- **Chapter 6, Example:** illustrates how to build transactional applications with JBossJTA- accessing a database with JDBC and managing recovery.

- **Chapter 7, Configuring JBossJTA:** describes how to configure JBossJTA features.
- **Chapter 8, Using JBossJTA in the application servers:** describes how to use and configure JBossJTA when embedded in an application server, particularly JBoss.

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: <code>persistPolicy (Never OnTimer OnUpdate NoMoreOftenThan)</code>
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 JBoss Transactions API 4.2 documentation set:

- JBoss Transactions API 4.2 *Release Notes*: Provides late-breaking information about JBoss Transactions API 4.2.
- JBoss Transactions API 4.2 *Installation Guide*: This guide provides instructions for installing JBoss Transactions API 4.2.

- JBoss Transactions API 4.2 *Administrators Guide*: Provides guidance for administering the system.

Contacting Us

Questions or comments about JBoss Transactions API 4.2 should be directed to our support team.

An introduction to the JTA

The Java Transaction API

The interfaces specified by the many transaction standards are typically too low-level for most application programmers. Therefore, Sun Microsystems has specified higher-level interfaces to assist in the development of distributed transactional applications. Note, these interfaces are still low-level, and require, for example, the programmer to be concerned with state management and concurrency for transactional application. In addition, they are geared more for applications which require XA resource integration capabilities, rather than the more general resources which the other APIs allow.

With reference to [JTA99], distributed transaction services typically involve a number of participants:

- *application server*: which provides the infrastructure required to support the application run-time environment which includes transaction state management, e.g., an EJB server.
- *transaction manager*: provides the services and management functions required to support transaction demarcation, transactional resource management, synchronisation and transaction context propagation.
- *resource manager*: (through a resource adapter¹) provides the application with access to resources. The resource manager participates in distributed transactions by implementing a transaction resource interface used by the transaction manager to communicate transaction association, transaction completion and recovery.
- *a communication resource manager (CRM)*: supports transaction context propagation and access to the transaction service for incoming and outgoing requests.

From the transaction manager's perspective, the actual implementation of the transaction services does not need to be exposed; only high-level interfaces need to be defined to allow transaction demarcation, resource enlistment, synchronization and recovery process to be driven from the users of the transaction services. The JTA is a high-level application interface

¹ A Resource Adapter is used by an application server or client to connect to a Resource Manager. JDBC drivers which are used to connect to relational databases are examples of Resource Adapters.

that allows a transactional application to demarcate transaction boundaries, and contains also contains a mapping of the X/Open XA protocol.

Note: the JTA support provided by *JBossJTA* is compliant with the 1.0.1 specification.

Chapter 2

Transactions

The API

The Java Transaction API consists of three elements: a high-level application transaction demarcation interface, a high-level transaction manager interface intended for application server, and a standard Java mapping of the X/Open XA protocol intended for transactional resource manager. All of the JTA classes and interfaces occur within the `javax.transaction` package, and the corresponding *JBossJTA* implementations within the `com.arjuna.ats.jta` package.

UserTransaction

The `UserTransaction` interface provides applications with the ability to control transaction boundaries. It has methods for beginning, committing, and rolling back top-level transactions: nested transactions are not supported, and `begin` throws the `NotSupportedException` when the calling thread is already associated with a transaction. `UserTransaction` automatically associates newly created transactions with the invoking thread.

Note: In *JBossJTA*, `UserTransactions` can be obtained from the static `com.arjuna.ats.jta.UserTransaction.userTransaction()` method.

In order to select the local JTA implementation it is necessary to perform the following steps:

1. make sure the property `com.arjuna.ats.jta.jtaTMIImplementation` is set to `com.arjuna.ats.internal.jta.transaction.arjunacore.TransactionManagerImple`.
2. make sure the property `com.arjuna.ats.jta.jtaUTImplementation` is set to `com.arjuna.ats.internal.jta.transaction.arjunacore.UserTransactionImple`.

TransactionManager

The `TransactionManager` interface allows the application server to control transaction boundaries on behalf of the application being managed.

Note: In *JBossJTA*, transaction manager implementations can be obtained from the static `com.arjuna.ats.jta.TransactionManager.transactionManager` method.

The Transaction Manager maintains the transaction context association with threads as part of its internal data structure. A thread's transaction context is either *null* or it refers to a specific

global transaction. Multiple threads may be associated with the same global transaction. As noted above, nested transactions are not supported.

Each transaction context is encapsulated by a `Transaction` object, which can be used to perform operations which are specific to the target transaction, regardless of the calling thread's transaction context.

The `begin` method of `TransactionManager` starts a new top-level transaction and associates the transaction context with the calling thread. If the calling thread is already associated with a transaction then it throws the `NotSupportedException`.

The `getTransaction` method returns the `Transaction` object that represents the transaction context currently associated with the calling thread. This object can be used to perform various operations on the target transaction, described later.

The `commit` method is used to complete the transaction currently associated with the calling thread. After it returns, the calling thread is associated with no transaction. If `commit` is called when the thread is not associated with any transaction context, the TM throws an exception. In some implementation, the commit operation is restricted to the transaction originator only. If the calling thread is not allowed to commit the transaction, the TM throws an exception. *JBossJTA* does not currently impose any restriction on the ability of threads to terminate transactions.

The `rollback` method is used to rollback the transaction associated with the current thread. After the `rollback` method completes, the thread is associated with no transaction.

Caution: Each Xid that JBossTS creates must have a unique node identifier encoded within it and JBossTS will only recover transactions and states that match a specified node identifier. The node identifier to use should be provided to JBossTS via the `com.arjuna.ats.arjuna.xa.nodeIdentifier` property. You must make sure this value is unique across your JBossTS instances. If you do not provide a value, then JBossTS will fabricate one and report the value via the logging infrastructure.

Suspending and resuming a transaction

The JTA supports the concept of a thread temporarily suspending and resuming transactions to enable it to perform non-transactional work. The `suspend` method is called to temporarily suspend the current transaction that is associated with the calling thread, i.e., so that the thread is no longer operating within its scope. If the thread is not associated with any transaction, a `null` object reference is returned; otherwise, a valid `Transaction` object is returned. The `Transaction` object can later be passed to the `resume` method to reinstate the transaction context.

The `resume` method associates the specified transaction context with the calling thread. If the transaction specified is a valid transaction, the transaction context is associated with the calling thread; otherwise, the thread is associated with no transaction.

Note: if `resume` is invoked when the calling thread is already associated with another transaction, the Transaction Manager throws the `IllegalStateException` exception.

```
Transaction tobj = TransactionManager.suspend();
..
TransactionManager.resume(tobj);
```

Note: some transaction manager implementations allow a suspended transaction to be resumed by a different thread. This feature is not required by JTA, but *JBossJTA* does support this.

When a transaction is suspended the application server must ensure that the resources in use by the application are no longer registered with the suspended transaction. When a resource is de-listed this triggers the Transaction Manager to inform the resource manager to disassociate the transaction from the specified resource object. When the application's transaction context is resumed, the application server must ensure that the resources in use by the application are again enlisted with the transaction. Enlisting a resource as a result of resuming a transaction triggers the Transaction Manager to inform the resource manager to re-associate the resource object with the resumed transaction.

The Transaction interface

The `Transaction` interface allows operations to be performed on the transaction associated with the target object. Every top-level transaction is associated with one `Transaction` object when the transaction is created. The `Transaction` object can be used to:

- enlist the transactional resources in use by the application.
- register for transaction synchronization call backs.
- commit or rollback the transaction.
- obtain the status of the transaction.

The `commit` and `rollback` methods allow the target object to be committed or rolled back. The calling thread is not required to have the same transaction associated with the thread. If the calling thread is not allowed to commit the transaction, the transaction manager throws an exception. At present *JBossJTA* does not impose restrictions on threads terminating transactions.

Resource enlistment

Transactional resources such as database connections are typically managed by the application server in conjunction with some resource adapter and optionally with connection pooling optimization. In order for an external transaction manager to co-ordinate transactional work performed by the resource managers, the application server must enlist and de-list the resources used in the transaction. These resources (participants) are enlisted with the transaction so that they can be informed when the transaction terminates, e.g., are driven through the two-phase commit protocol.

As stated previously, the JTA is much more closely integrated with the XA concept of resources than the arbitrary objects. For each resource in-use by the application, the application server invokes the `enlistResource` method with an `XAResource` object which identifies the resource in use. See Chapter 4 for details on how the implementation of the `XAResource` can affect recovery in the event of a failure.

The enlistment request results in the transaction manager informing the resource manager to start associating the transaction with the work performed through the corresponding resource. The transaction manager is responsible for passing the appropriate flag in its `XAResource.start` method call to the resource manager.

The `delistResource` method is used to disassociate the specified resource from the transaction context in the target object. The application server invokes the method with the two parameters: the `XAResource` object that represents the resource, and a flag to indicate whether the operation is due to the transaction being suspended (`TMSUSPEND`), a portion of the work has failed (`TMFAIL`), or a normal resource release by the application (`TMSUCCESS`).

The de-list request results in the transaction manager informing the resource manager to end the association of the transaction with the target `XAResource`. The flag value allows the application server to indicate whether it intends to come back to the same resource whereby the resource states must be kept intact. The transaction manager passes the appropriate flag value in its `XAResource.end` method call to the underlying resource manager.

Transaction synchronization

Transaction synchronization allows the application server to be notified before and after the transaction completes. For each transaction started, the application server may optionally register a `Synchronization` call back object to be invoked by the transaction manager:

- The `beforeCompletion` method is called prior to the start of the two-phase transaction complete process. This call is executed in the same transaction context of the caller who initiates the `TransactionManager.commit` or the call is executed with no transaction context if `Transaction.commit` is used.
- The `afterCompletion` method is called after the transaction has completed. The status of the transaction is supplied in the parameter. This method is executed without a transaction context.

Transaction equality

The transaction manager implements the `Transaction` object's `equals` method to allow comparison between the target object and another `Transaction` object. The `equals` method should return `true` if the target object and the parameter object both refer to the same global transaction.

```
Transaction txObj = TransactionManager.getTransaction();
Transaction someOtherTxObj = ..
..

boolean isSame = txObj.equals(someOtherTxObj);
```

The Resource Manager

The XAResource interface

Whereas some transaction specifications and systems define a generic resource which can be used to register arbitrary resources with a transaction, the JTA is much more XA specific. The `javax.transaction.xa.XAResource` interface is a Java mapping of the XA interface. The `XAResource` interface defines the contract between a Resource Manager and a Transaction Manager in a distributed transaction processing environment. A resource adapter for a resource manager implements the `XAResource` interface to support association of a top-level transaction to a resource such as a relational database.

The `XAResource` interface can be supported by any transactional resource adapter that is intended to be used in an environment where transactions are controlled by an external transaction manager, e.g., a database management system. An application may access data through multiple database connections. Each database connection is associated with an `XAResource` object that serves as a proxy object to the underlying resource manager instance. The transaction manager obtains an `XAResource` for each resource manager participating in a top-level transaction. It uses the `start` method to associate the transaction with the resource, and it uses the `end` method to disassociate the transaction from the resource.

The resource manager is responsible for associating the transaction with all work performed on its data between the `start` and `end` invocations. At transaction commit time, these transactional resource managers are informed by the transaction manager to prepare, commit, or rollback the transaction according to the two-phase commit protocol.

In order to be better integrated with Java, the `XAResource` differs from the standard XA interface in the following ways:

- The resource manager initialization is done implicitly by the resource adapter when the resource (connection) is acquired. There is no `xa_open` equivalent.
- `Rmid` is not passed as an argument. Each `Rmid` is represented by a separate `XAResource` object.
- Asynchronous operations are not supported because Java supports multi-threaded processing and most databases do not support asynchronous operations.
- Error return values that are caused by the transaction manager's improper handling of the `XAResource` object are mapped to Java exceptions via the `XAException` class.
- The DTP concept of "Thread of Control" maps to all Java threads that are given access to the `XAResource` and `Connection` objects. For example, it is legal for

two different threads to perform the start and end operations on the same `XAResource` object.

Extended XAResource control

By default, whenever an `XAResource` object is registered with a JTA compliant transaction service, you have no control over the order in which it will be invoked during the two-phase commit protocol, with respect to other `XAResource` objects. In *JBossTS*, however, there is support for controlling the order via the two interfaces `com.arjuna.ats.jta.resources.StartXAResource` and `com.arjuna.ats.jta.resources.EndXAResource`. By inheriting your `XAResource` instance from either of these interfaces, you control whether an instance of your class will be invoked first or last, respectively.

Note: Only one instance of each interface type may be registered with a specific transaction.

In the TxCore manual we discussed the Last Resource Commit optimization (LRCO), whereby a single resource that is only one-phase aware (does not support prepare), can be enlisted with a transaction that is manipulating two-phase aware participants. This optimization is also supported within the JTA aspect of *JBossTS*.

In order to use the LRCO, your `XAResource` implementation must extend the `com.arjuna.ats.jta.resources.LastResourceCommitOptimisation` marker interface (it provides no methods). When enlisting the resource via `Transaction.enlistResource`, *JBossTS* will ensure that only a single instance of this type of participant is used within each transaction. Your resource will be driven last in the commit protocol: no invocation of prepare will occur.

Note: An attempt to enlist more than one instance of a `LastResourceCommitOptimisation` class will fail and `false` will be returned from `Transaction.enlistResource`.

In order to utilize the LRCO in a distributed environment, it is necessary to disable interposition support. It is still possible to use implicit context propagation.

Opening a Resource Manager

The X/Open XA interface requires that the transaction manager initialize a resource manager (`xa_open`) prior to any other `xa_` calls. JTA requires initialization of a resource manager to be embedded within the resource adapter that represents the resource manager. The transaction manager does not need to know how to initialize a resource manager; it is only responsible for informing the resource manager about when to start and end work associated with a transaction and when to complete the transaction. The resource adapter is responsible for opening (initializing) the resource manager when the connection to the resource manager is established.

Closing a Resource Manager

A resource manager is closed by the resource adapter as a result of destroying the transactional resource. A transaction resource at the resource adapter level is comprised of two separate objects:

- An `XAResource` object that allows the transaction manager to start and end the transaction association with the resource in use and to coordinate transaction completion process.
- A connection object that allows the application to perform operations on the underlying resource (for example, JDBC operations on an RDBMS).

Once opened, the resource manager is kept open until the resource is released (closed) explicitly. When the application invokes the connection's `close` method, the resource adapter invalidates the connection object reference that was held by the application and notifies the application server about the close. The transaction manager should invoke the `XAResource.end` method to disassociate the transaction from that connection.

The `close` notification allows the application server to perform any necessary cleanup work and to mark the physical XA connection as free for reuse, if connection pooling is in place.

Threads of control

The X/Open XA interface specifies that the transaction association related `xa` calls must be invoked from the same thread context. This thread-of-control requirement is not applicable to the object-oriented component-based application run-time environment, in which application threads are dispatched dynamically at method invocation time. Different threads may be using the same connection resource to access the resource manager if the connection spans multiple method invocation. Depending on the implementation of the application server, different threads may be involved with the same `XAResource` object. The resource context and the transaction context may be operated independent of thread context. This means that it is possible for different threads to be invoking the `start` and `end` methods.

If the application server allows multiple threads to use a single `XAResource` object and the associated connection to the resource manager, it is the responsibility of the application server to ensure that there is only one transaction context associated with the resource at any point of time. Thus the `XAResource` interface requires that the resource managers be able to support the two-phase commit protocol from any thread context.

Transaction association

Transactions are associated with a transactional resource via the `start` method, and disassociated from the resource via the `end` method. The resource adapter is responsible for internally maintaining an association between the resource connection object and the `XAResource` object. At any given time, a connection is associated with a single transaction, or it is not associated with any transaction at all. Because JTA does not support nested

transactions it is an error for the `start` method to be invoked on a connection that is currently associated with a different transaction.

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

Externally controlled connections

For transactional application whose transaction states are managed by an application server, its resources must also be managed by the application server so that transaction association is performed properly. If an application is associated with a transaction, it is an error for the application to perform transactional work through the connection without having the connection's resource object already associated with the global transaction. The application server must ensure that the `XAResource` object in use is associated with the transaction. This is done by invoking the `Transaction.enlistResource` method.

If a server side transactional application retains its database connection across multiple client requests, the application server must ensure that before dispatching a client request to the application thread, the resource is enlisted with the application's current transaction context. This implies that the application server manages the connection resource usage status across multiple method invocations.

Resource sharing

When the same transactional resource is used to interleave multiple transactions, it is the responsibility of the application server to ensure that only one transaction is enlisted with the resource at any given time. To initiate the transaction commit process, the transaction manager is allowed to use any of the resource objects connected to the same resource manager instance. The resource object used for the two-phase commit protocol does not need to have been involved with the transaction being completed.

The resource adapter must be able to handle multiple threads invoking the `XAResource` methods concurrently for transaction commit processing. For example, with reference to the code below, suppose we have a transactional resource `r1`. Global transaction `xid1` was started and ended with `r1`. Then a different global transaction `xid2` is associated with `r1`. In the meanwhile, the transaction manager may start the two phase commit process for `xid1` using `r1` or any other transactional resource connected to the same resource manager. The resource adapter needs to allow the commit process to be executed while the resource is currently associated with a different global transaction.

```
XAResource xares = r1.getXAResource();  
  
xares.start(xid1); // associate xid1 to the connection  
  
..  
xares.end(xid1); // disassociate xid1 to the connection  
..
```

```
xares.start(xid2); // associate xid2 to the connection
..
// While the connection is associated with xid2,
// the TM starts the commit process for xid1
status = xares.prepare(xid1);
..
xares.commit(xid1, false);
```

Local and global transactions

The resource adapter must support the usage of both local and global transactions within the same transactional connection. Local transactions are transactions that are started and coordinated by the resource manager internally. The `XAResource` interface is not used for local transactions. When using the same connection to perform both local and global transactions, the following rules apply:

- The local transaction must be committed (or rolled back) before starting a global transaction in the connection.
- The global transaction must be disassociated from the connection before any local transaction is started.

Dynamic Registration

Dynamic registration is not supported in `XAResource` because of the following reasons:

- In the Java component-based application server environment, connections to the resource manager are acquired dynamically when the application explicitly requests a connection. These resources are enlisted with the transaction manager on a needed basis.
- If a resource manager requires a way to dynamically register its work to the global transaction, the implementation can be done at the resource adapter level via a private interface between the resource adapter and the underlying resource manager.

Transaction recovery

Failure recovery

During recovery, the Transaction Manager needs to be able to communicate to all resource managers that are in use by the applications in the system. For each resource manager, the Transaction Manager uses the `XAResource.recover` method to retrieve the list of transactions that are currently in a prepared or heuristically completed state. Typically, the system administrator configures all transactional resource factories that are used by the applications deployed on the system. An example of such a resource factory is the JDBC `XADataSource` object, which is a factory for the JDBC `XAConnection` objects.

Because `XAResource` objects are not persistent across system failures, the Transaction Manager needs to have some way to acquire the `XAResource` objects that represent the resource managers which might have participated in the transactions prior to the system failure. For example, a Transaction Manager might, through the use of JNDI lookup mechanism, acquire a connection from each of the transactional resource factories, and then obtain the corresponding `XAResource` object for each connection. The Transaction Manager then invokes the `XAResource.recover` method to ask each resource manager to return the transactions that are currently in a prepared or heuristically completed state.

Note: When running XA recovery it is necessary to tell JBossTS which types of Xid it can recover. Each Xid that JBossTS creates has a unique node identifier encoded within it and JBossTS will only recover transactions and states that match a specified node identifier. The node identifier to use should be provided to JBossTS via a property that starts with the name `com.arjuna.ats.jta.xaRecoveryNode`; multiple values may be provided. A value of `*` will force JBossTS to recover (and possibly rollback) all transactions irrespective of their node identifier and should be used with caution.

If using the *JBossJTA* JDBC 2.0 driver, then *JBossJTA* will take care of all `XAResource` crash recovery automatically. Otherwise one of the following recovery mechanisms will be used:

- If the `XAResource` is serializable, then the serialized form will be saved during transaction commitment, and used during recovery. It is assumed that the recreated `XAResource` is valid and can be used to drive recovery on the associated database.
- The `com.arjuna.ats.jta.recovery.XAResourceRecovery`, `com.arjuna.ats.jta.recovery.XARecoveryResourceManager` and `com.arjuna.ats.jta.recovery.XARecoveryResource` interfaces are used. These are documented in the *JDBC chapters* on failure recovery.

Note: In JBossTS 3.3 the interface `XAConnectionRecovery` was deprecated in favor of `XAResourceRecovery`. Users are encouraged to move to his new interface.

Recovering non-JBossJTA XAConnections

When recovering from failures, *JBossJTA* requires the ability to reconnect to databases that were in use prior to the failures in order to resolve any outstanding transactions. Most connection information will be saved by the transaction service during its normal execution, and can be used during recovery to recreate the connection. However, it is possible that not all such information will have been saved prior to a failure (for example, a failure occurs before such information can be saved, but after the database connection is used). In order to recreate those connections it is necessary to provide implementations of the following *JBossJTA* interface `com.arjuna.ats.jta.recovery.XAResourceRecovery`, one for each database that may be used by an application.

Note: if using the transactional JDBC 2.0 driver provided with JBossJTA, then no additional work is necessary in order to ensure that recovery occurs.

To inform the recovery system about each of the `XAResourceRecovery` instances, it is necessary to specify their class names through property variables. Any property variable found in the properties file, or registered at runtime, which starts with the name `com.arjuna.ats.jta.recovery.XAResourceRecovery` will be assumed to represent one of these instances, and its value should be the class name. For example:

```
com.arjuna.ats.jta.recovery.XAResourceRecoveryOracle=com.foo.barRecovery
```

Additional information that will be passed to the instance when it is created may be specified after a semicolon:

```
com.arjuna.ats.jta.recovery.XAResourceRecoveryOracle=com.foo.barRecovery;my
Data=hello
```

Any errors will be reported during recovery.

```
public interface XAResourceRecovery
{
    public XAResource getXAResource () throws SQLException;

    public boolean initialise (String p);

    public boolean hasMoreResources ();
};
```

Each method should return the following information:

- *initialise*: once the instance has been created, any additional information which occurred on the property value (anything found after the first semi-colon) will be passed to the object. The object can then use this information in an implementation specific manner to initialise itself, for example.

- *hasMoreResources*: each `XAResourceRecovery` implementation may provide multiple `XAResource` instances. Before any call to `getXAResource` is made, `hasMoreResources` is called to determine whether there are any further connections to be obtained. If this returns *false*, `getXAResource` will not be called and the instance will not be used further.
- *getXAResource*: returns an instance of the `XAResource` object. How this is created (and how the parameters to its constructors are obtained) is up to the `XAResourceRecovery` implementation. The parameters to the constructors of this class should be similar to those used when creating the initial driver or data source, and should obviously be sufficient to create new `XAResources` that can be used to drive recovery.

Chapter 5

JDBC and transactions

Using the transactional JDBC driver

JBossJTA supports the construction of both local and distributed transactional applications which access databases using the JDBC 2.0 APIs. JDBC 2.0 supports two-phase commit of transactions, and is similar to the XA X/Open standard. The JDBC 2.0 support is found in the `com.arjuna.ats.jdbc` package.

The JDBC 2.0 support has been tested with the following drivers and databases:

- Merant's Sequelink 5.1 with Oracle 8.1.6
- Cloudscape 3.6 with Cloudscape's own database.
- Oracle 8.1.6/8.1.7 native driver
- MS SQL Server 2000 native driver

Managing transactions

JBossJTA must be able to associate work performed on a JDBC connection with a specific transaction. Therefore, implicit transaction propagation and/or indirect transaction management must be used by applications, i.e., for each JDBC connection it must be possible for *JBossJTA* to determine the invoking thread's current transaction context.

Restrictions

The following restrictions are imposed by limitations in the JDBC specifications and by *JBossJTA* to ensure that transactional interactions with JDBC databases can be correctly managed:

Nested transactions are not supported by JDBC 2.0. If an attempt is made to use a JDBC connection within a subtransaction, *JBossJTA* will throw a suitable exception and no work will be allowed on that connection. However, if you wish to have nested transactions, then you can set the `com.arjuna.ats.jta.supportSubtransactions` property to YES.

Transactional drivers

The *JBossJTA* approach to incorporating JDBC connections within transactions is to provide transactional JDBC drivers through which all interactions occur. These drivers intercept all invocations and ensure that they are registered with, and driven by, appropriate transactions. There is a single type of transactional driver through which any JDBC driver can be driven; obviously if the database is not transactional then ACID properties cannot be guaranteed.

This driver is `com.arjuna.ats.jdbc.TransactionalDriver`, which implements the `java.sql.Driver` interface.

Loading drivers

The driver may be directly instantiated and used within an application. For example:

```
TransactionalDriver arjunaJDBC2Driver = new TransactionalDriver();
```

They can be registered with the JDBC driver manager (`java.sql.DriverManager`) by adding them to the Java system properties. The `jdbc.drivers` property contains a list of driver class names, separated by colons, that are loaded by the JDBC driver manager when it is initialized.

```
/*
 * Register the driver via the system properties variable
 * "jdbc.drivers"
 */

Properties p = System.getProperties();

switch (dbType)
{
case MYSQL:
    p.put("jdbc.drivers", "org.gjt.mm.mysql.Driver");
    break;
case CLOUDSCAPE:
    p.put("jdbc.drivers", "COM.cloudscape.core.JDBCdriver");
    break;
}

System.setProperties(p);
```

Alternatively, the `Class.forName()` method may be used to load the driver or drivers:

```
Class.forName("sun.jdbc.odbc.JdbcOdbcDriver");
```

Calling `Class.forName()` will automatically register the driver with the JDBC driver manager. It is also possible to explicitly create an instance of the JDBC driver:

```
sun.jdbc.odbc.JdbcOdbcDriver drv = new sun.jdbc.odbc.JdbcOdbcDriver();
DriverManager.registerDriver(drv);
```

When you have loaded a driver, it is available for making a connection with a DBMS.

Connections

In this section we shall discuss the notion of transactional JDBC connections, how they are managed within *JBossJTA* and the implications on using them within an application.

Making the connection

Because JDBC connectivity in *JBossJTA* works by simply providing a new JDBC driver, application code can remain relatively the same to that when not using transactions. Typically, the application programmer need only start and terminate transactions.

JDBC 2.0

Before describing the JDBC 2.0 support it is necessary to mention that the following properties can be set and passed to the *JBossJTA* driver (they are all located in the `com.arjuna.ats.jdbc.TransactionalDriver` class):

- *userName*: the user name to use when attempting to connect to the database.
- *password*: the password to use when attempting to connect to the database.
- *createDb*: if set to true, the driver will attempt to create the database when it connects. This may not be supported by all JDBC 2.0 implementations.
- *dynamicClass*: this specifies a class to instantiate to connect to the database, rather than using JNDI.

XADataSources

JDBC 2.0 connections are created from appropriate `DataSources`. Those connections which must participate within distributed transactions are obtained from `XADataSources`. Therefore, when using a JDBC 2.0 driver, *JBossJTA* will use the appropriate `DataSource` whenever a connection to the database is made. It will then obtain `XAResources` and register them with the transaction via the JTA interfaces. It is these `XAResources` which the transaction service will use when the transaction terminates in order to drive the database to either commit or rollback the changes made via the JDBC connection.

There are two ways in which the *JBossJTA* JDBC 2.0 support can obtain `XADataSources`. These will be explained in the following sections. Note, for simplicity we shall assume that the JDBC 2.0 driver is instantiated directly by the application.

Java Naming and Directory Interface (JNDI)

In order to allow a JDBC driver to use arbitrary `DataSources` without having to know specific details about their implementations, `DataSources` are typically obtained from JNDI. A specific `(XA)DataSource` can be created and registered with an appropriate JNDI implementation, and the application (or JDBC driver) can later bind to and use it. Since JNDI only allows the application to see the `(XA)DataSource` as an instance of the interface (e.g., `javax.sql.XADataSource`) rather than as an instance of the implementation class (e.g., `com.mydb.myXADataSource`), the application is not tied at build time to only use a specific `(XA)DataSource` implementation.

To get the `TransactionalDriver` class to use a JNDI registered `XADataSource` it is first necessary to create the `XADataSource` instance and store it in an appropriate JNDI

implementation. Details of how to do this can be found in the JDBC 2.0 tutorial available at JavaSoft. An example is show below:

```
XADatasource ds = MyXADatasource();
Hashtable env = new Hashtable();
String initialCtx =
PropertyManager.getProperty("Context.INITIAL_CONTEXT_FACTORY");

env.put(Context.INITIAL_CONTEXT_FACTORY, initialCtx);

initialContext ctx = new InitialContext(env);

ctx.bind("jdbc/foo", ds);
```

Where the `Context.INITIAL_CONTEXT_FACTORY` property is the JNDI way of specifying the type of JNDI implementation to use.

Then the application must pass an appropriate connection URL to the JDBC 2.0 driver:

```
Properties dbProps = new Properties();

dbProps.setProperty(TransactionalDriver.userName, "user");
dbProps.setProperty(TransactionalDriver.password, "password");

TransactionalDriver arjunaJDBC2Driver = new TransactionalDriver();
Connection connection = arjunaJDBC2Driver.connect("jdbc:arjuna:jdbc/foo",
dbProps);
```

The JNDI URL must be pre-pended with *jdbc:arjuna:* in order for the `ArjunaJDBC2Driver` to recognise that the `DataSource` must participate within transactions and be driven accordingly.

Dynamic class instantiation

Many JDBC 2.0 implementations provide proprietary implementations of `XADatasources` that provide non-standard extensions to the specification. In order to allow the application to remain isolated from the actual JDBC 2.0 implementation it is using and yet continue to be able to use these extensions, *JBossJTA* hides the details of these proprietary implementations using dynamic class instantiation. In addition, the use of JNDI is not required when using this mechanism because the actual implementation of the `XADatasource` will be directly instantiated, albeit in a manner which will not tie an application or driver to a specific implementation. *JBossJTA* therefore has several classes which are for specific JDBC 2.0 implementations, and these can be selected at runtime by the application setting the *dynamicClass* property appropriately:

Database Type	Property Name
Cloudscape 3.6	com.arjuna.ats.internal.jdbc.drivers.cloudscape_3_6
Sequelink 5.1	com.arjuna.ats.internal.jdbc.drivers.sequelink_5_1
Oracle 8.1.6	com.arjuna.ats.internal.jdbc.drivers.oracle_8_1_6

Table 1: Dynamic Class property values for specific databases.

The application code must specify which dynamic class the `TransactionalDriver` should instantiate when setting up the connection:

```
Properties dbProps = new Properties();

dbProps.setProperty(TransactionalDriver.userName, "user");
dbProps.setProperty(TransactionalDriver.password, "password");
dbProps.setProperty(TransactionalDriver.dynamicClass,

"com.arjuna.ats.internal.jdbc.drivers.sequelink_5_0");

TransactionalDriver arjunaJDBC2Driver = new TransactionalDriver();
Connection connection =
arjunaJDBC2Driver.connect("jdbc:arjuna:sequelink://host:port;databaseName=f
oo", dbProperties);
```

Using the connection

Once the connection has been established (for example, using the `java.sql.DriverManager.getConnection` method), all operations on the connection will be monitored by *JBossJTA*. Note, it is not necessary to use the transactional connection within transactions. If a transaction is not present when the connection is used, then operations will be performed directly on the database.

Note: JDBC does not support subtransactions.

Transaction timeouts can be used to automatically terminate transactions should the connection not be terminated within an appropriate period.

JBossJTA connections can be used within multiple different transactions simultaneously, i.e., different threads, with different notions of the current transaction, may use the same JDBC connection. *JBossJTA* does connection pooling for each transaction within the JDBC connection. So, although multiple threads may use the same instance of the JDBC connection, internally this may be using a different connection instance per transaction. With the exception of *close*, all operations performed on the connection at the application level will only be performed on this transaction-specific connection.

JBossJTA will automatically register the JDBC driver connection with the transaction via an appropriate resource. When the transaction terminates, this resource will be responsible for either committing or rolling back any changes made to the underlying database via appropriate calls on the JDBC driver.

Once created, the driver and any connection can be used in the same way as any other JDBC driver or connection.

```
Statement stmt = conn.createStatement();

try
{
    stmt.executeUpdate("CREATE TABLE test_table (a INTEGER,b INTEGER)");
}
catch (SQLException e)
{
    // table already exists
}

stmt.executeUpdate("INSERT INTO test_table (a, b) VALUES (1,2)");

ResultSet res1 = stmt.executeQuery("SELECT * FROM test_table");
```

Connection pooling

For each user name and password, *JBossJTA* will maintain a single instance of each connection for as long as that connection is in use. Subsequent requests for the same connection will get a reference to the originally created connection, rather than a new instance. Attempts to close the connection are allowed, but the connection will only actually be closed when all users (including transactions) have either finished with the connection, or issued close calls.

Reusing connections

Some JDBC drivers allow the reuse of a connection for multiple different transactions once a given transaction has completed. Unfortunately this is not a common feature, and other drivers require a new connection to be obtained for each new transaction. By default, the *JBossJTA* transactional driver will always obtain a new connection for each new transaction. However, if an existing connection is available and is currently unused, it is possible to make *JBossJTA* reuse this connection. In order to do this, the `reuseconnection=true` option must be specified on the JDBC URL. For example:

```
jdbc:arjuna:sequelink://host:port;databaseName=foo;reuseconnection=true
```

Terminating the transaction

Whenever a transaction terminates (either explicitly by the application programmer, or implicitly when any associated transaction timeout expires) that has a JDBC connection registered with it, *JBossJTA* will drive the database (via the JDBC driver) to either commit or roll back any changes made to it. This happens transparently to the application.

AutoCommit

If `AutoCommit` of the `java.sql.Connection` is set to `true` for JDBC 1.0 then the execution of every SQL statement is a separate top-level transaction, and grouping multiple statements to be managed within a single OTS transaction is not possible. Therefore, *JBossJTA* will disable `AutoCommit` on JDBC 1.0 connections before they can be used. If

auto commit is subsequently set to true by the application, *JBossJTA* will raise the `java.sql.SQLException`.

Setting isolation levels

When using the *JBossJTA* JDBC driver, it may be necessary to set the underlying transaction isolation level on the XA connection. By default, this is set to `TRANSACTION_SERIALIZABLE`, but you may want to set this to something more appropriate for your application. In order to do this, set the `com.arjuna.ats.jdbc.isolationLevel` property to the appropriate isolation level in string form, e.g., `TRANSACTION_READ_COMMITTED`, or `TRANSACTION_REPEATABLE_READ`.

Note: At present this property applies to all XA connections created in the JVM.

Examples

JDBC example

The following code illustrates many of the points described above (note that for simplicity, much error checking code has been removed).

```

public class JDBCTest
{
public static void main (String[] args)
{
    /*
     * Register the driver via the system properties variable
     * "jdbc.drivers"
     */

    Properties p = System.getProperties();

    p.put("jdbc.drivers", "com.merant.sequelink.jdbc.SequeLinkDriver");

    System.setProperties(p);

    /*
     * Now register the JBossJTA 2.x transactional driver.
     */

    try
    {
        DriverManager.registerDriver(new
com.arjuna.ats.jdbc.TransactionalDriver());
    }
    catch (Exception e)
    {
        e.printStackTrace();
        System.exit(0);
    }

    Connection conn = null;
    Connection conn2 = null;
    Statement stmt = null;          // non-tx statement
    Statement stmtx = null;        // will be a tx-statement
    Properties dbProperties = new Properties();

    try
    {
        System.out.println("\nCreating connection to database: "+url);

        dbProperties.put(TransactionalDriver.userName, "foo");
        dbProperties.put(TransactionalDriver.password, "bar");
        dbProperties.put(TransactionalDriver.dynamicClass,

"com.arjuna.ats.internal.jdbc.drivers.sequelink_5_0");

        conn = DriverManager.getConnection(url, dbProperties);
    }
}
}

```

```

conn2 = DriverManager.getConnection(url, dbProperties);
stmt = conn.createStatement(); // non-tx statement

try
{
    stmt.executeUpdate("DROP TABLE test_table");
    stmt.executeUpdate("DROP TABLE test_table2");
}
catch (Exception e)
{
    // assume not in database.
}

try
{
    stmt.executeUpdate("CREATE TABLE test_table (a INTEGER,b
INTEGER)");
    stmt.executeUpdate("CREATE TABLE test_table2 (a INTEGER,b
INTEGER)");
}
catch (Exception e)
{
}

try
{
    System.out.println("Starting top-level transaction.");

com.arjuna.ats.jta.UserTransaction.transactionManager().begin();

    stmtx = conn.createStatement(); // will be a tx-statement

    System.out.println("\nAdding entries to table 1.");

    stmtx.executeUpdate("INSERT INTO test_table (a, b) VALUES
(1,2)");

    ResultSet res1 = null;

    System.out.println("\nInspecting table 1.");

    res1 = stmtx.executeQuery("SELECT * FROM test_table");
    while (res1.next())
    {
        System.out.println("Column 1: "+res1.getInt(1));
        System.out.println("Column 2: "+res1.getInt(2));
    }

    System.out.println("\nAdding entries to table 2.");

    stmtx.executeUpdate("INSERT INTO test_table2 (a, b) VALUES
(3,4)");

    res1 = stmtx.executeQuery("SELECT * FROM test_table2");
    System.out.println("\nInspecting table 2.");

    while (res1.next())
    {
        System.out.println("Column 1: "+res1.getInt(1));
        System.out.println("Column 2: "+res1.getInt(2));
    }

    System.out.print("\nNow attempting to rollback changes.");
}

```

```

com.arjuna.ats.jta.UserTransaction.transactionManager().rollback();

com.arjuna.ats.jta.UserTransaction.transactionManager().begin();
    stmtx = conn.createStatement();
    ResultSet res2 = null;

    System.out.println("\nNow checking state of table 1.");

    res2 = stmtx.executeQuery("SELECT * FROM test_table");
    while (res2.next())
    {
        System.out.println("Column 1: "+res2.getInt(1));
        System.out.println("Column 2: "+res2.getInt(2));
    }

    System.out.println("\nNow checking state of table 2.");

    stmtx = conn.createStatement();
    res2 = stmtx.executeQuery("SELECT * FROM test_table2");
    while (res2.next())
    {
        System.out.println("Column 1: "+res2.getInt(1));
        System.out.println("Column 2: "+res2.getInt(2));
    }

com.arjuna.ats.jta.UserTransaction.transactionManager().commit(true);
    }
    catch (Exception ex)
    {
        ex.printStackTrace();
        System.exit(0);
    }
    }
    catch (Exception sysEx)
    {
        sysEx.printStackTrace();
        System.exit(0);
    }
    }
}

```

Failure recovery example

This class implements the XAResourceRecovery interface for XAResources. The parameter supplied in setParameters can contain arbitrary information necessary to initialize the class once created. In this instance it contains the name of the property file in which the db connection information is specified, as well as the number of connections that this file contains information on (separated by ;).

Caution: This is only an *example* of the sorts of things an XAResourceRecovery implementer could do. This implementation uses a property file which is assumed to contain sufficient information to recreate connections used during the normal run of an application so that we can perform recovery on them. It is not recommended that information such as user name and password appear in such a raw text format as it opens up a potential security hole.

The db parameters specified in the property file are assumed to be in the format:

- DB_x_DatabaseURL=
- DB_x_DatabaseUser=
- DB_x_DatabasePassword=
- DB_x_DatabaseDynamicClass=

where x is the number of the connection information.

```
public class BasicXARecovery implements XAResourceRecovery
{
    /*
     * Some XAResourceRecovery implementations will do their startup work
     * here, and then do little or nothing in setDetails. Since this one
needs
     * to know dynamic class name, the constructor does nothing.
     */

    public BasicXARecovery () throws SQLException
    {
        numberOfConnections = 1;
        connectionIndex = 0;
        props = null;
    }

    /**
     * The recovery module will have chopped off this class name already.
     * The parameter should specify a property file from which the url,
     * user name, password, etc. can be read.
     */

    public boolean initialise (String parameter) throws SQLException
    {
        int breakPosition = parameter.indexOf(BREAKCHARACTER);
        String fileName = parameter;

        if (breakPosition != -1)
        {
            fileName = parameter.substring(0, breakPosition -1);

            try
            {
                numberOfConnections =
Integer.parseInt(parameter.substring(breakPosition +1));
            }
            catch (NumberFormatException e)
            {
                return false;
            }
        }

        jdbcPropertyManager.propertyManager.addPropertyFile(fileName);

        try
        {
            jdbcPropertyManager.propertyManager.loadProperties(true);

            props = jdbcPropertyManager.propertyManager.getProperties();
        }
    }
}
```

```

    }
    catch (Exception e)
    {
        return false;
    }

    return true;
    }

    public synchronized XAResources getXAResource () throws SQLException
    {
        JDBC2RecoveryConnection conn = null;

        if (hasMoreConnections())
        {
            connectionIndex++;

            conn = getStandardConnection();

            if (conn == null)
                conn = getJNDIConnection();

            if (conn == null)
            {
            }
        }
    }

    return conn.recoveryConnection().getConnection().getXAResource();
    }

    public synchronized boolean hasMoreResources ()
    {
        if (connectionIndex == numberOfConnections)
            return false;
        else
            return true;
    }

    private final JDBC2RecoveryConnection getStandardConnection () throws
    SQLException
    {
        String number = new String(""+connectionIndex);
        String url = new String(dbTag+number+urlTag);
        String password = new String(dbTag+number+passwordTag);
        String user = new String(dbTag+number+userTag);
        String dynamicClass = new String(dbTag+user+dynamicClassTag);

        Properties dbProperties = new Properties();

        String theUser = props.getProperty(user);
        String thePassword = props.getProperty(password);

        if (theUser != null)
        {
            dbProperties.put(TransactionalDriver.userName, theUser);
            dbProperties.put(TransactionalDriver.password, thePassword);

            String dc = props.getProperty(dynamicClass);

            if (dc != null)
                dbProperties.put(TransactionalDriver.dynamicClass, dc);

            return new JDBC2RecoveryConnection(url, dbProperties);
        }
    }

```

```

    }
    else
        return null;
    }

    private final JDBC2RecoveryConnection getJNDIConnection () throws
SQLException
    {
        String number = new String(""+connectionIndex);
        String url = new String(dbTag+jndiTag+number+urlTag);
        String password = new String(dbTag+jndiTag+number+passwordTag);
        String user = new String(dbTag+jndiTag+number+userTag);

        Properties dbProperties = new Properties();

        String theUser = props.getProperty(user);
        String thePassword = props.getProperty(password);

        if (theUser != null)
        {
            dbProperties.put(TransactionalDriver.userName, theUser);
            dbProperties.put(TransactionalDriver.password, thePassword);

            return new JDBC2RecoveryConnection(url, dbProperties);
        }
        else
            return null;
    }

    private int        numberOfConnections;
    private int        connectionIndex;
    private Properties props;

    private static final String dbTag = "DB_";
    private static final String urlTag = "_DatabaseURL";
    private static final String passwordTag = "_DatabasePassword";
    private static final String userTag = "_DatabaseUser";
    private static final String dynamicClassTag = "_DatabaseDynamicClass";
    private static final String jndiTag = "JNDI_";

    /*
     * Example:
     *
     * DB2_DatabaseURL=jdbc\:arjuna\:sequelink\://qa02\:20001
     * DB2_DatabaseUser=tester2
     * DB2_DatabasePassword=tester
     *
     DB2_DatabaseDynamicClass=com.arjuna.ats.internal.jdbc.drivers.sequelink_5_1
     *
     * DB_JNDI_DatabaseURL=jdbc\:arjuna\:jndi
     * DB_JNDI_DatabaseUser=tester1
     * DB_JNDI_DatabasePassword=tester
     * DB_JNDI_DatabaseName=empay
     * DB_JNDI_Host=qa02
     * DB_JNDI_Port=20000
     */

    private static final char BREAKCHARACTER = ';'; // delimiter for
parameters
}

```

The class `com.arjuna.ats.internal.jdbc.recovery.JDBC2RecoveryConnection` may be used to create a new connection to the database using the same parameters that were used to create the initial connection.

Chapter 7

Configuring JBossJTA

Configuration options

The following table shows the configuration features, with default values shown in italics. For more detailed information, the relevant section numbers are provided.

Configuration Name	Possible Values	Relevant Section
<code>com.arjuna.ats.jta.supportSubtransactions</code>	YES/NO	Chapter 5
<code>com.arjuna.ats.jta.jtaTMImplementation</code>	<code>com.arjuna.ats.internal.jta.transaction.arjunacore.TransactionManagerImple/com.arjuna.ats.internal.jta.transaction.jts.TransactionManagerImple</code>	Chapter 2
<code>com.arjuna.ats.jta.jtaUTImplementation</code>	<code>com.arjuna.ats.internal.jta.transaction.arjunacore.UserTransactionImple/com.arjuna.ats.internal.jta.transaction.jts.UserTransactionImple</code>	Chapter 1, Chapter 2
<code>com.arjuna.ats.jta.xaBackoffPeriod</code>		
<code>com.arjuna.ats.jdbc.isolationLevel</code>	Any supported JDBC isolation level.	Chapter 5

Table 2: JBossJTA configuration options.

Using JBossJTA in application servers

JBOSS Application Server

Configuration

Service Configuration

The Arjuna Transaction Service is configured primarily via the XML files stored in the `etc` directory, but when run as a JBOSS service there are a number of configurable attributes available. They are as follows:

TransactionTimeout – The default transaction timeout to be used for new transactions. Specified as an integer in seconds.

StatisticsEnabled – This determines whether or not the transaction service should gather statistical information. This information can then be viewed using the `PerformanceStatistics` MBean. Specified as a Boolean. The default is to not gather this information.

PropagateFullContext – This determines whether a full transactional context is propagated by context importer/exporter. If set to false only the current transaction context is propagated. If set to true the full transaction context (including parent transactions) is propagated.

These attributes are specified as MBean attributes in the `jboss-service.xml` file located in the `server/all/conf` directory, e.g.

```
<mbean code="com.arjuna.ats.jbossatx.jts.TransactionManagerService"
name="jboss:service=TransactionManager">

  <attribute name="TransactionTimeout">300</attribute>
  <attribute name="StatisticsEnabled">true</attribute>

</mbean>
```

The transaction service is configurable also via the standard Arjuna Transaction Service property files. These are located in the JBossTS install location under the `etc` sub-directory. These files can be edited manually or through JMX. Each property file is exposed via an object with the name `com.arjuna.ts.properties` and an attribute of `module` where `module` is equal to the name of the module to be configured, e.g. `com.arjuna.ts.properties:module=arjuna`.

The services

There is currently one service offered by the JBOSS integration. In this section we shall discuss what this service does.

TransactionManagerService

The transaction manager service's main purpose is to ensure the recovery manager is started. It also binds the JBossTS JTA transaction manager to `java:/TransactionManager` name with the JNDI provider. This service depends upon the existence of the CORBA ORB Service and it must be using JacORB as the underlying ORB implementation.

There are two instances of this service:

- **distributed:** this uses the JTS enabled transaction manager implementation and hence supports distributed transactions and recovery. To configure this use the `com.arjuna.ats.jbossatx.jts.TransactionManagerService` class. This is the default configuration.
- **local:** this uses the purely local JTA implementation. To configure this use the `com.arjuna.ats.jbossatx.jta.TransactionManagerService` class.

Ensuring Transactional Context is Propagated to the Server

It is possible to coordinate transactions from a coordinator which is not located within the JBoss server (e.g. using transactions created by an external OTS server). To ensure the transaction context is propagated via JRMP invocations to the server, the transaction propagation context factory needs to be explicitly set for the JRMP invoker proxy. This is done as follows:

```
JRMPInvokerProxy.setTPCFactory( new  
com.arjuna.ats.internal.jbossatx.jts.PropagationContextManager() );
```

Index

- AutoCommit, **30**
- Configurable options, **39**
- DataSource, **27**
- Dynamic class instantiation, **28**
- Enlisting resources, **13**
- Example, *32*
- Failure recovery, **22**
- JBoss Transactions compliance, *10*
- JBoss Transactions restrictions, *24*
- JDBC 2.0, **24**
 - properties, *27*
- JDBC2 Recovery, **22**
- JNDI string extension, **28**
- Loading drivers, *25*
- Making connections, *27*
- Obtaining TransactionManager, **11**
- Obtaining UserTransaction, **11**
- Package, **11**
- Property variables
 - XAConnectionRecovery, *22*
- Relationship to JTA, **27**
- Supported JDBC drivers, **24**
- Thread restrictions, **12**
 - suspending and resuming transactions, *13*
 - terminating transactions, *13*
- Transaction synchronization, **14**
- Transactional drivers, *24*
- Transactional Objects for Java
 - configuration, *39*
- TransactionManager, **11**
- UserTransaction, **11**
- Using connections, *30*
- XAConnectionRecovery, **21, 22**
- XADatasource, **27**
- XAResource, **16, 27**
 - failure recovery, *21*