



Orbix 6.3.8

A decorative graphic consisting of several overlapping, wavy blue lines that curve and flow across the lower half of the page, creating a sense of motion and depth.

**CORBA OTS Guide:
C++**

Micro Focus
The Lawn
22-30 Old Bath Road
Newbury, Berkshire RG14 1QN
UK

<http://www.microfocus.com>

Copyright © Micro Focus 2015. All rights reserved.

MICRO FOCUS, the Micro Focus logo and Micro Focus Licensing are trademarks or registered trademarks of Micro Focus IP Development Limited or its subsidiaries or affiliated companies in the United States, United Kingdom and other countries.

All other marks are the property of their respective owners.

12/16/15

Contents

Preface	v
Contacting Micro Focus	vii
Transaction Service	1
About Transactions.....	1
Transaction Managers.....	2
OMG OTS and X/Open XA Interfaces	5
Transaction Interfaces	5
OTS Interfaces	6
The X/Open XA Interface.....	8
Getting Started with Transactions	11
Application Overview	11
Transaction Demarcation.....	13
Transaction Propagation and POA Policies	14
XA Resource Manager Integration.....	16
Application-Specific Resources	18
Configuration Issues.....	19
Transaction Demarcation and Control	21
The OTS Current Object	21
Direct Transaction Demarcation	26
Propagation and Transaction Policies	29
Implicit Propagation Policies	29
Shared and Unshared Transactions.....	30
Policy Meanings	30
Example Use of an OTSPolicy.....	32
Example Use of a NonTxTargetPolicy	34
Use of the ADAPTS OTSPolicy	36
Orbix-Specific OTSPolicies	37
Migrating from TransactionPolicies.....	40
Explicit Propagation	41
Using XA Resource Managers with OTS	43
The XA Interface	43
XA and Multi-Threading.....	45
Using the Orbix XA Plug-In	46
Associations between Transactions and Connections	47
Association State Diagram.....	48
Using a Remote Resource Manager	49
Transaction Management	53
Synchronization Objects.....	53
Transaction Identity Operations	55
Transaction Status	56
Transaction Relationships	58

Recreating Transactions	59
Writing Recoverable Resources	63
The Resource Interface	63
Creating and Registering Resource Objects	65
Resource Protocols	67
Responsibilities and Lifecycle of a Resource Object	75
Interoperability	79
Use of InvocationPolicies	79
Use of the TransactionalObject Interface	79
Interoperability with Orbix 3 OTS Applications	80
Using the Orbix 3 otstf with Orbix Applications	83
OTS Plug-Ins and Deployment Options	85
The OTS Plug-In	86
The OTS Lite Plug-In	88
The Encina Transaction Manager	89
The itotstm Transaction Manager Service	90
Appendix OTS Management	93
Introduction to OTS Management	93
TransactionManager Entity	95
Transaction Entity	97
Encina Transaction Log Entity	99
Encina Volume Entity	100
Management Events	101
Glossary	103
Index	107

Preface

Orbix OTS is a full implementation from Micro Focus of the interoperable transaction service as specified by the Object Management Group. Orbix OTS complies with the following specifications:

- CORBA 2.6
- OTS 1.2
- GIOP 1.2 (default), 1.1, and 1.0

Audience

This guide is intended to help you become familiar with the transaction service, and shows how to develop applications with it. This guide assumes that you are familiar with CORBA concepts, and with C++.

This guide does not discuss every interface and its operations in detail, but gives a general overview of the capabilities of the transaction service and how various components fit together. For detailed information about individual operations, refer to the *CORBA Programmer's Reference*.

Organization of this Guide

This guide is divided as follows:

[“Transaction Service”](#) provides a brief overview of the basic concepts involved in using the transactions service.

[“OMG OTS and X/Open XA Interfaces”](#) provides an overview of the transaction service's interfaces. It also provides information on the X/Open XA interfaces and how to use them to interact with compliant resources.

[“Getting Started with Transactions”](#) is a simple example of the steps involved in developing a client that uses the transaction service. It discusses the basic steps required to use transactions and the concepts behind them.

[“Transaction Demarcation and Control”](#) covers transaction demarcation. It covers both using the transactions `Current` object, which is convenient but limited, and using the `TransactionFactory` and the `Terminator` interfaces to directly manipulate demarcation.

[“Propagation and Transaction Policies”](#) covers how to control how the transaction is propagated to its target object through the use of POA policies.

[“Using XA Resource Managers with OTS”](#) provides a detailed discussion how to implement `CosTransactions::Resource` objects on top of the standard X/Open XA interface to manage transactional resources.

["Transaction Management"](#) covers some additional areas of transaction management. This includes synchronization objects, transaction identity and status operations, relationships between transactions and recreating transactions.

["Writing Recoverable Resources"](#) describes the `CosTransactions::Resource` interface; how resource objects participate in the transaction protocols and the requirements for implementing resource objects.

["Interoperability"](#) describes how the Orbix OTS interoperates with older releases of Orbix and with other OTS implementations including the Orbix 3 OTS.

["OTS Plug-Ins and Deployment Options"](#) discusses the plugins that implement the transaction service and options for deploying them.

Typographical Conventions

This guide uses the following typographical conventions:

`Constant width` Constant width (courier font) in normal text represents portions of code and literal names of items such as classes, functions, variables, and data structures. For example, text might refer to the `CORBA::Object` class.

Constant width paragraphs represent code examples or information a system displays on the screen. For example:

```
#include <stdio.h>
```

Italic Italic words in normal text represent *emphasis* and *new terms*.

Italic words or characters in code and commands represent variable values you must supply, such as arguments to commands or path names for your particular system. For example:

```
% cd /users/your_name
```

Note: Some command examples may use angle brackets to represent variable values you must supply. This is an older convention that is replaced with *italic* words or characters.

Keying Conventions

This guide may use the following keying conventions:

No prompt	When a command's format is the same for multiple platforms, a prompt is not used.
%	A percent sign represents the UNIX command shell prompt for a command that does not require root privileges.
#	A number sign represents the UNIX command shell prompt for a command that requires root privileges.
>	The notation > represents the DOS, Windows NT, Windows 95, or Windows 98 command prompt.
...	Horizontal or vertical ellipses in format and syntax descriptions indicate that material has been eliminated to simplify a discussion.
[]	Brackets enclose optional items in format and syntax descriptions.
{ }	Braces enclose a list from which you must choose an item in format and syntax descriptions.
	A vertical bar separates items in a list of choices enclosed in { } (braces) in format and syntax descriptions.

Contacting Micro Focus

Our Web site gives up-to-date details of contact numbers and addresses.

Further Information and Product Support

Additional technical information or advice is available from several sources.

The product support pages contain a considerable amount of additional information, such as:

- The *Product Updates* section of the Micro Focus SupportLine Web site, where you can download fixes and documentation updates.
- The *Examples and Utilities* section of the Micro Focus SupportLine Web site, including demos and additional product documentation.

To connect, enter <http://www.microfocus.com> in your browser to go to the Micro Focus home page, then click *Support*.

Note:

Some information may be available only to customers who have maintenance agreements.

If you obtained this product directly from Micro Focus, contact us as described on the Micro Focus Web site, <http://www.microfocus.com>. If you obtained the product from another source, such as an authorized distributor, contact them for help first. If they are unable to help, contact us.

Also, visit:

- The Micro Focus Community Web site, where you can browse the Knowledge Base, read articles and blogs, find demonstration programs and examples, and discuss this product with other users and Micro Focus specialists.
- The Micro Focus YouTube channel for videos related to your product.

Information We Need

However you contact us, please try to include the information below, if you have it. The more information you can give, the better Micro Focus SupportLine can help you. But if you don't know all the answers, or you think some are irrelevant to your problem, please give whatever information you have.

- The name and version number of all products that you think might be causing a problem.
- Your computer make and model.
- Your operating system version number and details of any networking software you are using.
- The amount of memory in your computer.
- The relevant page reference or section in the documentation.
- Your serial number. To find out these numbers, look in the subject line and body of your Electronic Product Delivery Notice email that you received from Micro Focus.

Contact information

Our Web site gives up-to-date details of contact numbers and addresses.

Additional technical information or advice is available from several sources.

The product support pages contain considerable additional information, including the WebSync service, where you can download fixes and documentation updates. To connect, enter <http://www.microfocus.com> in your browser to go to the Micro Focus home page.

If you are a Micro Focus SupportLine customer, please see your SupportLine Handbook for contact information. You can download it from our Web site or order it in printed form from your sales representative. Support from Micro Focus may be available only to customers who have maintenance agreements.

You may want to check these URLs in particular:

- <http://www.microfocus.com/products/corba/orbix/orbix-6.aspx> (trial software download and Micro Focus Community files)
- https://supportline.microfocus.com/productdoc.aspx_ (documentation updates and PDFs)

To subscribe to Micro Focus electronic newsletters, use the online form at:

<http://www.microfocus.com/Resources/Newsletters/infocus/newsletter-subscription.asp>

Transaction Service

This chapter describes the transaction processing capabilities of Orbix, showing how to use the Object Transaction Service (OTS) for transaction demarcation, propagation and integration with resource managers. Integration with X/Open XA compliant resource managers is described.

About Transactions

What is a transaction?

Orbix gives separate software objects the power to interact freely even if they are on different platforms or written in different languages. Orbix adds to this power by permitting those interactions to be transactions.

What is a transaction? Ordinary, non-transactional software processes can sometimes proceed and sometimes fail, and sometimes fail after only half completing their task. This can be a disaster for certain applications. The most common example is a bank fund transfer: imagine a failed software call that debited one account but failed to credit another. A transactional process, on the other hand, is secure and reliable as it is guaranteed to succeed or fail in a completely controlled way.

Transaction support in Orbix

To support the development of object-oriented, distributed, transaction-processing applications, Orbix offers:

- An implementation of the Object Management Group's Object Transaction Service (OMG OTS).
- Integration with resource managers supporting the X/Open XA interface.
- A pluggable architecture that supports both a lightweight OTS implementation and a full recoverable two-phase-commit (2PC) implementation.

Example

The classical illustration of a transaction is that of funds transfer in a banking application. This involves two operations: a debit of one account and a credit of another (perhaps after extracting an appropriate fee). To combine these operations into a single unit of work, the following properties are required:

- If the debit operation fails, the credit operation should fail, and vice-versa; that is, they should both work or both fail.
- The system goes through an inconsistent state during the process (between the debit and the credit). This inconsistent state should be hidden from other parts of the application.

- It is implicit that committed results of the whole operation are permanently stored.

Properties of transactions

The following points illustrate the so-called ACID properties of a transaction.

Atomic	A transaction is an all or nothing procedure – individual updates are assembled and either committed or aborted (rolled back) simultaneously when the transaction completes.
Consistent	A transaction is a unit of work that takes a system from one consistent state to another.
Isolated	While a transaction is executing, its partial results are hidden from other entities accessing the transaction.
Durable	The results of a transaction are persistent.

Thus a transaction is an operation on a system that takes it from one persistent, consistent state to another.

Transaction Managers

Purpose of a Transaction Manager

Most resource managers, for example databases and message queues, provide support for native transactions. However, when an application wants two or more resource managers to be part of the same transaction some third party must provide the necessary coordination to ensure the ACID properties are guaranteed for the distributed transaction. This is where the concept of a transaction manager that is independent of the individual resource manager comes in.

The application uses the transaction manager to create the transaction. Each resource manager accessed during the transaction becomes a participant in the transaction. Then when the application completes the transaction, either with a commit or rollback request, the transaction manager communicates with each resource manager.

Two-phase commit protocol

When there are two or more participants involved in a transaction the transaction manager uses a two-phase-commit (2PC) protocol to ensure that all participants agree on the final outcome of the transaction despite any failures that may occur. Briefly the 2PC protocol works as follows:

- In the first phase, the transaction manager sends a “prepare” message to each participant. Each participant responds to this message with a vote indicating whether the transaction should be committed or rolled back.
- The transaction manager collects all the prepare votes and makes a decision on the outcome of the transaction. If all participants voted to commit the transaction may commit. However if a least one participant voted to rollback the transaction is rolled back. This completes the first phase.
- In the second phase the transaction manager sends either commit or rollback messages to each participant.

The 2PC protocol guarantees the ACID properties despite any failures that may occur. Usually the transaction manager uses a log to record the progress of the 2PC protocol so that messages can be replayed during recovery.

One-phase-commit protocol

If there is only one participant in the transaction the transaction manager can use a one-phase-commit (1PC) protocol instead of the 2PC protocol which can be expensive in terms of the number of messages sent and the data that must be logged. The 1PC protocol essentially delegates the transaction completion to the single resource manager. Orbix supports this 1PC protocol which allows developers to make use of the Orbix transaction manager without suffering the overheads associated with the 2PC protocol. By making use of the OTS and XA interfaces an application can be easily extended to support multiple resource managers within a transaction easily.

OMG OTS and X/Open XA Interfaces

The OMG OTS provides interfaces to manage the demarcation of transactions and the propagation of transaction contexts. With the X/Open XA interface, integration with compliant resource managers such as databases and message queues is provided.

Transaction Interfaces

Purpose

The OMG OTS provides interfaces to manage the demarcation of transactions (creation and completion), the propagation of transaction contexts to the participants of the transaction and interfaces to allow applications to participate in the transaction.

With the X/Open XA interface, integration with compliant resource managers such as databases and message queues is provided.

Figure 1 shows these areas of transaction management.

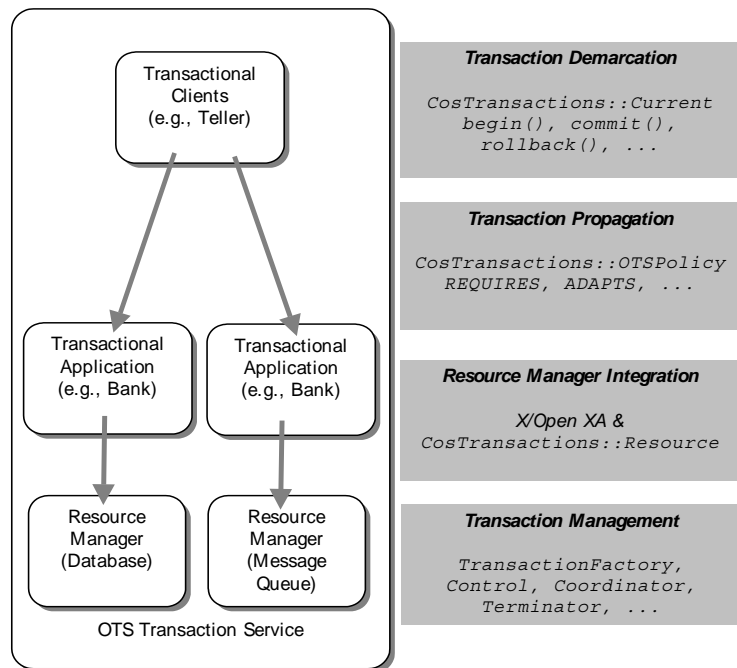


Figure 1: OTS and XA

Transaction Demarcation

Transaction demarcation is where the application sets the boundaries of the transaction. Typically this is done using the `OTS Current` interface; invoking the `begin()` operation at the start of the transaction and either `commit()` or `rollback()` at the end of the transaction. An alternative to using the `Current` interface is to create transactions directly using the `TransactionFactory` interface and commit or rollback the transactions using the `Terminator` interface.

Transaction Propagation

Propagation refers to the passing of information related to the transaction to the application objects that are participants in the transaction. When the `Current` interface is used for transaction demarcation this propagation takes place transparently and is controlled by a number of POA policies. Transactions created using the `TransactionFactory` interface must be propagated by adding an extra parameter to the operation.

Resource Manager Integration

Integration with resource managers such as databases is done using the XA interface. Alternatively an application may use the `OTS Resource` interface to provide integration with proprietary resource managers.

Transaction Management

The `OTS` interfaces also provide operations for general transaction management. These include, setting timeouts, registering resource objects and synchronization objects, comparing transactions and getting transaction names

OTS Interfaces

Supported OTS Interfaces

The following is a list of the main interfaces supported by the `OTS`. All interfaces are part of the IDL module `CosTransactions`. For more details on these interfaces, refer to the *CORBA Programmer's Reference*.

Table 1: *OTS Interfaces*

Interface	Purpose
Control	The return type of <code>TransactionFactory::create()</code> . It provides access to the two controllers of the transactions, the <code>Coordinator</code> and the <code>Terminator</code> .
Coordinator	Provides operations to register objects that participate in the transaction.
Current	A local interface that provides the concept of a transaction to the current thread of control. The <code>Current</code> interface supports a subset of the operations provided by the <code>Coordinator</code> and <code>Terminator</code> interfaces.
RecoveryCoordinator	Used in certain failure cases to complete the transaction completion protocol for a registered resource object.
Resource	Represents a recoverable participant in a transaction. Objects supporting this interface are registered with a transaction's coordinator, and are then invoked at key points in the transaction's completion.
SubtransactionAware Resource	Represents a participant that is aware of nested transactions. Nested transactions are not supported in this release.
Synchronization	Represents a non-recoverable object allowing application specific operations to occur both before and after transaction completion.
Terminator	Provides a means of directly committing or rolling back a transaction.
TransactionalObject	This interface has been deprecated and replaced with transaction policies (see " Propagation and Transaction Policies ").
TransactionFactory	Provides a means of directly creating top-level transactions.

OTS Transaction Modes

When using the OTS interfaces for transaction demarcation and propagation, there are two modes of use:

Indirect/Implicit	In the indirect/implicit mode transactions are created, committed and rolled back using the <code>Current</code> interface. Propagation takes place automatically depending on the policies in the target object's POA.
Direct/Explicit	In the direct/explicit mode transactions are created using the <code>TransactionFactory</code> and committed or rolled back using the <code>Terminator</code> object. Propagation is done by adding a parameter (for example, the transaction's control object) to each IDL operation.

The preferred mode for most applications is the indirect/implicit mode. The direct/explicit provides more flexibility but is more difficult to manage (see ["Direct Transaction Demarcation" on page 26](#) and ["Explicit Propagation" on page 41](#)) for more details.

The X/Open XA Interface

XA Interfaces

The X/Open XA interface is a C API between a transaction manager and a resource manager (for example, a database). The C API provides functions for opening and closing connections to the resource manager (`xa_open()` and `xa_close()`), managing associations between the current connection and global transactions (`xa_start()` and `xa_end()`), transaction protocols (`xa_prepare()`, `xa_commit()`, `xa_rollback()` and `xa_forget()`), and functions to support recovery (`xa_recover()`).

Integration with OTS

Integration between XA compliant resource managers and the OTS is provided by several interfaces in the XA module, as detailed in the following table.

Table 2: *XA interfaces.*

Interface	Purpose
Connector	Provides a means of getting <code>CurrentConnection</code> and <code>ResourceManager</code> objects.
<code>CurrentConnection</code>	Represents the current XA connection to a resource manager.

Table 2: *XA interfaces.*

Interface	Purpose
BeforeCompletionCallback	Allows an application to be called before the completion of a transaction.
ResourceManager	Use to register and unregister BeforeCompletionCallback objects.

Getting Started with Transactions

This chapter illustrates the Object Transaction Service (OTS) by way of an example application. It includes the basic steps needed to develop an application with the OTS.

Application Overview

Funds transfer application

The example application is that of funds transfer between two bank accounts. [Figure 2](#) shows the application. The client has a reference to two objects representing two accounts. The account objects are implemented directly on top of an XA-compliant database and use SQL to access the database. This example shows the source and destination accounts using different databases, however they could both be using the same database.




Figure 2: Example OTS Application – Funds Transfer

Interface definition

The interface for the account objects is defined in IDL as follows:

```
// IDL
module Bank
{
    typedef float CashAmount;
    interface Account
    {
        exception InsufficientFunds {};
        void deposit(in CashAmount amt);
        void withdraw(in CashAmount amt)
            raises (InsufficientFunds);
    };
    ...
};
```

TransactionalObject interface deprecated

Readers familiar with version 1.1 of the OTS specification (used by OrbixOTM and Orbix 3) will notice that the `Account` interface does not inherit from the `CosTransactions::TransactionalObject` interface. The use of that interface to mark objects as transactional has been deprecated in favor of using POA policies in version 1.2 of the specification. The `TransactionalObject` interface is still supported for backward compatibility with the OTS in OrbixOTM and Orbix 3. See [“Use of the TransactionalObject Interface” on page 79](#) for more details.

Since the `TransactionalObject` interface is deprecated, application developers no longer have to change the IDL used by their applications when adding transactional capabilities.

Transferring funds

Given a source and destination account, the funds transfer is performed by invoking the `withdraw()` operation on the source account followed by invoking the `deposit()` operation on the destination account. The application will look something like the following:

```
// C++
Bank::Account_var src_acc = ...
Bank::Account_var dest_acc = ...
Bank::CashAmount amount = 100.0;
src_acc->withdraw(amount);
dest_acc->deposit(amount);
```

Completing the application

To make this a transactional application we need three more steps:

1. The funds transfer application needs to be wrapped in a transaction to ensure the ACID properties. This is covered in [“Transaction Demarcation” on page 13](#).
2. The application must make sure the transaction is propagated to the two account objects during the invocations of the `deposit()` and `withdraw()` operations. This is covered in [“Transaction Propagation and POA Policies” on page 14](#)
3. The implementation of the account objects must be integrated with an XA compliant database. This is covered in [“XA Resource Manager Integration” on page 16](#).

Transaction Demarcation

Demarcation using OTS current object

Transaction demarcation refers to setting the boundaries of the transaction. The simplest way to do this is to use the OTS current object. The following are the steps involved:

1. Obtain a reference to the OTS current object from the ORB.
2. Create a new transaction.
3. Perform the funds transfer.
4. Complete the transaction by either committing it or rolling it back.

More information on transaction demarcation including other ways of creating, committing and rolling back transactions is covered in ["Transaction Demarcation and Control"](#).

Obtain a reference to the OTS current object from the ORB

The OTS current object supports the `CosTransactions::Current` interface and a reference to the object is obtained by calling the ORB operation `resolve_initial_references("TransactionCurrent")`.

Note that the file `CosTransactions.hh` must be included to use the interfaces defined in the `CosTransactions` module. Error handling has been omitted for clarity:

```
// C++
...
#include <CosTransactions.hh>
...
int main(int argc, char** argv)
{
    CORBA::ORB_var orb = CORBA::ORB_init(argc, argv);

    CORBA::Object_var obj =

        orb->resolve_initial_references("TransactionCurrent");
    CosTransactions::Current_var tx_current =
        CosTransactions::Current::_narrow(obj);
    ...
}
```

Create a new transaction

The next step is the creation of a new top-level transaction. This is done by invoking `begin()` on the OTS current object:

```
// C++
tx_current->begin();
```

If the `begin()` succeeds, a new transaction is associated with the current thread of control.

Perform the funds transfer

The funds transfer is the same as shown in the application overview. There are no changes for transaction management. The code is reproduced here for completeness:

```
// C++
Bank::Account_var src_acc = ...
Bank::Account_var dest_acc = ...
Bank::CashAmount amount = 100.0;
src_acc->withdraw(amount);
dest_acc->deposit(amount);
```

Complete the transaction by either committing it or rolling it back

Once the work has been done, we need to complete the transaction. Most of the time the application simply wants to attempt to commit the changes made: this is done by invoking the `commit()` operation on the OTS current object:

```
// C++
try {
    tx_current->commit(IT_false)
} catch (CORBA::TRANSACTION_ROLLEDBACK&) {
    // Transaction has been rolled back.
}
```

The `commit()` operation only attempts to commit the transaction. It may happen that due to system failures or other reasons the transaction cannot be committed; in this case the `TRANSACTION_ROLLEDBACK` system exception is raised.

The parameter passed to `commit()` is a boolean specifying whether heuristics outcomes should be reported to the client (see [“Heuristic Outcomes” on page 71](#) for details on heuristic outcomes). In this example we do not wait for heuristic outcomes.

If instead of attempting a commit the application wants to roll back the changes made, the operation `rollback()` is invoked on the OTS current object:

```
// C++
tx_current->rollback()
```

Transaction Propagation and POA Policies

Propagating the transaction

The funds transfer application invokes the `withdraw()` and `deposit()` operations within the context of a transaction associated with the current thread of control. However the transaction needs to be propagated to the target objects to ensure that any updates they make are done in the context of the application's transaction.

POA Policies

To ensure propagation of transaction contexts the target objects must be placed in a POA with specific OTS POA policies. In particular the POA must use one of the OTSPolicy values `REQUIRES` or `ADAPTS`. The following code shows the creation of a POA with the `REQUIRES` OTSPolicy and the activation of an account object in the POA.

```
// C++

CORBA::ORB_var orb = ...

// Create a policy object for the REQUIRES OTS Policy.
CORBA::Any policy_val;
policy_val <<= CosTransactions::REQUIRES;
CORBA::Policy_var tx_policy =
    orb->create_policy(CosTransactions::OTS_POLICY_TYPE,
                     policy_val);

// Add OTS policy to policy list (just 1 policy in this case).
CORBA::PolicyList policies(1);
policies.length(1);
policies[0] = CORBA::Policy::_duplicate(tx_policy);

// Get a reference to the root POA.
CORBA::Object_var obj =
    orb->resolve_initial_references("RootPOA");
PortableServer::POA_var root_poa =
    PortableServer::POA::_narrow(obj);

// Create a new POA with the OTS Policy.
PortableServer::POA_var POA_tx_poa =
    root_poa->create_POA("REQUIRES TX",
                       root_poa->the_POAManager(),
                       policies);

// Create object using the transactional POA. This example
// uses servant_to_reference() to create the object
//
// AccountImpl is the servant class implementing the
// IDL interface Account.
AccountImpl* servant = new AccountImpl(...);
PortableService::ObjectId_var id =
    tx_poa->activate_object(servant);
obj = tx_poa->servant_to_reference(servant);
Bank::Account_var account = Bank::Account::_narrow(obj);
```

OTSPolicy values

There are three OTSPolicy values: `REQUIRES`, `ADAPTS` and `FORBIDS`. `REQUIRES` specifies that the object must be invoked within a transaction; `ADAPTS` allows the object to be invoked both within and without a transaction; `FORBIDS` specifies that the object must not be invoked within a transaction. See [“Propagation and Transaction Policies”](#) for a full discussion of POA and client policies relating to transaction propagation. Support for the deprecated

`TransactionalObject` interface is discussed in [“Use of the TransactionalObject Interface”](#) on page 79.

XA Resource Manager Integration

Process of using an XA Resource Manager

Integrating an XA compliant resource manager with OTS managed transactions involves three steps:

1. Setting up configuration variables for the resource manager.
2. Application initialization.
3. Accessing the database during an OTS transaction.

Full details are in [“Using XA Resource Managers with OTS”](#).

Resource Manager Configuration

Each resource manager used by an application requires configuration. The configuration is placed in a namespace that is passed to the `create_resource_manager()` operation during application initialization. The minimum configuration is the specification of the resource manager’s open-string. This is a resource manager specific string that is passed to the `xa_open()` call and contains sufficient information to create an XA connection to the database. For example this can contain user name and password details.

The following example shows the configuration for an Oracle database using the `xa_resource_managers:oracle` namespace. The `thread_model` configuration variable specifies scope of an XA connection (either thread or process):

```
xa_resource_managers:oracle:thread_model = "PROCESS";  
xa_resource_managers:oracle:open_string =  
    "Oracle_XA+Acc=P/SCOTT/TIGER+SesTm=60+SqlNet=osol"
```

Application Initialization

Applications using XA resource managers must include the file `omg/XA.hh` to access the interfaces in the XA module. During application initialization `ResourceManager` and `CurrentConnection` objects are created to represent the resource manager being integrated. This is done by getting a reference to the Connector

object (by passing "XAConnector" to `resolve_initial_references()`) and calling `create_resource_manager()`:

```
// C++
...
CORBA::ORB_var orb = ...

// Get reference to the XAConnector object.
CORBA::Object_var xa_connector_obj =
    orb->resolve_initial_references("XAConnector");
XA::Connector_var xa_connector =
    XA::Connector::_narrow(xa_connector_obj);

// Get XA Connection object for the resource manager.
XA::CurrentConnection_var current_connection;
XA::ResourceManager_var rm =
    xa_connector->create_resource_manager(
        "xa_resource_managers:oracle",
        xaosw, "",
        current_connection);
```

The `create_resource_manager()` operation is passed the resource manager's name, XA switch (xaosw is Oracle's XA switch), open-string and close string as well as flags that affect the behavior of the resource manager. It returns a reference to the `ResourceManager` object and a reference to the `CurrentConnection` object (as an out parameter).

Accessing the Database within an OTS Transaction

The application code used to read and write to the database is the same as for a normal application with the following exceptions:

1. Before each access to the database the `start()` operation must be called on the XA Connection object to associate the connection with the current transaction.
2. After the database access the `end()` operation must be called on the XA Connection object to remove the association with the current transaction.
3. Resource manager operations related to transaction management such as the embedded SQL operations `BEGIN`, `COMMIT`, or `ROLLBACK` must not be used.

The following shows how integration with an XA-compliant database is achieved using embedded SQL:

```
// C++
void AccountImpl::deposit(float amt)
{
    // Get the coordinator and otid for the current transaction.
    CosTransactions::Current_var tx_current = ...
    CosTransactions::Control_var control =
        tx_current->get_control();
    CosTransactions::Coordinator_var tx =
        control->get_coordinator();
    CosTransactions::PropagationContext_var ctx =
        tx->get_txcontext();
    const CosTransactions::otid_t& otid = ctx->current.otid;

    // Associate current transaction with the XA connection
    // to the database.
    XA::CurrentConnection_var current_connection = ...
    current_connection->start(tx, otid);

    EXEC SQL BEGIN DECLARE SECTION
        unsigned long acc_id = m_accId;
        float          balance = 0.0;
    EXEC SQL END DECLARE SECTION

    // Get the current balance.
    EXEC SQL SELECT BALANCE
        INTO :balance
        FROM ACCOUNTS
        WHERE ACC_ID = :acc_id;

    // Update balance.
    balance += amt;
    EXEC SQL UPDATE ACCOUNTS
        SET BALANCE = :balance
        WHERE ACC_ID = :acc_id;

    // Dissociate the current transaction from the XA
    // connection to the database.
    current_connection->end(tx, otid, IT_true);
}
```

Application-Specific Resources

Resource interface operations

The `CosTransactions::Resource` interface provides a mechanism for applications to become involved in the commit and rollback protocol of a transaction. The `Resource` interface provides five operations that are called at key points during the commit or rollback protocols:

- `prepare()`
- `commit()`
- `rollback()`
- `commit_one_phase()`
- `forget()`

Implementing resource objects

An application implements a resource object that supports the `Resource` interface and registers an instance of the object with a transaction using the `register_resource()` operation provided by the `Coordinator` interface. Resource object implementations are responsible for cooperating with the OTS to ensure the ACID properties for the whole transaction. In particular resource objects must be able to recover from failures.

The implementation of resource objects is discussed in detail in ["Writing Recoverable Resources"](#).

Configuration Issues

Issues

Before an application using OTS can run there are a number of configuration issues. These are concerned with loading the appropriate plug-ins and setting up the client and server bindings to enable implicit propagation of transactions.

Loading the OTS plug-in

For server applications, the OTS plug-in must be loaded explicitly by including it in the `orb_plugins` configuration variable. For example:

```
orb_plugins = [ ..., "ots"];
```

The client and server bindings are controlled with the configuration variables `binding:client_binding_list` and `binding:server_binding_list` respectively. The settings for both variables need to take account of the OTS for potential bindings. For example, to be considered for the IIOP/GIOP and collocated-POA bindings the variables must be set as follows:

```
binding:client_binding_list = ["OTS+POA_Coloc",  
                              "OTS+GIOP+IIOP",  
                              "POA_Coloc",  
                              "GIOP+IIOP"];
```

```
binding:server_binding_list = ["OTS", ""];
```

Other configuration variables can be used to alter the characteristics of your application. These are covered in ["Configuration Summary"](#).

Transaction Demarcation and Control

The most convenient means of demarcating transactions is to use the OTS Current object. Direct transaction demarcation using the TransactionFactory and Terminator interfaces provide more flexibility but is more difficult to manage.

The OTS Current Object

Current Interface

The OTS Current object maintains associations between the current thread of control and transactions. The Current interface is defined as follows:

```
// IDL (in module CosTransactions)
local interface Current : CORBA::Current {

    void begin()
        raises(SubtransactionsUnavailable);

    void commit(in boolean report_heuristics)
        raises(NoTransaction, HeuristicMixed,
              HeuristicHazard);

    void rollback()
        raises(NoTransaction);

    void rollback_only()
        raises(NoTransaction);

    Status get_status();

    string get_transaction_name();

    void set_timeout(in unsigned long seconds);
    unsigned long get_timeout();

    Control get_control();

    Control suspend();

    void resume(in Control which)
        raises(InvalidControl);
};
```

Threads and transactions

The OTS `Current` object maintains the association between threads and transactions. This means the same OTS `Current` object can be used by several threads. Figure 3 shows the relationship between threads, the OTS `Current` object, and the three objects that represent a transaction (`Control`, `Coordinator` and `Terminator`).




Figure 3: *Thread and Transaction Associations*

Getting a Reference to the OTS Current Object

A reference to the OTS `Current` object is obtained by calling `resolve_initial_references()` passing "TransactionCurrent" as the parameter and narrowing the result to `CosTransactions::Current`. For example:

```
// C++
CosTransactions::Current_var tx_current;
try {
    CORBA::ORB_var orb = ...
    CORBA::Object_var obj =

        orb->resolve_initial_references("TransactionCurrent");

    tx_current = CosTransactions::Current::_narrow(obj);
}
catch (CORBA::SystemException& ex)
{
    // Error handling.
    ...
}
```

The `Current` interface is declared as local which means references to the `Current` object cannot be passed as parameters to IDL operations or passed to operations such as `object_to_string()`.

Creating Transactions

The `begin()` operation is used to create a new transaction and associate the new transaction with the current thread of control. If there is no current transaction a top-level transaction is created; otherwise a nested transaction is created (see ["Nested Transactions" on page 25](#)).

The following code creates a new transaction:

```
// C++
CosTransactions::Current_var tx_current = ...
try
{
    tx_current->begin();
}
catch (CosTransactions::SubtransactionsUnavailable& ex)
{
    // Already in a transaction and nested transaction are
    // not
    // supported.
}
catch (CORBA::SystemException& ex)
{
    // Error handling...
}
```

Committing the Current Transaction

The `commit()` operation attempts to commit the current transaction, if any, and removes the current thread/transaction association. If the `commit()` operation returns normally the transaction was successfully committed. However if the `TRANSACTION_ROLLEDBACK` system exception is raised the transaction has been rolled back. In both cases the transaction is disassociated with the current thread of control.

For example, the following code attempts to commit the current transaction:

```
// C++
CosTransactions::Current_var tx_current = ...
try
{
    // Attempt to commit the current transaction.
    tx_current->commit(IT_false);
}
catch (CORBA::TRANSACTION_ROLLEDBACK&)
{
    // The transaction was rolled back.
}
catch (CORBA::SystemException& ex)
{
    // Error handling...
}
catch (CosTransactions::NoTransaction& ex)
{
    // There was no transaction to commit.
}
```

If there is no current transaction the `CosTransactions::NoTransaction` exception is raised.

The `commit()` operation takes a boolean parameter that indicates whether reporting of heuristic exceptions is permitted. Heuristic exceptions occur when there is a conflict or potential conflict between the outcome decided by the transaction coordinator and the outcome performed by one or more resource managers (see [“Heuristic Outcomes” on page 71](#) for more details). If a value of `true` is passed, the application must be prepared to catch the `HeuristicMixed` and `HeuristicHazard` exceptions; if a value of `false` is passed these exceptions are never raised.

Rolling Back the Current Transaction

The `rollback()` operation rolls back the current transaction, if any, and removes the current thread/transaction association. For example, the following code rolls back the current transaction:

```
// C++
CosTransactions::Current_var tx_current = ...
try
{
    tx_current->rollback();
}
catch (CORBA::SystemException& ex)
{
    // Error handling...
}
catch (CosTransactions::NoTransaction& ex)
{
    // There was no transaction to commit.
}
```

If there is no current transaction the `CosTransactions::NoTransaction` exception is raised.

The `rollback_only()` operation may also be used to mark a transaction to be rolled back. This operation does not actively rollback the transaction, but instead prevents it from ever being committed. This can be useful, for example, to ensure the current transaction will be rolled back during a remote operation. Again, the `NoTransaction` exception is raised if there is no current transaction.

Nested Transactions

Nested transactions, also known as sub-transactions, provide a way of composing applications from a set of transactions each of which can fail independently of each other. Nested transactions form a hierarchy known as a transaction family. No updates are made permanent until the top-level transaction commits.

When using the `Current` object, a nested transaction is created by calling `begin()` when there is already a transaction associated with the current thread of control. When nested transaction is committed or rolled back, the thread transaction association reverts back to the parent transaction.

Note:

Nested transactions are not supported in this release of Orbix.

Timeouts

The `set_timeout()` operation sets the timeout in seconds for subsequent top-level transactions. It does not set the timeout for the current transaction. Passing a value of 0 means subsequent top-level transactions will never timeout.

If `set_timeout()` is not called the default timeout is taken from the `plugins:ots:default_transaction_timeout` configuration variable.

The `get_timeout()` operation returns the current timeout in seconds for subsequent top-level transactions. It does not return the timeout for the current transaction.

For example, the following code sets the timeout for subsequent top level transactions to 30 seconds:

```
// C++
CosTransactions::Current_var tx_current = ...
tx_current->set_timeout(30);
```

Suspending and Resuming Transactions

The `suspend()` operation temporarily removes the association between the current thread of control and the current transaction if any. Calling `suspend()` returns a reference to a control object for the transaction. The transaction can be resumed later by calling the `resume()` operation passing in the reference to the control object.

Suspending a transaction is useful if it is necessary to perform work outside of the current transaction. For example:

```
// C++
CosTransactions::Current_var tx_current = ...
tx_current->begin();
account->deposit(100.0);

// Suspend the current transaction.
CosTransactions::Control_var control =
    tx_current->suspend();

// Do some non-transactional work.
...

// Resume the transaction.
tx_current->resume(control);

tx_current->commit(IT_true);
```

The `resume()` operation raises the `CosTransactions::InvalidControl` exception if the transaction represented by the control object cannot be resumed.

Sometimes the work done during the transaction's suspend state can be work on a different transaction. Thus, `suspend()` and `resume()` give you a way to work on multiple transactions within the same thread of control.

Miscellaneous Operations

The `get_status()` and `get_transaction_name()` operations provide information on the current transaction. The `get_control()` operation returns the `Control` object for the current transaction or `nil` if there is no current transaction. This is used to provide access to the `Coordinator` and `Terminator` objects for more advanced control. See "[Transaction Management](#)" for more details

Direct Transaction Demarcation

Using the transaction factory to create transactions

The alternative to using the `OTS Current` object is to use the transaction factory directly to create transactions.

Example

The following code shows the creation of a new top-level transaction:

```
// C++
//
// Get a reference to the transaction factory.
CORBA::ORB_var orb = ...
CORBA::Object_var obj =
    orb->resolve_initial_references("TransactionFactory");
CosTransactions::TransactionFactory_var tx_factory =
    CosTransactions::TransactionFactory::_narrow(obj);

// Create a transaction with a timeout of 60 seconds.
CosTransactions::Control_var control =
    tx_factory->create(60);
```

The first step is to obtain a reference to the transaction factory object. This is done by calling `resolve_initial_references()` passing a value of "TransactionFactory" and narrowing the result to `CosTransactions::TransactionFactory`.

The `create()` operation creates a new top-level transaction and returns a control object representing the new transaction. `create()` is passed the timeout in seconds for the transaction. A value of 0 means there is no timeout.

To complete a transaction created using the transaction factory, the terminator object is used. The terminator object is obtained by calling `get_terminator()` on the control object. The `Terminator` interface provides the `commit()` and `rollback()` operations. These are the same as the ones provided by the `Current` interface except they do not raise the `NoTransaction` exception.

Example of a commit

The following shows the attempted commit of a transaction using the direct approach:

```
// C++
//
try {
    CosTransactions::Terminator_var term =
        control->get_terminator();
    term->commit(IT_true);
} catch (CORBA::TRANSACTION_ROLLEDBACK&) {
    // Transaction has been rolled back.
}
```


Propagation and Transaction Policies

This chapter describes how to control transfer of the transaction to the target object using POA policies or explicitly.

Implicit Propagation Policies

Implicit and Explicit Propagation

Propagation refers to the transfer of the transaction to the target object during an invocation.

For transactions created using the OTS Current object , propagation is implicit. That is, the application does not have to change the way the object is invoked in order for the transaction to be propagated. Implicit propagation is controlled using POA policies.

For transactions created directly via the TransactionFactory reference, explicit propagation must be used.

Policies for implicit propagation

For implicit propagation, there are two POA policies and one client policy that affect the behavior of invocations with respect to transactions.

The POA policies are:

- OTSPolicy
- InvocationPolicy

Both policies allow an object to set requirements on whether the object is invoked in the context of a transaction and transaction model being used.

The client OTS policy is:

- NonTxTargetPolicy

This alters the client's behavior when invoking on objects that do not permit transactions.

Note: These three policies replace the deprecated TransactionPolicy and the use of the deprecated TransactionalObject interface both of which are still supported in this release. See ["Migrating from TransactionPolicies" on page 40](#) and ["Use of the TransactionalObject Interface" on page 79](#) for more details.

Shared and Unshared Transactions

InvocationPolicy transaction models

The InvocationPolicy deals with the transaction model supported by the target object. There are two transaction models:

- shared
- unshared

Shared model

The shared model is the familiar end-to-end transaction where the client and the target object both share the same transaction. That is, an invocation on an object within a shared transaction is performed within the context of the transaction associated with the client.

Unshared model

An unshared transaction is used for asynchronous messaging where different transactions are used along the invocation path between the client and the target object. Here, the target object invocation is performed within the context of a different transaction than the transaction associated with the client. Hence, the client and target object does not share the same transaction. This model is required since with asynchronous messaging it is not guaranteed that the client and server are active at the same time.

Orbix does not support unshared transactions in this release. They are included in the following discussion for completeness only.

Policy Meanings

The three standard OTSPolicy values

The OTSPolicy has three possible standard values plus additional two values specific to Orbix. The Orbix-specific values are discussed in ["Orbix-Specific OTSPolicies" on page 37](#); the standard values and their meanings are:

REQUIRES This policy is used when the target object always expects to be invoked within the context of a transaction. If there is no transaction the `TRANSACTION_REQUIRED` system exception is raised. This policy guarantees that the target object is always invoked within a transaction.

FORBIDS	This policy is used when the target object does not permit invocations performed within the context of a transaction. If a transaction is present the <code>INVALID_TRANSACTION</code> system exception is raised. This policy guarantees that the target object is never invoked within a transaction. This is the default policy.
ADAPTS	This policy is used when the target object can accept both the presence and absence of a transaction. If the client is associated with a transaction, the target object is invoked in the context of the transaction; otherwise the target object is invoked without a transaction. This policy guarantees that the target object is invoked regardless of whether there is a transaction or not. Here, the target object adapts to the presence or not of a transaction.

Objects with the `REQUIRES` or `ADAPTS` `OTSPolicy` are also known as transactional objects since they support invocations within transactions; objects with the `FORBIDS` `OTSPolicy` or no `OTSPolicy` at all are known as non-transactional objects since they do not support invocations within transactions.

For an example of using an `OTSPolicy` see [“Example Use of an `OTSPolicy`” on page 32](#) below.

The two `NonTxTargetPolicy` values

The default behavior for a client that invokes on an object within a transaction where the target object has the `FORBIDS` `OTSPolicy` (or where the object does not have any `OTSPolicy`, since `FORBIDS` is the default) is for the `INVALID_TRANSACTION` exception to be raised. This behavior can be altered with the `NonTxTargetPolicy`. The policy values and their meanings are:

PREVENT	The invocation is prevented from proceeding and the <code>INVALID_TRANSACTION</code> system exception is raised. This is the default behavior
PERMIT	The invocation proceeds but the target object is not invoked within the context of the transaction. This satisfies the target object's requirements and allows the client to make invocations on non-transactional objects within a transaction.

Setting the policies

As with all client policies, there are four ways in which they may be set:

1. Using configuration. For the `NonTxTargetPolicy` the variable to set is `policies:non_tx_target_policy`.
2. Set the policy on the ORB using the `CORBA::PolicyManager` interface.

3. Set the policy for the current invocation using the `CORBA::PolicyCurrent` interface.
4. Set the policy on the target object using the `CORBA::Object::_set_policy_overrides()` operation.

For more information on client policies see the chapter “Using Policies” in the *CORBA Programmer’s Guide*. For an example of using a `NonTxTargetPolicy` see [“Example Use of a NonTxTargetPolicy” on page 34](#) below.

Note that since the default `OTSPolicy` is `FORBIDS`, using the `PREVENT NonTxTargetPolicy` could result in previously working code becoming unworkable due to invocations been denied. The `PREVENT` policy should be used with care.

The three InvocationPolicy values

Finally, the choice of which transaction model (shared or unshared) that an object supports is done using the `InvocationPolicy`. This has three values:

SHARED	The target object supports only shared transactions. This is the default. An asynchronous invocation results in the <code>TRANSACTION_MODE</code> system exception being raised.
UNSHARED	The target object supports only unshared transactions. A synchronous invocation results in the <code>TRANSACTION_MODE</code> system exception begin raised.
EITHER	The target object supports both shared and unshared transactions.

Note that the `UNSHARED` and `EITHER` `InvocationPolicies` cannot be used in combination with the `FORBIDS` and `ADAPTS` `OTSPolicies`. Attempting to create a POA with these policy combinations results in the `PortableServer::InvalidPolicy` exception being raised.

Example Use of an OTSPolicy

Steps to create an object with an OTSPolicy

The following are the steps to create an object with a particular `OTS` policy:

1. Create a `CORBA Policy` object that represents the desired `OTS` policy. This is done by calling the `ORB` operation `create_policy()` passing in the value `CosTransactions::OTS_POLICY_VALUE` as the first parameter and the policy value (encoded as an `any`) as the second parameter.
2. Create a POA that includes the `OTSPolicy` in its policy list. This is done by calling `create_POA()`.
3. Create an object using the new POA.

Example

The following code sample shows an object being created in a POA that uses the `ADAPTS` `OTSPolicy`. For clarity, the POA is created off the root POA and only one new policy is added.

```
// C++
//
// Create CORBA policy object for ADAPTS OTSPolicy
CORBA::Any tx_policy_value;
tx_policy_value <<= CosTransactions::ADAPTS;

CORBA::ORB_var orb = ...
CORBA::Policy_var tx_policy = orb->create_policy(
    CosTransactions::OTS_POLICY_TYPE, tx_policy_value);

// Create a POA using the transactional policy.
CORBA::PolicyList policies(1);
policies.length(1);
policies[0] = CORBA::Policy::_duplicate(tx_policy)

// Get a reference to the root POA.
CORBA::Object_var obj =
orb->resolve_initial_references("RootPOA");
PortableServer::POA_var root_poa =
    PortableServer::_narrow(obj);

// Set up nil POAManager reference.
PortableServer::POAManager_var nil_mgr =
PortableServer::POAManager::_nil();

PortableServer::POA_var tx_poa =
root_poa->create_POA("TX ADAPTS", nil_mgr, policies);

// Create object using the transactional POA. This example
// uses servant_to_reference() to create the object

// AccountImpl is the servant class implementing the
// IDL interface Account.
AccountImpl* servant = new AccountImpl(...);

PortableServer::ObjectId_var id =
    tx_poa->activate_object(servant);

obj = tx_poa->servant_to_reference(servant);
Account_var account = Account::_narrow(obj);
```

Example Use of a NonTxTargetPolicy

Steps to use a NonTxTargetPolicy

The following are the steps for a client to use a `NonTxTargetPolicy` when invoking on a non-transactional object:

1. Get a reference to the `PolicyCurrent` or `PolicyManager` object.
2. Create a CORBA `Policy` object that represents the desired `NonTxTargetPolicy`. This is done by calling the `CORBA::ORB::create_policy()` operation passing in the value `CosTransactions::NON_TX_TARGET_POLICY_TYPE` as the first parameter and the policy value (encoded as an `any`) as the second parameter.
3. Call the `set_policy_overrides()` operation on the `PolicyCurrent` or `PolicyManager` object passing in a policy list containing the `NonTxTargetPolicy`. Alternatively call the `_set_policy_overrides()` operation on the target object itself.
4. Invoke on the non-transaction object (from within a transaction).

Example

The following code shows a client using the `PERMIT` `NonTxTargetPolicy` to invoke on a non-transactional object within a transaction. The client uses the `PolicyCurrent` object to set the policy. Assume that the `Account` object is using the `REQUIRES` or `ADAPTS` `OTSPolicy` and the `AuditLog` object is using the `FORBIDS` `OTSPolicy` or no `OTSPolicy` at all:

```

// C++
//
// Get reference to PolicyCurrent object.
CORBA::ORB_var orb = ...
CORBA::Object_var obj =
    orb->resolve_initial_references("PolicyCurrent");

CORBA::PolicyCurrent_var policy_current =
    CORBA::PolicyCurrent::_narrow(obj);

// Create PERMIT NonTxTarget policy.
CORBA::PolicyList policy_list(1);
policy_list.length(1);

CORBA::Any tx_policy_value;
tx_policy_value <<= CosTransactions::PERMIT;

policy_list[0] = orb->create_policy(
    CosTransactions::NON_TX_TARGET_POLICY_TYPE,
    tx_policy_value);

// Set policy overrides.
policy_current->set_policy_overrides(policy_list,
    CORBA::ADD_OVERRIDE);

// Invoke on target object
CosTransactions::Current_var tx_current = ...
Account_var account = ...
AuditLog_var log = ...

tx_current->begin();
account->deposit(100.00);
log->append("User ... deposited 100 to account ...");
tx_current->commit(IT_true);

```

Specifying the default NonTxTargetPolicy

The default `NonTxTargetPolicy` value is taken from the `policies:non_tx_target_policy` configuration variable, which can be set to "prevent" and "permit" to represent the `PREVENT` and `PERMIT` policy values. If this configuration variable is not set, the default is `PREVENT`.

Use of the ADAPTS OTSPolicy

Using the ADAPTS OTSPolicy

The ADAPTS OTSPolicy is useful for implementing services that must work whether or not the client is using OTS transactions. If the client is using transactions, the target object simply executes in the same transaction context and its work will be either committed or rolled back when the client completes the transaction.

However, if there is no transaction the target object can choose to create a local transaction for the duration of the invocation.

Example

The following code shows how a servant might be implemented to take advantage of the ADAPTS OTSPolicy (error handling has been omitted):

```
// C++
void AccountImpl::deposit(float amount)
{
    CosTransactions::Current_var tx_current = ...

    // Test if a transaction was propagated from the client.
    CosTransactions::Control_var control =
        tx_current->get_control();

    if (CORBA::is_nil(control))
    {
        // No current transaction, so create one.
        tx_current->begin();
    }

    // Do the transactional work
    ...

    // If a local transaction was created, commit it.
    if (CORBA::is_nil(control))
    {
        tx_current->commit(IT_true);
    }
}
```

This approach allows clients to selectively bracket operations with transactions based on how much work is done. For example, if only a single server operation is performed then no client transaction needs to be created. However, if more than one operation is performed the client creates a transaction to ensure ACID properties for all of the operations.

For example (error handling omitted):

```
// C++
// Deposit money into a single account (no transaction
// needed).
Account_var acc = ...
acc->deposit(100.00);

// Transfer money between two account (this requires a
// transaction)
Account_var src_acc = ...
Account_var dest_acc = ...
CosTransactions::Current_var tx_current = ...

tx_current->begin();
src_acc->withdraw(200.00);
dest_acc->deposit(200.00);
tx_current->commit(IT_true);
```

For this example the servant created an OTS transaction. However, it could just create a local database transaction instead or not create any transaction at all.

Orbix-Specific OTSPolicies

The two proprietary OTSPolicy values

Orbix extends the set of OTSPolicies with two proprietary values to support automatically created transactions and optimizations. The values and their meanings are:

- | | |
|-------------|---|
| AUTOMATIC | This policy is used when the target object always expects to be invoked within the context of a transaction. If there is no transaction a transaction is created for the duration of the invocation. This policy guarantees that the target object is always invoked within a transaction. See "Automatic Transactions" . |
| SERVER_SIDE | This policy is used in conjunction with just-in-time transaction creation to optimize the number of network messages in special cases. See "Just-In-Time Transaction Creation" . |

Automatic Transactions

The ADAPTS OTSPolicy (see ["Use of the ADAPTS OTSPolicy" on page 36](#)) is useful for implementing servants that can be invoked both with and without transactions. A useful pattern to use is for the servant to check for the existence of a transaction and create one for the duration of the invocation if there is none. The AUTOMATIC OTSPolicy provides this functionality without having to code it into the servant implementation.

From the target object's point of view the `AUTOMATIC` OTSPolicy is the same as `REQUIRES` since the target object is always invoked in the context of a transaction. However, from the clients point of view, the `AUTOMATIC` policy is the same as `ADAPTS` since the client can choose whether to invoke on the object within a transaction or not. In fact, object references created in a POA with the `AUTOMATIC` OTSPolicy contain the `ADAPTS` policy so they can be used by other OTS implementations that do not support the `AUTOMATIC` OTSPolicy.

For the case were the client does not use a transaction and the automatically created transaction fails to commit, the standard `TRANSACTION_ROLLEDBACK` system exception is raised. Reporting of heuristic exceptions is not supported.

Just-In-Time Transaction Creation

Orbix provides three extensions to support the concept of just-in-time (JIT) transaction creation to eliminate network messages in special conditions. These extensions are:

1. A configuration option to enable JIT transaction creation, which allows the creation of a transaction to be delayed until it is really needed.
2. The `SERVER_SIDE` OTSPolicy which allows a transaction to be created just before a target object is invoked.
3. A additional operation `commit_on_completion_of_next_call()` that allows the next invocation on an object to also commit the transaction.

The use of JIT transaction creation is useful when invocations between a client and an object involve using a network connection. This is because it can reduce the number of network messages that are exchanged to create, propagate and commit a transaction.

Enabling JIT Transaction Creation

JIT transaction creation is enabled by setting the `plugins:ots:jit_transactions` configuration variable to "true". When enabled a call to `Current::begin()` does not create a transaction; instead, it remembers that the client requested to create one. The client is said to be in the context of an empty transaction. At this stage a call to `Current::get_status()` would return `StatusActive` event though a real transaction has not been created. Likewise, calls to `Current::commit()` and `Current::rollback()` would succeed. A real transaction is only created at the following points:

1. When any of the following `CosTransactions::Current` operations are invoked: `rollback_only()`, `get_control()`, `get_transaction_name()` or `suspend()`.
2. When an object with any of the standard OTSPolicies is invoked.

If the target object's OTSPolicy is `SERVER_SIDE`, a real transaction is not created until the invocation has reached the object's POA. Note that unlike the `AUTOMATIC` OTSPolicy, this transaction is not terminated when the invocation has completed. Instead, the client adopts the newly created transaction.

When JIT transactions are not enabled, the `SERVER_SIDE` OTSPolicy behaves the same as the `ADAPTS` OTSPolicy, except that unlike the `AUTOMATIC` policy, other OTS implementations will not recognize the new policy.

A final optimization is possible when JIT transaction creation and the `SERVER_SIDE` OTSPolicy are used. The OTS current object in Orbix provides an additional operation that allows a transaction to be committed within the context of the target object rather than by the client:

```
// IDL
module IT_CosTransactions
{
    interface Current : CosTransactions::Current
    {
        void
        commit_on_completion_of_next_call()
            raises (CosTransactions::NoTransaction)
    };
};
```

The `commit_on_completion_of_next_call()` operation causes the current transaction to be committed after the completion of the next object invocation (so long as the target object is using the `SERVER_SIDE` OTSPolicy). The transaction commit is performed by the target object's POA, which means that the transaction will have been created and committed in the context of the target object rather than by the client.

To use the operation the client must include the file `<orbix/cos_transactions.hh>` and narrow the OTS current object to the `IT_CosTransactions::Current` interface.

```
// C++
CosTransactions::Current_var tx_current = ...

IT_CosTransactions::Current_var it_tx_current =
    IT_CosTransactions::Current::_narrow(tx_current);

Account_var account = ...
it_tx_current->begin();

account->deposit(100.00);

it_tx_current->commit_on_completion_of_next_call();
account->deposit(50.00);

it_tx_current->commit(IT_true);
```

Note that the client still must call the `commit()` operation, though this will not result in any network messages.

Migrating from TransactionPolicies

Mapping from TransactionPolicy values

Previous releases of Orbix used the deprecated `CosTransaction::TransactionPolicy` which provided seven standard policy values and two Orbix extensions. Below is a table that provides the mapping from `TransactionPolicy` values to their `OTSPolicy` and `InvocationPolicy` equivalent.

Table 3: *Mapping from TransactionPolicy values*

TransactionPolicy Value	OTSPolicy Value	InvocationPolicy Value
Allows_shared	ADAPTS	SHARED
Allows_none	FORBIDS	SHARED
Requires_shared	REQUIRES	SHARED
Allows_unshared	ADAPTS	Not supported
Allows_either	ADAPTS	Not supported
Requires_unshared	REQUIRES	UNSHARED
Requires_either	REQUIRES	EITHER or none
Automatic_shared	AUTOMATIC	SHARED
Server_side_shared	SERVER_SIDE	SHARED

Combining Policy Types

It is possible to create a POA that combines all three policy types to support interoperability with earlier versions of Orbix. However, invalid combinations result in the `PortableServer::InvalidPolicy` exception being raised when `PortableServer::POA::create_POA()` is called. An invalid combination is any combination not in [Table 3](#); for example combining `Requires_shared` with `ADAPTS` and `SHARED`.

The mappings for the `Allows_unshared` and `Allows_either` `TransactionPolicies` are not supported since this would lead to an invalid combination of `OTSPolicies` and `InvocationPolicies`.

Note: Support for the `TransactionPolicy` type may be discontinued in a future Orbix release. It is recommended that only `OTSPolicies` and `InvocationPolicies` be used.

Explicit Propagation

Altering the IDL to propagate explicitly

When a transaction is created directly using the `TransactionFactory` interface the transaction must be propagated explicitly to target objects. This means altering the IDL for the application to add an extra parameter for the transaction's Control object.

Example

The following is the `Account` IDL interface modified to support explicit propagation:

```
// IDL (in module Bank)
#include <CosTransactions.idl>
...
interface Account
{
    exception InsufficientFunds {};

    void deposit(in CashAmount amt,
                in CosTransactions::Control ctrl);

    void withdraw(in CashAmount amt,
                 in CosTransactions::Control ctrl)
        raises (InsufficientFunds);
};
```

Each invocation on the account object must now take a reference to a transaction control as its last parameter:

```
// C++
CosTransactions::TransactionFactory_var tx_factory = ...
CosTransactions::Control_var control =
    tx_factory->create(60);

Bank::Account_var src_acc = ...
Bank::Account_var dest_acc = ...
Bank::CashAmount amount = 100.0;
src_acc->withdraw(amount, control);
dest_acc->deposit(amount, control);

CosTransactions::Terminator_var term =
    control->get_terminator();
term->commit(IT_true);
```

It is also possible to pass a reference to the transaction's coordinator object instead of its control object.

Using XA Resource Managers with OTS

This chapter describes how to integrate with transactional systems by implementing `CosTransactions::Resource` objects on top of the standard X/Open XA interface.

The XA Interface

Resource objects

To use a transactional system (such as a database system) with the transaction service, you must "connect" the transactions provided by the transactional system to the distributed transactions managed by the transaction service. With the transaction service, this is achieved by implementing `CosTransactions::Resource` objects — each resource represents a local transaction in the transactional system — and registering these Resource objects with the distributed transactions.

Because many systems provide a standard interface to their transactional capabilities — the X/Open XA interface — you can implement `CosTransactions::Resource` objects on top of the XA interface, and provide an easy-to-use integration with the transaction service. This is precisely what the Orbix XA plug-in provides.

XA Overview

XA (X/Open CAE Specification, Distributed Transaction Processing: The XA specification, December 1991, ISBN: 1 872630 24 3) specifies a standard C API provided by transactional systems (called Resource Managers in the XA specification) that want to participate in distributed transactions managed by transaction

managers developed by other vendors. XA defines a set of C-function pointers, and a C-struct that holds these function pointers, `xa_switch_t` (see `orbix_sys/xa.h`):

```
struct xa_switch_t
{
    char name[RMNAMESZ]; /* name of resource manager */
    long flags; /* resource manager specific options */
    long version; /* must be 0 */
    int (*xa_open_entry) /* xa_open function pointer */
    (char *, int, long);
    int (*xa_close_entry) /* xa_close function pointer */
    (char *, int, long);
    int (*xa_start_entry) /* xa_start function pointer */
    (XID *, int, long);
    int (*xa_end_entry) /* xa_end function pointer */
    (XID *, int, long);
    int (*xa_rollback_entry) /* xa_rollback function
pointer */
    (XID *, int, long);
    int (*xa_prepare_entry) /* xa_prepare function pointer
*/
    (XID *, int, long);
    int (*xa_commit_entry) /* xa_commit function pointer
*/
    (XID *, int, long);
    int (*xa_recover_entry) /* xa_recover function pointer
*/
    (XID *, long, int, long);
    int (*xa_forget_entry) /* xa_forget function pointer
*/
    (XID *, int, long);
    int (*xa_complete_entry) /* xa_complete function
pointer */
    (int *, int *, int, long);
};
```

Function pointers

Each XA Resource Manager must provide a global instance of `xa_switch_t`. For example, Oracle's global `xa_switch_t` instance is called `xaosw`.

The function pointers provided by this `xa_switch_t` instances can be divided into four categories:

- Functions to connect and disconnect to the XA Resource Manager: `xa_open()` and `xa_close()`. The string passed to `xa_open()` typically contains connection information, e.g. a database name and a username and password.
- Transaction completion functions `xa_prepare()`, `xa_commit()`, `xa_rollback()`, `xa_forget()` correspond to the `CosTransactions::Resource` operations.
- Recovery function `xa_recover()` is currently not used by the XA plug-in.
- Functions used to start and end associations between connections and a transactions: `xa_start()`, `xa_end()`

In order to use an XA connection to do some work within a distributed transaction, it is necessary to create an association between this connection and the distributed transaction.

`xa_start()` is used to create such an association;
`xa_end(TMSUSPEND)` suspends the association, without releasing the connection; `xa_start(TMRESUME)` resumes a suspended association; `xa_end(TMSUCCESS)` terminates an association with success; and `xa_end(TMFAIL)` terminates an association and marks the transaction rollback-only.

Note: `xa_complete()` is only used for asynchronous XA, an optional part of XA which is not supported by any popular XA implementation.

XA and Multi-Threading

In the XA specification, the scope of an XA connection is called "thread-of-control". Each thread-of-control can only use the connections that it has established (using `xa_open()`). The XA specification maps thread-of-control to operating system process (2.2.8). Each thread in a process has access to all the XA connections established by this process. This is clearly specified in the JTA specification (XA for Java).

Unfortunately, for the C XA API, most vendors implement the following:

- a thread-unsafe mode, in which the scope of each XA connection is the process (XA thread-of-control maps to process)
- a thread-safe mode, in which the scope of each XA connection is the thread by which it was created (XA thread-of-control maps to thread)

For example, with Oracle, the "+threads={true,false}" option of the OracleXA open string lets the application programmer choose between these two modes. The thread-of-control equal thread model sometimes simplifies the API used to access the data. For example, Oracle embedded SQL in C/C++ (Pro*C/C++) has a notion of a default database connection for each thread of control.

When the model is thread-of-control equal process, and a process has a pool of connections to the same database, it is necessary to explicitly specify which connection to use (with an Oracle AT clause):

```
EXEC SQL AT :db_name INSERT VALUES(123, 43, 3.49) INTO  
    SALE_DETAILS;
```

But when the model is thread-of-control equal thread, and each thread has one connection to a given database, there is no need to explicitly specify the connection to use (no AT clause):

```
EXEC SQL INSERT VALUES(123, 43, 3.49) INTO SALE_DETAILS;
```

The EXEC SQL statements used in a multi-threaded multi-connection application look very much like the EXEC SQL statement used in a single-threaded single-connection application.

The main drawback of tying connection and threads is flexibility since it prevents the application from managing connections independently of threads, which limits the kind of connection pooling that can be implemented. Also, a CORBA server typically dispatches different requests to different threads: the

thread-of-control equal thread model prevents the use of `xa_end(TMSUSPEND)` at the end of a request and `xa_start(TMRESUME)` at the beginning of the next request in the same transaction, since an association must be resumed by the thread of control from which it was suspended.

Using the Orbix XA Plug-In

The Orbix XA plug-in implements and manages `CosTransactions::Resource` objects on behalf of the application. It supports the two thread-of-control models described in the previous paragraph: when the thread model is `XA::PROCESS`, it uses a single-threaded persistent POA to host its `CosTransactions::Resource` servants. When the thread model is `XA::THREAD`, it uses a multi-threaded persistent POA.

You access the XA plug-in by obtaining a reference to the `XA::Connector` local object through `resolve_initial_references()`:

```
#include <omg/xa.hh>
CORBA::Object_var xa_connector_obj =
    orb->resolve_initial_references("XAConnector");
XA::Connector_var xa_connector =
    XA::Connector::_narrow(xa_connector_obj);
```

Then you create an `XA::ResourceManager`, by calling `create_resource_manager` on the connector. This operation creates a persistent POA that hosts the resource manager's servant and will host the `CosTransactions::Resource` servants. The `create_resource_manager` operation also returns an `XA::CurrentConnection` local object, which establishes (with `xa_open()`) connections when needed, and lets you start, suspend, resume, and end associations between any transaction and the current XA thread of control's connection.

```
XA::CurrentConnection_var current_connection;
XA::ResourceManager_var rm =
    xa_connector->create_resource_manager(
        "xa_resource_managers:oracle",
        // the name of an Orbix configuration namespace
        xaosw, // XA switch
        "", // empty open-string, i.e. the unsecured
            // open-string is specified in
        configuration
        "", // empty close-string, i.e. the
        unsecured
            // close-string is specified in the
        configuration
        XA::PROCESS, // thread-model
        false, // no automatic association
        false, // do not use dynamic registration
        current_connection // (out) current connection local
        object
    );
```

The first parameter of `create_resource_manager` is the name of an Orbix configuration namespace; this configuration namespace defines the name of the resource manager persistent POA (defaults to the given namespace name), the open string when

the `open_string` parameter is empty, the close string when the `close_string` parameter is empty, and various other properties. The resource manager id can also be set in the configuration using the `ruid` variable. When the `ruid` variable is set, the XA integration uses the value as the `ruid` passed to `xa_open()` and all subsequent `xa_` calls. When the `ruid` variable is not set, the XA integration generates a new `ruid` value for each `CurrentConnection` object.

Associations between Transactions and Connections

The `CurrentConnection` local interface is defined in the `XA` module as follows:

```
enum ThreadModel { PROCESS, THREAD };
local interface CurrentConnection
{
    void
    start(
        // xa_start(TMNOFLAGS) or xa_start(TMJOIN)
        in CosTransactions::Coordinator tx,
        in CosTransactions::otid_t otid
    );
    void
    suspend(
        // xa_end(TMSUSPEND)
        in CosTransactions::Coordinator tx,
        in CosTransactions::otid_t otid
    );
    void resume(
        // xa_start(TMRESUME)
        in CosTransactions::Coordinator tx,
        in CosTransactions::otid_t otid
    );
    void end(
        // xa_end(TMSUCCESS) or xa_end(TMFAIL)
        in CosTransactions::Coordinator tx,
        in CosTransactions::otid_t otid,
        in boolean success
    );
    ThreadModel thread_model();
    long ruid();
};
```

When the thread model is `PROCESS`, `xa_open()` is called by the first `start` call or the first operation performed by a Resource servant; and `xa_close()` is called during shutdown. When the thread model is `THREAD`, `xa_open()` is called the first time a thread calls `CurrentConnection::start`, or any operation on a Resource servant; `xa_close()` is called when this thread exits.

In order to do some work within a distributed transaction with a given resource manager, you have to associate the resource manager's current connection with this transaction, by calling `CurrentConnection::start`:

```
// assuming the OTS transaction is associated with the
current
// thread
CosTransactions::Control_var control =
    tx_current->get_control();
CosTransactions::Coordinator_var tx =
    control->get_coordinator();
CosTransactions::PropagationContext_var ctx =
    tx->get_txcontext();
const CosTransactions::otid_t& otid = ctx->current.otid;
current_connection->start(tx, otid);
```

The first time `CurrentConnection::start()` is called with a given transaction, the XA plug-in creates a `CosTransactions::Resource` persistent object and registers this object with the transaction coordinator.

Once you have finished using a connection, it is critical to end the association with the transaction for two reasons:

- It releases the connection, and makes it available for other transactions
- As long as any connection is associated with a transaction, this transaction cannot be committed. Some systems (e.g. Oracle) don't even allow to roll back a transaction while it is associated with any connection.

The recommended way to start and end (or start/suspend/resume/suspend...) an association is to use a helper C++ class: the helper class constructor creates the association by calling `start`, and the helper class destructor ends the association. The multi-threaded transfer demo provides a helper `Association` class which uses `start` and `end`; the single-threaded farm demo provides a helper `Association` class which uses `start`, `suspend` and `resume`.

Association State Diagram

Figure 4 shows the state diagram of an association between a transaction and an XA connection. In this diagram all `start`, `suspend`, `resume`, and `end` calls are successful (they do not raise any exception). When `start`, `suspend`, `resume` or `end` raises `CORBA::INTERNAL` with the minor code `IT_XA_MinorCodes::INTERNAL::XAER_RMFAIL`, the new state is "non existant". When `resume`, `suspend` or `end` raises `CORBA::TRANSACTION_ROLLEDBACK` with the minor code `IT_XA_MinorCodes::TRANSACTION_ROLLEDBACK::XA_RB`, the new state is "non existant". When `end` raises `CORBA::TRANSACTION_ROLLEDBACK` with the minor code

IT_XA_MinorCodes::TRANSACTION_ROLLEDBACK::DEFERRED_ROLLBACK, the new state is "non existent". For every other exception raised by start, suspend, resume and end, there is no state transition.




Figure 4: Association State Diagram

Using a Remote Resource Manager

The Resource servants and the application logic that performs the transactional data access (for example, through embedded SQL in C/C++ calls) do not need to be in the same process. You use the operation `Connector::connect_to_resource_manager` to connect to a remote `XA::ResourceManager`:

```

XA::CurrentConnection_var current_connection =
  xa_connector->connect_to_resource_manager(
    "xa_resource_managers:oracle",
    // the name of an Orbix configuration namespace
    rm,      // object reference to an XA::ResourceManager
            object
    xaosw,   // XA switch
    "",     // open string (empty string means that the
            actual open
            // string is in configuration)
    "",     // close string (empty string means that the
            actual close
            // string is in configuration)
    XA::PROCESS, // thread-model
    false,      // no automatic association
    false,      // do not use dynamic registration
  );
  
```

Some systems (e.g. Oracle) even allow you to create associations between a given transaction and connections to the same database established by different processes: this is referred to as "tightly coupled threads" in the XA specification.

Using a remote resource manager is particularly useful for single-threaded servers, since it allows you to make a data-access server available for other transactions as soon as the transaction has finished with this server (before the completion of the transaction). See the farm demo.

Before Completion Callback

You can register with a resource manager any number of `BeforeCompletionCallback` objects:

```
interface BeforeCompletionCallback
{
    void
    before_completion(
        in CosTransactions::Coordinator tx,
        in CosTransactions::otid_t otid,
        in boolean success
    );
};
interface ResourceManager
{
    unsigned long register_before_completion_callback(
        in BeforeCompletionCallback bcc);
    void unregister_before_completion_callback(
        in unsigned long key);
};
```

The before completion callbacks objects are called by the Resource servant before `prepare`, `commit_one_phase`, and `rollback` on a non-prepared transaction. If any of these before completion callbacks calls raise an exception, the transaction is rolled back. A typical use of the `BeforeCompletionCallback` is to end a suspended association in a single-threaded server. See the farm demo.

Asynchronous Rollback Support

An XA implementation may or may not support asynchronous rollbacks, that is `xa_rollback()` may or may not be called on a transaction while this transaction is actively associated with some connection. This is typically not documented by the XA implementation — OracleXA does not support asynchronous rollbacks, while SybaseXA does.

When you set `supports_async_rollback` to "false" and use a remote resource manager, the XA plug-in uses a transient object to handle asynchronous rollbacks (by deferring them until the association is ended). This transient object is hosted by the root POA, so you have to activate the root POA manager.

Ping Period

The Resource Manager can periodically check that the transactions with which the Resource servants it manages were registered are still alive by calling `get_status` on their respective coordinators. When a call to `get_status` fails (that is, it raises any exception), and the associated Resource is not prepared, this Resource is immediately rolled back.

Transaction Management

This chapter covers some additional areas of transaction management. This includes Synchronization objects, transaction identity and status operations, relationships between transactions and recreating transactions.

Synchronization Objects

Synchronization interface

The transaction service provides a `Synchronization` interface to allow an object to be notified before the start of a transaction's completion and after it is finished. This is useful, for example, for applications integrated with an XA compliant resource manager where the data is cached inside the application. By registering a synchronization object with the transaction the cache can be flushed to the resource manager before the transaction starts to commit. Without the synchronization object any updates made by the application could not be moved from the cache to the resource manager. The `Synchronization` interface is as follows:

```
// IDL (in module CosTransactions)
interface Synchronization :
    CosTransactions::TransactionalObject {
    void before_completion();

    void (in Status s);
};
```

`before_completion()`

This operation is invoked during the commit protocol before any 2PC or 1PC operations have been called, that is before any XA or Resource prepare operations.

An implementation may flush all modified data to the resource manager to ensure that when the commit protocol begins, the data in the resource is up to date.

Raising a system exception causes the transaction to be rolled back as does invoking the `rollback_only()` operation on the `Current` or `Coordinator` interfaces.

The `before_completion()` operation is only called if the transaction is to be committed. If the transaction is being rolled back for any reason this operation is not called.

after_completion()

This operation is invoked after the transaction has completed, that is after all XA or Resource commit or rollback operations have been called. The operation is passed the status of the transaction so it is possible to determine the outcome. It is possible that `before_completion()` has not been called, so the implementation must be able to deal with this possibility.

An implementation can use this operation to release locks that were held on behalf of the transaction or to clean up caches. Raising an exception in this operation has no effect on the outcome of the transaction as this has already been determined. All system exceptions are silently ignored.

register_synchronization()

A synchronization object is registered with a transaction by calling the `register_synchronization()` operation on the transaction's coordinator. Assuming the `SynchronizationImpl` class supports the `Synchronization` interface the following code may be used:

```
// C++
//
// Get the control and coordinator object for the
// current transaction.
//
CosTransactions::Current_var tx_current = ...
CosTransactions::Control_var control =
    tx_current->get_control();
CosTransactionsCoordinator_var coordinator =
    control->get_coordinator();

//
// Create a synchronization servant and activate it in a
// transactional POA. The OTS Policy should be ADAPTS
//
SynchronizationImpl servant = new SynchronizationImpl();
PortableServer::POA_var poa = ...
CosTransactions::Synchronization_var obj =
    sync_servant->activate(poa);

//
// Register the synchronization once with the transaction
//
coord->register_synchronization(obj);
```


The `register_synchronization()` operation raises the `Inactive` exception if the transaction has started completion or has already been prepared. A synchronization object must only be registered once per transaction, this is the application's responsibility.

Note: Unlike resource objects, synchronization objects are not recoverable. The transaction service does not guarantee that either operation on the interface will be called in the event of a failure. It is imperative that applications use a resource object if they need guarantees in these situations (to release persistent locks for example).

Transaction Identity Operations

Coordinator interface identity operations

The `Coordinator` interface provides a number of operations related to the identify of transactions. Some of these operations are also available in the `Current` interface:

```
// IDL (in module CosTransactions)
interface Coordinator {
    boolean is_same_transaction(in Coordinator tc);
    unsigned long hash_transaction();
    unsigned long hash_top_level_tran();
    string get_transaction_name();
    PropagationContext get_txcontext();
    ...
};
```

Table 4: *Coordinator interface identity operations*

Operation	Description
<code>is_same_transaction()</code>	Takes a transaction coordinator as a parameters and returns true if both coordinator objects represent the same transaction; otherwise returns false.
<code>hash_transaction()</code>	Returns a hash code for the transaction represented by the target coordinator object. Hash codes are uniformly distributed over the range of a CORBA unsigned long and are not guaranteed to be unique for each transaction.

Table 4: *Coordinator interface identity operations*

Operation	Description
<code>get_transaction_name()</code>	Returns a string representation of the transaction's identify. This string is not guaranteed to be unique for each transaction so it is only useful for display and debugging purposes. This operation is also available on the <code>Current</code> interface.
<code>get_txcontext()</code>	Returns the <code>PropagationContext</code> structure for the transaction represented by the target coordinator object. Amongst other information, the <code>PropagationContext</code> structure contains the transaction identifier in the <code>current.otid</code> field. See "Recreating Transactions" on page 59 for more information on the structure of the <code>PropagationContext</code> .

Maintaining information in individual transactions

The `is_same_transaction()` and `hash_transaction()` operations are useful when it is necessary for an application to maintain data on a per transaction basis (for example, for keeping track of whether a particular transaction has visited the application before to determine whether a `Resource` or `Synchronization` object needs to be registered). The `hash_transaction()` operation can be used to implement an efficient hash table while the `is_same_transaction()` operation can be used for comparison within the hash table.

For nested transaction families the `hash_top_level_transaction()` is provided. This returns the hash code for the top level transaction.

Transaction Status

Coordinator interface status operations

The `Coordinator::get_status()` operation returns the current status of a transaction. This operation is also provided by `Current::get_status()` for the current transaction. The status returned may be one of the following values:

<code>StatusActive</code>	The transaction is active. This is the case after the transaction has started and before the transaction has started to be committed or rolled back.
<code>StatusCommitted</code>	The transaction has successfully completed its commit protocol.

StatusCommitting	The transaction is in the process of committing.
StatusMarkedRollback	The transaction has been marked to be rolled back.
StatusNoTransaction	There is no transaction. This can only be returned from the <code>Current::get_status()</code> operation and occurs when there is no transaction associated with the current thread of control.
StatusPrepared	The transaction has completed the first phase of the 2PC protocol.
StatusPreparing	The transaction is in the process of the first phase of the 2PC protocol.
StatusRolledBack	The transaction has completed rolling back.
StatusRollingBack	The transaction is in the process of being rolled back.
StatusUnknown	The exact status of the transaction is unknown at this point.

The following code shows how to obtain the status of a transaction from the transaction's coordinator object:

```
// C++
CosTransactions::Coordinator_var coord = ...
CosTransactions::Status status = coord->get_status();
if (status == CosTransactions::StatusActive)
{
    ...
} else if (status == CosTransactions::StatusRollingBack)
{
    ...
} else if ...
```

There are two additional status operations for use within nested transaction families:

- `get_top_level_status()` returns the status of the top-level transaction.
- `get_parent_status()` returns the status of a transaction's parent.

Transaction Relationships

Coordinator interface relationship operations

The `Coordinator` interface provides several operations to test the relationship between transactions. Each operation takes as a parameter a reference to another transaction's coordinator object:

```
// IDL (in module CosTransactions)
interface Coordinator {
    boolean is_same_transaction(in Coordinator tc);
    boolean is_related_transaction(in Coordinator tc);
    boolean is_ancestor_transaction(in Coordinator tc);
    boolean is_descendant_transaction(in Coordinator tc);
    boolean is_top_level_transaction();
    ...
};
```

Table 5: *Coordinator interface relationship operations*

Operation	Description
<code>is_same_transaction()</code>	Returns true if both coordinator objects represent the same transaction; otherwise returns false.
<code>is_related_transaction()</code>	Returns true if both coordinator objects represent transactions in the same nested transaction family; otherwise returns false.
<code>is_ancestor_transaction()</code>	Returns true if the transaction represented by the target coordinator object is an ancestor of the transaction represented by the coordinator parameter; otherwise returns false. A transaction is an ancestor to itself and a parent transaction is an ancestor to its child transactions.
<code>is_descendant_transaction()</code>	Returns true if the transaction represented by the target coordinator object is a descendant of the transaction represented by the coordinator parameter; otherwise returns false. A transaction is a descendant of itself and is a descendent of its parent.

Table 5: *Coordinator interface relationship operations*

Operation	Description
is_top_level_transaction()	Returns true if the transaction represented by the target coordinator object is a top-level transaction; otherwise returns false.

Example

The following code tests if the transaction represented by the coordinator c1 is an ancestor of the transaction represented by the coordinator c2:

```
// C++
CosTransactions::Coordinator_var c1 = ...
CosTransactions::Coordinator_var c2 = ...
if (c1->is_ancestor_transaction(c2))
{
    // c1 is an ancestor of c2
}
else
{
    // c1 is not an ancestor of c2
}
```

Recreating Transactions

TransactionFactory interface

The TransactionFactory interface provides the create() operation for creating new top-level transactions. The interface also provides a recreate() operation to import an existing transaction into the

local context. The `recreate()` is passed a `PropagationContext` structure and returns a `Control` object representing the recreated transaction. The interfaces and types are declared as follows:

```
// IDL (in module CosTransactions)
struct otid_t {
    long formatID;
    long bqual_length;
    sequence <octet> tid;
};

struct TransIdentity {
    Coordinator coord;
    Terminator term;
    otid_t otid;
};

struct PropagationContext {
    unsigned long timeout;
    TransIdentity current;
    sequence <TransIdentity> parents;
    any implementation_specific_data;
};

interface TransactionFactory
{
    Control recreate(in PropagationContext ctx);
    ...
};

interface Coordinator
{
    PropagationContext get_txcontext();
    raises (Unavailable);
    ...
};
```

The `PropagationContext` is a structure that encodes sufficient information about the transaction to successfully recreate it. To get the `PropagationContext` for a transaction use the `get_txcontext()` operation provided by the `Coordinator` interface.

Example

The following code shows how to use the `get_txcontext()` and `recreate()` operations to explicitly import a transaction given a reference to the `Control` object for a foreign transaction:

```
// C++
CosTransactions::Control_var foreign_control = ...
CosTransactions::Coordinator_var foreign_coord =
    foreign_control->get_coordinator();
CosTransactions::PropagationContext_var ctx =
    foreign_coord->get_txcontext();

CosTransactions::TransactionFactory_var tx_factory = ...
CosTransactions::Control_var control =
    tx_factory->recreate(ctx);
```

The PropagationContext structure contains the transaction's global identifier in the current.otid field. This is essentially a sequence of octets divided into two parts: a global transaction identifier and a branch qualifier. This structure is indented to match the XID transaction identifier format for the X/Open XA specification.

Writing Recoverable Resources

The OTS supports resource objects to allow applications to participate in transactions. For example, an application might maintain some data for which ACID properties are required. This chapter describes the `CosTransactions::Resource` interface; how resource objects participate in the transaction protocols and the requirements for implementing resource objects.

The Resource Interface

Resource interface transaction operations

The `CosTransactions::Resource` interface provides a means for applications to participate in an OTS transaction. The interface is defined as follows:

```
// IDL (in module CosTransactions)
interface Resource
{
    void commit_one_phase()
        raises (HeuristicHazard);

    Vote prepare()
        raises (HeuristicMixed,
              HeuristicHazard);

    void rollback()
        raises (HeuristicCommit,
              HeuristicMixed,
              HeuristicHazard);

    void commit()
        raises (NotPrepared,
              HeuristicRollback,
              HeuristicMixed,
              HeuristicHazard);

    void forget();
};
```

Resource object implementations cooperate with the OTS, through these five operations, to ensure the ACID properties are satisfied for the whole transaction. Each resource object represents a single participant in a transaction and throughout the lifecycle of the resource it must respond to the invocations by the OTS until the resource object is no longer needed. This may include surviving the failure of the process or node hosting the resource object or the failure of the process or node hosting the OTS implementation.

Overview of the use of resource objects

Figure 5 shows a high level picture of how clients, applications, the OTS and resource objects interoperate to achieve the ACID properties.

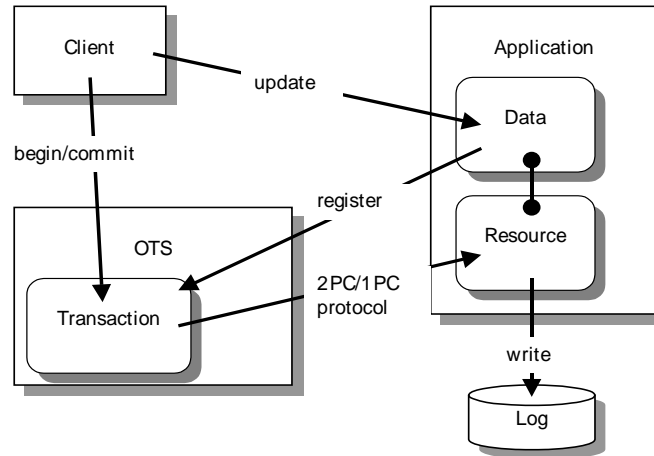


Figure 5: *Relationship between resources and transactions*

The steps involved are:

1. The client contacts the OTS implementation and creates a transaction.
2. The client makes invocations on the application within the context of the transaction and updates some data.
3. The application detects that the data is being updated and creates a resource object. The resource object is registered with the transaction.
4. The client completes by contacting the OTS implementation and attempting to commit the transaction.
5. The transaction initiates the commit protocol. The choice of which protocol to use (either 1PC or 2PC) depends on the number of resource objects registered with the transaction and whether the OTS supports the 1PC optimization.
6. Assuming the 2PC protocol is being used, the OTS sends a prepare message to the resource. The resource stably stores enough information to recover in case of a crash (for example, by writing the changes to a log file). The resource object votes to commit the transaction.
7. The OTS gathers the votes of all resource objects and decides the outcome of the transaction. This decision is send to all registered resource objects.
8. The resource object upon receiving the commit or rollback message makes the necessary changes and saves the decision to the log.
9. The OTS returns the outcome to the client.

Creating and Registering Resource Objects

Implementing servants for resource objects

Implementing servants for resource objects is similar to any servant implementation. The resource servant class needs to inherit from the `POA_CosTransactions::Resource` class to extend the `ResourcePOA` class and provide implementations for the five resource operations. For example, the following class can be used to implement a resource servant:

```
// C++
class ResourceImpl : public POA_CosTransactions::Resource
{
public:

    ResourceImpl();

    virtual ~ResourceImpl();

    CosTransactions::Vote
    prepare()
    throw (CORBA::SystemException,
          CosTransactions::HeuristicMixed,
          CosTransactions::HeuristicHazard);

    void
    rollback()
    throw (CORBA::SystemException,
          CosTransactions::HeuristicCommit,
          CosTransactions::HeuristicMixed,
          CosTransactions::HeuristicHazard);

    void
    commit()
    throw (CORBA::SystemException,
          CosTransactions::NotPrepared,
          CosTransactions::HeuristicRollback,
          CosTransactions::HeuristicMixed,
          CosTransactions::HeuristicHazard);

    void
    commit_one_phase()
    throw (CORBA::SystemException,
          CosTransactions::HeuristicHazard);

    void
    forget()
    throw (CORBA::SystemException);
};
```

Creating resource objects

Resource objects, once prepared, must survive failures until the 2PC protocol has completed. During recovery any resource objects requiring completion must be recreated using the same identifier so the transaction coordinator can deliver the outcome. This means that resource objects must be created within a POA with a `PERSISTENT` lifespan policy and a `USER_ID` ID assignment policy. See the sections "Setting Object Lifespan" and "Assigning Object IDs" in the chapter "Managing Server Objects" in the *CORBA Programmer's Guide* for more details.

Tracking resource objects

Each resource object can only be used once and may only be registered with one transaction. It is up to the application to keep track of whether it has seen a particular transaction before. This can be done efficiently using the `hash_transaction()` and `is_same_transaction()` operations provided by the `Coordinator` interface to implement a hash map (see "[Transaction Identity Operations](#)" on page 55 for details).

Some form of unique identifier must be used for the resource object's `ObjectId`. One possibility is to use the transaction identifier (obtained from the `otid` field in the transaction's propagation context).

Registering resource objects

Registration of a resource object with a transaction is done by the `register_resource()` operation provided by the transaction's coordinator object. For example, the following code sample shows a resource servant and object being created and registered with a transaction:

```
// C++
CosTransactions::Current_var tx_current = ...

// Get the transaction's coordinator object.
CosTransactions::Control_var control =
    tx_current->get_control();
CosTransactions::Coordinator_var coord =
    control->get_coordinator();

// Create resource servant.
ResourceImpl* servant = new ResourceImpl();
```

```

// Create resource object. The POA referenced by
resource_poa
// has the PERSISTENT lifespan policy and the USER_ID ID
// assignment policy.
PortableServer::POA_var resource_poa = ...
PortableServer::ObjectId_var oid = ...

resource_poa->activate_object_with_id(oid, servant);

CORBA::Object_var obj =
    resource_poa->servant_to_reference(servant);

CosTransactions::Resource_var resource =
    CosTransactions::Resource::_narrow(obj);

// Register the resource with the transaction
// coordinator.
CosTransactions::RecoveryCoordinator_var rec_coord =
    coord->register_resource(resource);

```

The `register_resource()` operation returns a reference to a recovery coordinator object:

```

// IDL (in module CosTransactions)
interface Coordinator
{
    RecoveryCoordinator register_resource(in Resource r)
        raises(Inactive);
    ...
};

interface RecoveryCoordinator
{
    Status replay_completion(in Resource r)
        raises(NotPrepared);
};

```

The recovery coordinator object supports a single operation, `replay_completion()`, that is used for certain failure scenarios (see [“Failure of the Transaction Coordinator” on page 73](#)). Resource objects must hold onto the recovery coordinator reference.

The `register_resource()` operation raises the `Inactive` exception if the transaction is no longer active.

Resource Protocols

Protocols supported by resource objects

Resource object implementations cooperate with the transaction coordinator to achieve the ACID properties. This section examines the protocols that resource objects are required to support:

- Rolling back a transaction.
- The 2-phase-commit protocol.
- Read-only resources.

- The 1-phase-commit protocol.
- Heuristic outcomes.
- Failure and recovery

Transaction Rollbacks

Up until the time the coordinator makes the decision to commit a transaction, the transaction may be rolled back for a number of reasons. These include:

- A client calling the `rollback()` operation.
- Attempting to commit the transaction after the transaction has been marked to be rolled-back with the `rollback_only()` operation.
- The transaction being timed-out.
- The failure of any participant in the transaction.

When the transaction is rolled-back all registered resource are rolled-back via the `rollback()` operation. [Figure 6](#) shows a transaction with two registered resource objects being rolled back after a timeout.

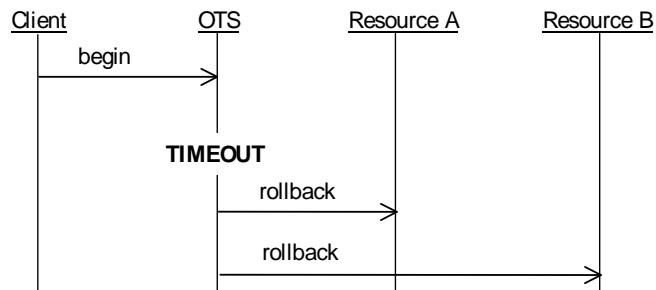


Figure 6: *Rollback after a timeout*

Rollbacks may also occur during the 2PC protocol (see below).

The 2-Phase-Commit Protocol

The 2-phase-commit (2PC) protocol is designed so that all participants within a transaction know the final outcome of the transaction. The final outcome is decided by the transaction coordinator but each resource object participating can influence this decision.

During the first phase, the transaction coordinator invokes the `prepare()` operation on each resource asking it to prepare to commit the transaction. Each resource object returns a vote which may be one of three possible values: `VoteCommit` indicates the resource is prepared to commit its part of the transaction; `VoteRollback` indicates the transaction must be rolled-back; and `VoteReadOnly` indicates the resource is no longer interested in the outcome of the transaction (see ["Read-Only Resources" on page 69](#)).

The coordinator makes a decision on whether to commit or rollback the transaction based on the votes of the resource objects. Once a decision has been reached the second phase commences where the resource objects are informed of the transaction outcome.

In order for the coordinator to decide to commit the transaction, each resource object must have either voted to commit the transaction or indicated that it is no longer interested in the outcome. Once a resource has voted to commit, it must wait for the outcome to be delivered via either the `commit()` or `rollback()` operation. The resource must also survive failures. This means that sufficient information must be stable stored so that during recovery the resource object and its associated state can be reconstructed. [Figure 7](#) shows a successful 2PC protocol with two resources objects. Both resources return `VoteCommit` from the `prepare()` operation and the coordinator decides to commit the transaction resulting in the `commit()` operations being invoked on the resources.

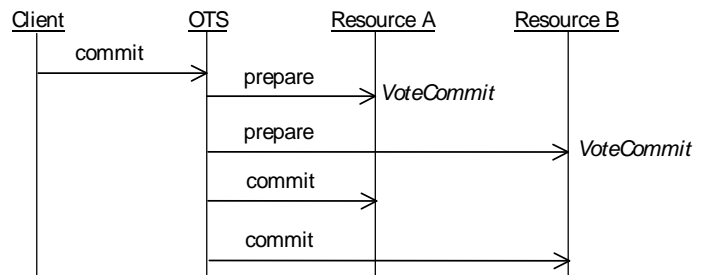


Figure 7: *Successful 2PC protocol with two resources*

If one resource returns `VoteRollback` the whole transaction is rolled back. Resources which have already been prepared and which voted to commit and resources which have not yet been prepared are told to rollback via the `rollback()` operation. [Figure 8](#) shows `VoteRollback` being returned by one resource which results in the other resource being told to rollback.




Figure 8: *Voting to rollback the transaction.*

Read-Only Resources

A resource can return `VoteReadOnly` from the `prepare()` operation which means the resource is no longer interested in the outcome of the transaction. This is useful, for example, when the

application data associated with the resource was not modified during the transaction. Here it does not matter whether the transaction is committed or rolled back. By returning `VoteReadOnly` the resource is opting out of the 2PC protocol and the resource object will not be contacted again by the transaction coordinator.

Figure 9 shows the 2PC protocol with two resource objects. In the first phase, the first resource returns `VoteReadOnly` and the second resource returns `VoteCommit`. During the second phase only the second resource is informed of the outcome (commit in this case).




Figure 9: A resource returning `VoteReadOnly`.

The 1-Phase-Commit Protocol

The 1-phase-commit (1PC) protocol is an optimization of the 2PC protocol where the transaction only has one participant. Here the OTS can short circuit the 2PC protocol and ask the resource to commit the transaction directly. This is done by invoking the `commit_one_phase()` operation rather than the `prepare()` operation.

When the 1PC protocol is used the OTS is delegating the commit decision to the resource object. If the resource object decides to commit the transaction, the `commit_one_phase()` operation returns successfully. However, if the resource decides to rollback the transaction it must raise the `TRANSACTION_ROLLEDBACK` system exception. Figure 10 shows a successful 1PC protocol.




Figure 10: A successful 1PC protocol.

Figure 11 shows a 1PC protocol resulting in the transaction being rolled-back.




Figure 11: The 1PC protocol resulting in a roll-back.

It is possible for the `commit_one_phase()` operation to be called even when more than one resource is registered with a transaction when resources return `VoteReadOnly` from `prepare()`. Assume for example there are three resources registered with a transaction. If the first two resources both return `VoteReadOnly` the third resource does not need to be prepared and the `commit_one_phase()` operation can be used instead.

Heuristic Outcomes

Heuristics outcomes occur when at least one resource object unilaterally decides to commit or rollback its part of the transaction and this decision is in conflict with the eventual outcome of the transaction. For example, a resource may have a policy that, once prepared, it will decide to commit if no outcome has been delivered within a certain period. This might be done to free up access to shared resources.

Any unilateral decisions made must be remembered by the resource. When the eventual outcome is delivered to the resource it must reply according to the compatibility of the decisions. For example, if the resource decides to commit its part of the transaction and the transaction is eventually rolled back, the resource's `rollback()` operation must raise the `HeuristicCommit` exception. The following table lists the resource's response for the various possible outcomes.

Table 6: *Heuristic Outcomes*

Resource Decision	Transaction Outcome	Resource's Response
Commit	Commit	<code>commit()</code> returns successfully.
Commit	Rollback	<code>rollback()</code> raises <code>HeuristicCommit</code>
Rollback	Rollback	<code>rollback()</code> returns successfully
Rollback	Commit	<code>commit()</code> raises <code>HeuristicRollback</code>

Once a resource has raised a heuristic exception it must remember this until the `forget()` operation has been called by the OTS (see [Figure 12](#)). For example, after a failure the OTS might invoke the rollback operation again in which case the resource must re-raise the `HeuristicCommit` exception. Once the `forget()` operation has been called the resource object is no longer required and can be deleted.

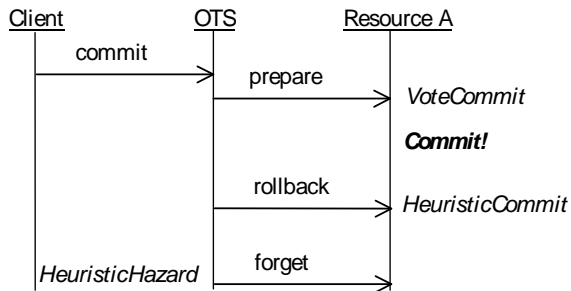


Figure 12: Raising the `HeuristicCommit` exception

Heuristic outcome are reported to the client only if `true` is passed to the `commit()` operation provided by the OTS Current object. They are reported by raising one of the exceptions: `HeuristicMixed` or `HeuristicHazard`. `HeuristicMixed` means a heuristic decision has been made resulting in some updates being committed and some being rolled back. `HeuristicHazard` indicates that a heuristic decision may have been made.

If the `commit_one_phase()` operation is called by the transaction coordinator, the commit decision is delegated to the resource implementation. This means that if the operation fails (that is results in a system exception other than `TRANSACTION_ROLLEDBACK` being raised) then the coordinator cannot know the true outcome of the transaction. For this case, the OTS raises the `HeuristicHazard` exception.

Failure and Recovery

Resource objects need to be able to deal with the failure of the process or node hosting the resource and the failure of the process or node hosting the OTS implementation.

Failure of the Resource

If the process or node hosting the resource object fails after the resource has been prepared, the resource object must be recreated during recovery so that the outcome of the transaction can be delivered to the resource. [Figure 13](#) shows a crash occurring sometime after the resource has been prepared but before the coordinator invokes the `commit()` operation. When the coordinator does invoke the `commit()` operation the resource object is not active and the coordinator will attempt to commit later. In the meantime the resource object is recreated and waits

for the `commit()` operation to be invoked. The next time the coordinator calls `commit()` the resource receives the invocation and proceeds as normal.




Figure 13: Recovery after the failure of a resource object

If the failure occurs before the resource has been prepared, there is no need to recreate the resource during recovery. When the 2PC protocol starts the OTS will not be able to contact the resource and the transaction will be rolled back.

Failure of the Transaction Coordinator

If the process or node hosting the transaction coordinator fails there are two possible ways in which the failure is resolved:

1. The transaction coordinator recovers and eventually sends the outcome to the resource. Here, the resource does not need to participate in the recovery; either the `commit()` or `rollback()` operation will be invoked as normal.
2. The resource detects that no outcome has been delivered and asks the transaction coordinator to complete the transactions. This is done using the `replay_completion()` operation provided by the recovery coordinator object.

The second way of resolving the failure of the OTS is required because the OTS supports a behavior called presumed rollback. With presumed rollback, if a transaction is rolled back the coordinator is not required to stably store this fact. Instead, on recovery if there is no information available on a transaction, the transaction is presumed to have rolled back. This saves on the amount of data that must be stably stored but means the resource object must check to see if the transaction has been rolled back.

Recall from [“Creating and Registering Resource Objects” on page 65](#) when a resource is registered with the coordinator a reference to a recovery coordinator object is returned. The recovery coordinator supports the `RecoveryCoordinator` interface:

```
// IDL (in module CosTransactions)
interface RecoveryCoordinator
{
    Status replay_completion(in Resource r)
        raises (NotPrepared);
};
```

The sole operation, `replay_completion()`, takes a resource object and returns the status of the transaction. If the transaction has not been prepared the `NotPrepared` exception is raised. The `replay_completion()` operation is meant to hint to the coordinator that the resource is expecting the transaction to be completed.

To support detecting presumed rolled-back transactions, the `replay_completion()` operation is used to detect if the transaction still exists. If the transaction still exists the operation will either return a valid status or the `NotPrepared` exception. However, if the transaction no longer exists the `OBJECT_NOT_EXIST` system exception will be raised (other system exceptions should be ignored).

By periodically calling `replay_completion()` and checking for the `OBJECT_NOT_EXIST` exception, the resource object can detect rolled-back transactions (see [Figure 14](#)). This periodic calling of `replay_completion()` must be done before the resource has been prepared, after the resource has been prepared and after recovery of the resource due to a crash. To implement the latter, the resource object needs to stably store the recovery coordinator reference (for example using a stringified IOR) so that after a failure, the recovery coordinator can be contacted.

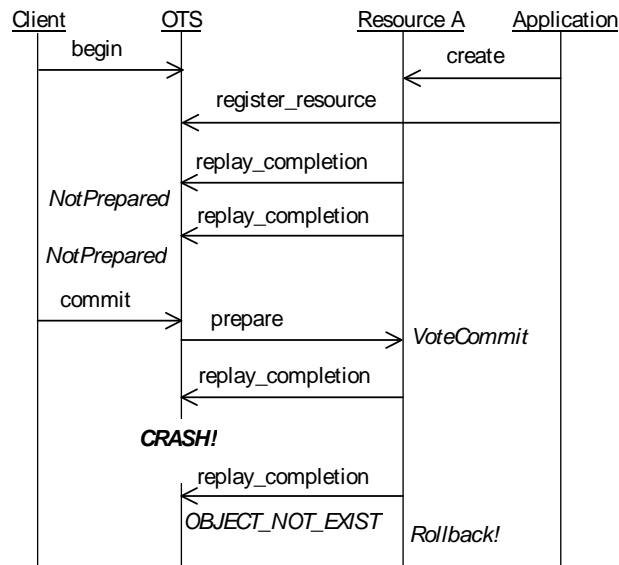


Figure 14: Use of the `replay_completion()` operation

Responsibilities and Lifecycle of a Resource Object

This section details the responsibilities of a resource object for each operation and shows the lifecycle of a resource object.

prepare()

`Vote prepare()` raises `(HeuristicMixed, HeuristicHazard)`;

The `prepare()` operation is called during the first phase of the 2PC protocol allowing the resource to vote in the transaction's outcome and if necessary prepare for eventual commitment.

Voting is done by returning one of the three values `VoteCommit`, `VoteRollback` and `VoteReadOnly`:

<code>VoteCommit</code>	This indicates that the resource is willing to commit its part of the transaction and has fully prepared itself for the eventual outcome of the transaction. The next invocation on the resource will be either <code>commit()</code> or <code>rollback()</code> .
<code>VoteRollback</code>	This indicates that the resource has decided to rollback the transaction. This ensures that the transaction will be rolled back. The resource object can forget about the transaction and no further operations will be invoked on the resource object.
<code>VoteReadOnly</code>	This indicates that the resource does not want to be further involved in the 2PC protocol. This does not affect the transaction outcome and the resource object can forget about the transaction. No further operations will be invoked on the resource object.

If a resource object returns `VoteCommit` it must stably store sufficient information so that in the event of a failure, the resource object and its state can be reconstructed and continue to participate in the 2PC protocol. The actual information that is saved depends on the application, but typically it will include the following:

- The identity of the transaction. This can be obtained from the `otid` field in the transaction's propagation context which in turn is obtained by the `get_txcontext()` operation on the transaction's coordinator.
- The `ObjectID` for the resource.
- The reference for the recovery coordinator object associated with the resource. This can be saved as a stringified IOR obtained by the `object_to_string()` operation.
- Sufficient information to redo or undo any modifications made to application data by the transaction.

The `prepare()` operation can raise two exceptions dealing with heuristic outcomes: `HeuristicMixed` and `HeuristicHazard`. These exceptions may be used internally in an OTS implementation; most resource implementations do not need to raise these exceptions.

commit()

`void commit()` raises (`NotPrepared`, `HeuristicRollback`,
`HeuristicMixed`, `HeuristicHazard`)

The `commit()` operation is called during the second phase of the 2PC protocol after the coordinator has decided to commit the transaction. The `commit()` operation may be invoked multiple times due to various failures such as a network error, failure of the OTS and failure of the application.

Typically the `commit()` operation does the following:

- Make permanent any modifications made to the data associated with the resource.
- Cleans up all traces of the transaction, including information stably stored for recovery.

The `commit()` operation can raise one of four user exceptions: `NotPrepared`, `HeuristicRollback`, `HeuristicMixed`, `HeuristicHazard`. The `NotPrepared` exception must be raised if `commit()` is invoked before the resource has been prepared (that is, returned `VoteCommit` from the `prepare()` operation).

The `HeuristicRollback` exception must be raised if the resource had decided to rollback its part of the transaction after being prepared and prior to the `commit()` operation being invoked. If this exception is raised it must be raised on future invocations of the `commit()` operation and the resource must wait for the `forget()` operation to be invoked before cleaning up the transaction.

The `HeuristicMixed` and `HeuristicHazard` exceptions may be used internally in an OTS implementation; most resource implementations do not need to raise these exceptions.

rollback()

`void rollback()` raises (`HeuristicCommit`, `HeuristicMixed`,
`HeuristicHazard`)

There are two occasions when the `rollback()` operation is called:

1. During the second phase of the 2PC protocol after the coordinator has decided to commit the transaction.
2. When the transaction is rolled back prior to the start of the 2PC protocol. This may occur for several reasons including the client invoking the `rollback()` operation on the OTS Current object, the transaction begin timed-out, and an attempt to commit a transaction that has been marked for rollback.

The `rollback()` operation may be invoked multiple times due to various failures such as a network error, failure of the OTS and failure of the application.

Typically the `rollback()` operation does the following:

- Undo any modifications made to the data associated with the resource.
- Cleans up all traces of the transaction, including information stably stored for recovery.

The `rollback()` operation can raise one of three user exceptions: `HeuristicCommit`, `HeuristicMixed`, `HeuristicHazard`. The `HeuristicCommit` exception must be raised if the resource had decided to commit its part of the transaction after being prepared and prior to the `rollback()` operation being invoked. If this exception is raised it must be raised on future invocations of the `rollback()` operation and the resource must wait for the `forget()` operation to be invoked before cleaning up the transaction. Heuristic exceptions can only be raised if the resource has been prepared.

The `HeuristicMixed` and `HeuristicHazard` exceptions may be used internally in an OTS implementation; most resource implementations do not need to raise these exceptions.

commit_one_phase()

`void commit_one_phase()` raises (`HeuristicHazard`)

The `commit_one_phase()` operation may be invoked when there is only one resource registered with the transaction. The resource decides whether to commit or rollback the transaction. Typically the `commit_one_phase()` operation does the following:

- An attempt is made to commit any changes made to the application data. If this succeeds the operation returns normally; otherwise the changes are undone and the `TRANSACTION_ROLLEDBACK` system exception is raised.
- Cleans up all traces of the transaction.

The `HeuristicHazard` exception must be raised if the resource cannot determine whether the commit attempt was successful or not. If this exception is raised the resource must wait for the `forget()` operation to be invoked before cleaning up the transaction.

forget()

`void forget()`

The `forget()` operation is called after the resource object raised a heuristic exception from either `commit()`, `rollback()` or `commit_one_phase()`. The `forget()` operation may be invoked multiple times due to various failures such as a network error, failure of the OTS and failure of the application. Typically the resource cleans up all traces of the transaction, including information stably stored for recovery.

Resource Object Checklist

The following is a list of things to remember when implementing recoverable resource objects:

- A resource object can only be registered with one transaction. At the end of the resource's lifecycle the resource must be deactivated.
- Resource objects need unique identifiers. This means they must be created in a POA with a `USER_ID` ID assignment policy.
- Resource objects must be able to be recreated after a failure. This means they must be created in a POA with a `PERSISTENT` lifecycle policy.
- Resource objects must implement both the 2PC operations (`prepare()`, `commit()`, `rollback()` and `forget()`) as well as the 1PC operation (`commit_one_phase()`).
- Only return `VoteCommit` from the `prepare()` operation if the resource can commit the transaction and has stably stored sufficient state to be recreated after a failure.
- If a resource object wants to opt out of the 2PC protocol, it should return `VoteReadOnly` from the `prepare()` operation.
- If the resource takes heuristic decisions, the decisions must be remembered and reported to the OTS.
- Periodically call the `replay_completion()` operation to check for presumed rollback transactions.
- Resources are expensive in terms of 2PC messages and stable storage for recovery. Design your applications to minimize the number of resources used.

Interoperability

This chapter describes how the Orbix OTS interoperates with older releases of Orbix and with other OTS implementations including the Orbix 3 OTS.

Use of InvocationPolicies

Deprecated policies

This release of Orbix introduces the OTSPolicies, InvocationPolicies and NonTxTargetPolicies that replace the deprecated TransactionPolicies. The deprecated TransactionPolicies (for example, Requires_shared and Allows_shared) are supported allowing interoperability between different releases of Orbix.

When creating Orbix transactional POAs that must interoperate with previous releases, the policies for the POA must include the deprecated TransactionPolicy as well as the OTSPolicy and InvocationPolicy. See ["Migrating from TransactionPolicies" on page 40](#) for more details.

Note: Support for the TransactionPolicy type may be discontinued in a future Orbix release. It is recommended that only OTSPolicies and InvocationPolicies be used.

Use of the TransactionalObject Interface

Enabling support for the TransactionalObject interface

Version 1.1 of the OTS specification uses inheritance from the empty `CosTransactions:TransactionalObject` interface to indicate the transactional requirements of an object. For example, the Orbix 3 OTS only supports the `TransactionalObject` interface and not the policies.

Orbix provides support for the `TransactionalObject` interface, allowing different behaviors to be configured. This support needs to be enabled by setting the `plugins:ots:support_ots_v11` configuration variable to "true" (by default this support is not enabled). Once enabled, an object which supports the `TransactionalObject` interface is interpreted as having an effective OTSPolicy which depends on the value of the `plugins:ots:ots_v11_policy` configuration variable. [Table 7](#) details

this mapping:

Table 7: *Mapping TransactionalObject to OTSPolicies*

Inherits from Transactional Object	Value of plugins:ots:ots_v11_policy	Effective OTSPolicy Value
No	n/a	FORBIDS
Yes	"requires"	REQUIRES
Yes	"adapts"	ADAPTS

The default value for the `plugins:ots:ots_v11_policy` is "requires" since this is the default behavior for the Orbix 3 OTS. For backward compatibility with previous Orbix releases a value of "allows" is interpreted as "adapts".

It is recommended that when support for `TransactionalObject` is enabled, the `NonTxTargetPolicy PERMIT` should be used.

If an object supports `TransactionalObject` and also uses OTSPolicies, the OTSPolicies take priority; compatibility checks are not done.

To summarize, to enable support for the `TransactionalObject` interface the following is required:

1. Set the `plugins:ots:support_ots_v11` configuration variable to "true".
2. Set the `plugins:ots:ots_v11_policy` configuration variable to either "requires" (the default) or "adapts".
3. Use the `PERMIT NonTxTargetPolicy` (for example, by setting the `policies:non_tx_target_policy` configuration variable to "permit").

Interoperability with Orbix 3 OTS Applications

This section details how an Orbix client can interoperate with an existing Orbix 3 OTS application. Since Orbix 3 supports only the `TransactionalObject` interface this section is an extension of the previous section ["Use of the TransactionalObject Interface" on page 79](#)

Details on using the Encina OTS are covered in ["The Encina Transaction Manager" on page 89](#).

Orbix 3 OTS Interoperability

[Figure 15](#) shows an Orbix client working with an existing Orbix 3 OTS application. The first thing to note is that the Orbix 3 OTS always requires a full 2PC transaction manager such as that provided by the Encina OTS (see ["The Encina Transaction Manager" on page 89](#)) or the `otstf` provided with Orbix 3. A 1PC-only transaction created by the OTS Lite transaction manager

will not be usable by the Orbix 3 OTS. This means that the Orbix client must be configured to use an external transaction factory to create transactions.

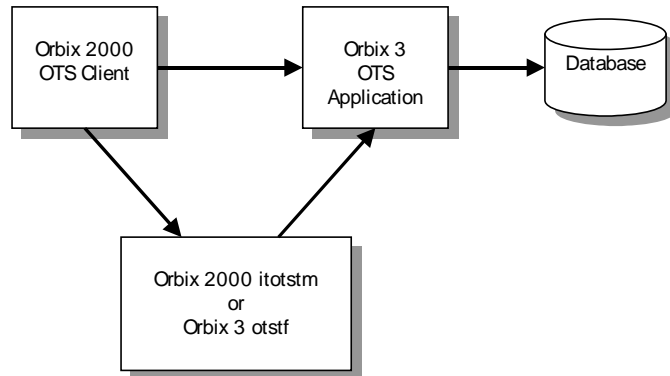


Figure 15: *Interoperability with Orbix 3 OTS Applications*

Using otstf as transaction manager

To get the Orbix client to use the Orbix 3 otstf server as its transaction manager, the

`initial_references:TransactionFactory:reference` configuration variable must be set to the reference of the otstf's transaction factory object. This can be done by passing the `-T` switch to the otstf and copying the IOR reference output. Alternatively the otstf can publish its name to the name service using the `-t` switch and a suitable corbaname URL can be used as the reference value (see the section "Resolving Names with corbaname" in the chapter "Naming Service" in the *CORBA Programmer's Guide*).

The Orbix 3 OTS application must be enabled to import standard transaction contexts. This is done by setting the Orbix 3 `OrbixOTS.INTEROP` configuration variable to `"TRUE"`.

The final consideration is the mapping from inheritance from `TransactionalObject` to the effective `OTSPolicy`. The Orbix 3 OTS provides a proprietary policy mechanism which mimics the behavior of the `OTSPolicies` `REQUIRES` and `ADAPTS` (the default being `REQUIRES`). Therefore, when selecting the value for the `plugins:ots:ots_v11_policy` configuration variable, make sure it matches the policy expected by the Orbix 3 application.

Bypassing otstf

It is possible to bypass the use of the otstf server and use the transaction factory provided by the Orbix 3 OTS application. This is done by modifying the Orbix 3 application to publish its internal transaction factory reference. This is illustrated in the following code:

```
// Orbix 3 OTS C++ Application Code
CORBA::ORB_var orb = ...
OrbixOTS::Server_var ots = ...

// Get reference to the local transaction factory.
CosTransactions::TransactionFactory_var tx_factory =
    ots->get_transaction_factory_reference();

// Publish reference (eg, to the name service or a file)
```

Summary

The following is a checklist for enabling interoperability between Orbix clients and Orbix 3 OTS applications.

1. Set the `plugins:ots:support_ots_v11` configuration variable to `"true"`.
2. Set the `plugins:ots:ots_v11_policy` configuration variable to match the equivalent Orbix 3 OTS policy for the `TransactionalObject` interface.
3. Use the `PERMIT NonTxTargetPolicy`.
4. Set the `initial_references:TransactionFactory:reference` configuration variable to refer to either the Orbix 3 otstf's transaction factory or another transaction factory that supports 2PC.
5. Set the Orbix 3 `OrbixOTS.INTEROP` configuration variable to `"TRUE"`.

For more information on the use of the otstf server and setting Orbix 3 transaction policies, refer to the Orbix 3 OTS manual.

Using the Orbix 3 otstf with Orbix Applications

Using Orbix 3 otstf transaction manager

Another possible use of Orbix 3 is to use the 2PC otstf transaction manager with an Orbix OTS application. This setup is shown in [Figure 16](#).

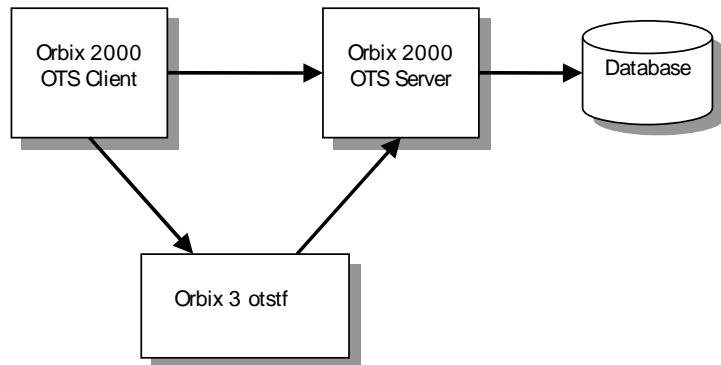


Figure 16: *Using and alternative OTS Implementation*

This setup is achieved by setting the `initial_references:TransactionFactory:reference` configuration variable to refer to the otstf's transaction factory.

OTS Plug-Ins and Deployment Options

Orbix provides a generic OTS plugin that provides an implementation of the OTS Current object including transaction propagation. In addition there are two OTS transaction manager implementations: OTS Lite, which provides a lightweight transaction coordinator supporting only the 1PC protocol, and OTS Encina, which provides full recoverable 2PC support. This chapter discusses deployment options.

OTS Plug-ins

Orbix provides a generic OTS plugin that provides an implementation of the OTS Current object including transaction propagation.

There are two OTS transaction manager implementations:

- OTS Lite
- OTS Encina.

OTS Lite

OTS Lite provides lightweight transaction coordinator supporting only the 1PC protocol. It is available as an application plug-in and requires minimal configuration and administration but can only be used by applications with only a single resource manager.

OTS Encina

OTS Encina provides full recoverable 2PC support allowing it to be used by applications that are using one or more resource managers. It is available as a standalone service and as a application plug-in.

Note:

- OTS Encina is only available in the Orbix Enterprise Edition.
- OTS Encina is not supported on 64 bit windows platforms.

Features in OTS

Table 8 shows the features supported by these pieces.

Table 8: *Features in OTS Implementation*

Feature	Generic OTS	OTS Lite	OTS Encina
Current Object	Y		
Transaction Policies	Y		
Old Transaction Policies	Y		
TransactionalObject	Y		
1PC Protocol		Y	Y
2PC Protocol		N	Y
Resource Objects		Y	Y
Synchronization Objects		Y	Y
Nested Transactions		N	N
iPAC Management		N	Y
XA Support		Y	Y
Application Plug-In	Y	Y	Y

The OTS Plug-In

Purpose of the OTS plug-in

Any application using the OTS Current object needs to load the OTS plug-in. This plug-in provides an implementation of the OTS Current object which provides the thread/transaction association, propagation of the current transaction to transactional objects and the policies OTSPolicy, InvocationPolicy and NonTxTargetPolicy. In addition the OTS plug-in provides the client stubs for the CosTransactions module, so applications need to link with the OTS plug-in library.

In OTS plug-in does not provide any transaction manager functionality. Instead the OTS plug-in delegates elsewhere using the standard CosTransactions module APIs (see [Figure 17](#)). This allows different deployment options to be easily supported through configuration.




Figure 17: *The Generic OTS Plug-In*

Loading the OTS plug-in

There are two ways in which the OTS plug-in can be loaded:

1. Explicitly adding the plug-in name "ots" to the `orb_plugins` configuration variable. For example: `orb_plugins = [..., "ots"];`
2. Setting the `initial_references:TransactionCurrent:plugin` configuration variable to the value "ots". This causes the OTS plug-in to be loaded when `resolve_initial_references("TransactionCurrent")` is called. When using this way, `resolve_initial_references()` should be called immediately after `ORB_init()` has been called and before any transaction POAs are created.

When the OTS plug-in is initialized it obtains a reference to a transaction factory object by calling `resolve_initial_references("TransactionFactory")`. So changing which transaction manager to use is just a matter of using configuration to change the outcome of `resolve_initial_references()`.

Deployment scenarios

The remainder of this section describes three possible deployment scenarios for C++:

- Using the OTS Lite plug-in when only 1PC transactions are required.
- Using the `itotstm` service with the OTS Encina plug-in where recoverable 2PC transactions are required.
- Using the OTS Encina plug-in loaded into the application itself.

For more information, see the *Orbix Deployment Guide*.

The OTS Lite Plug-In

The OTS Lite plug-in is a lightweight transaction manager that only supports the 1PC protocol. This plug-in allows applications that only access a single transactional resource to use the OTS APIs without incurring a large overhead, but allows them to migrate easily to the more powerful 2PC protocol by switching to a different transaction manager. [Figure 18](#) shows a client/server deployment that uses the OTS Lite plug-in.




Figure 18: *Deployment using the OTS Lite Plug-In*

As usual both the client and server applications must load the OTS plug-in. In addition the client application loads the OTS Lite plug-in, allowing the client to create 1PC transaction locally.

Note: When using the Orbix configuration tool, `itconfigure`, the OTS Lite plug-in is deployed by default.

Loading the OTS Lite plug-in

As with the OTS plug-in the OTS Lite plug-in can be loaded in two ways:

1. Adding the plug-in name "ots_lite" to the orb_plugins configuration variable. For example: orb_plugins = [..., "ots", "ots_lite"];
2. Setting the initial_references:TransactionFactory:plugin configuration variable to "ots_lite". This causes the OTS Lite plug-in to be loaded by the OTS plug-in when resolve_initial_references("TransactionFactory") is called.

The server application does not need to load the OTS Lite plug-in except when standard interposition is used (that is, when the plugins:ots:interposition_style configuration variable is set to "standard"). In this case when the OTS plug-in imports the transaction from the client a transaction manager is required to create the sub-coordinated transaction.

This deployment should be used when the application only accesses on transactional resource (for example, updates a single database).

The Encina Transaction Manager

The Encina OTS Transaction Manager provides full recoverable 2PC transaction coordination implemented on top of the industry proven Encina Toolkit from IBM/Transarc.

There are two ways in which the Encina OTS may be used:

1. By configuring the itotstm service to load the Encina OTS plug-in.
2. By loading the Encina OTS plug-in directly into the application.

Configuring the OTS Encina Plug-In

Whether the OTS Encina plug-in is used in the itotstm service or directly in the application, there are a number of administration steps required to successfully use it.

Note: If you selected Distributed Transaction services when running the Orbix configuration tool, itconfigure, the administration steps outlined in this subsection are done automatically.

1. Two transient POAs must be created. These serve as namespace POAs off which the OTS Encina plug-in creates its persistent POAs. The first POA is called "iOTS" and the second is a child POA whose name is set by the plugins:ots_encina:namespace_poa. The default value of this configuration variable is "otstm" for the itotstm service and "Encina" for an application loading the plug-in. The POAs should be created using itadmin as follows:

```
itadmin poa create -transient -allowdynamic iOTS
itadmin poa create -transient -allowdynamic iOTS/otstm
```

2. The Encina OTS is fully recoverable and requires a transaction log to write the state of its transactions. Assuming the log file is to be located in `"/local/logs/ots.log"` the log is created and initialized using `itadmin` as follows:

```
itadmin encinalog create /local/logs/ots.log
itadmin encinalog init /local/logs/ots.log
```

The effect of initializing the log is to create a restart file. This a file that contains sufficient information for the OTS Encina plug-in to restart and includes the location of the transaction log. In this example, the restart file is called `/local/logs/ots_restart`. The name of the restart file must be passed to the OTS Encina plug-in by setting the `plugins:ots_encina:restart_file` configuration variable.

The minimum configuration required to load the OTS Encina plug-in into an applications is:

```
<app-scope> {
  initial_references:TransactionFactory:plugin = "ots_encina";
  plugins:ots_encina:namespace_poa = "<name>";
  plugins:ots_encina:restart_file = "<path>";
}
```

The itotstm Transaction Manager Service

The `itotstm` program is a standalone transaction manager service which can be configured to load any transaction manager plug-in. This section shows how it can be used along with the Encina OTS plug-in to provide 2PC transactions for an application. The `itotstm` service is deployed if you select the Distributed Transaction service when running the Orbix configuration tool, `itconfigure`.

Using itconfigure

If you select the Distributed Transaction service when running the Orbix configuration tool, `itconfigure`, the OTS Lite plug-in and the `itotstm` service are deployed. By default the OTS Lite plug-in is configured to be used by all clients and servers. To make use of the `itotstm` service, however, clients need to pick up the `initial_references:TransactionFactory:reference` configuration variable that is set in the `iona_services.otstm` client configuration scope. This can be done this by passing `"-ORBname iona_services.otstm.client"` to the `ORB_init()` operation or by adding a copy of the variable to the application's configuration scope.

Example client/server deployment

[Figure 19](#) shows a client/server deployment where the `itotstm` in conjunction with the OTS Encina plug-in is used to provide 2PC transaction management. Here, neither the client nor the server needs to load any transaction manager plug-in. Instead the client OTS is configured to pick up its transaction factory reference from the OTS Encina plug-in loaded into the `itotstm` standalone service.




Figure 19: Using the OTS Encina plug-in with the itotstm service

There are two parts to setting up such a deployment.

- Configuring the itotstm to load the OTS Encina plug-in.
- Configuring the OTS plug-in to pickup the reference to the OTS Encina transaction factory within the itotstm service.

Configuring itotstm

The itotstm service uses the configuration scope "otstm" by default. This can be changed by using a different ORB name using the -ORBname command line option. Configuring itotstm to load the OTS Encina plug-in can be done in two ways:

1. Adding the OTS plug-in name "ots_encina" to the orb_plugins configuration variable. For example, orb_plugins = [..., "ots", "ots_encina"];
2. Setting the initial_references:TransactionFactory:plugin configuration variable to the name of the OTS Encina plug-in "ots_encina".

Note that in both cases the orb_plugins configuration variable must contain "ots" since the OTS plug-in is required for synchronization objects.

The remainder of the otstm scope should contain the configuration necessary for the OTS Encina plug-in.

Configuring the OTS plug-in

Next the OTS plug-in loaded into the applications needs to pick up the transaction factory reference of the OTS Encina plug-in. Essentially this means setting the

`initial_references:TransactionFactory:reference` configuration variable in the applications configuration scope to any suitable reference. Three possible ways of achieving this are:

1. Get the OTS Encina plug-in to export its transaction factory reference to the name service and use a corbaname style URL for the initial reference. This is done by setting the `plugins:ots_encina:transaction_factory_ns_name` configuration variable to the name for the object reference in the name service. For example, if this is set to "ots/encina" a URL of the form "corbaname:rir:#ots/encina" can be used.
2. Get the itotstm to publish the transaction factory IOR to a file using the "prepare" and "-publish_to_file" command-line switches. Then use the IOR in the file as the transaction factory reference.

The deployment should be used when the application requires or might require full recoverable 2PC transactions. For example, the application make use of ore or more resource managers.

Loading the OTS Encina Plug-In into the Application

An alternative to loading the OTS Encina plug-in into the itotstm service is to load the plug-in directly into the application. This deployment is shown in [Figure 20](#).

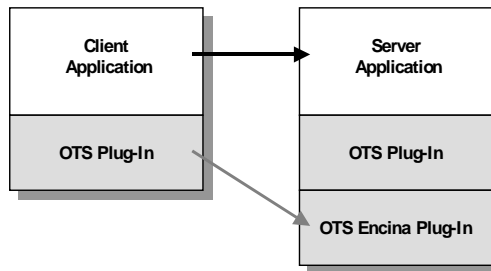


Figure 20: Loading the OTS Encina Plug-In into the Application

This deployment options should be used when the application requires full recoverable 2PC transactions and also wants to improve performance by eliminating some of the network messages that are necessary when the standalone itotstm service is used.

To configure this deployment, follow the instructions for configuring the OTS Encina plug-in, making sure the configuration is done within the application's scope.

OTS Management

This appendix describes the OTS server features that have been exposed for management. It explains all the managed entities, attributes, and operations. These can be managed using the Administrator management consoles.

Introduction to OTS Management

This section provides an introduction to the OTS management model and the Administrator management consoles.

OTS Management Model

Figure 21 shows the main components of the OTS management model.

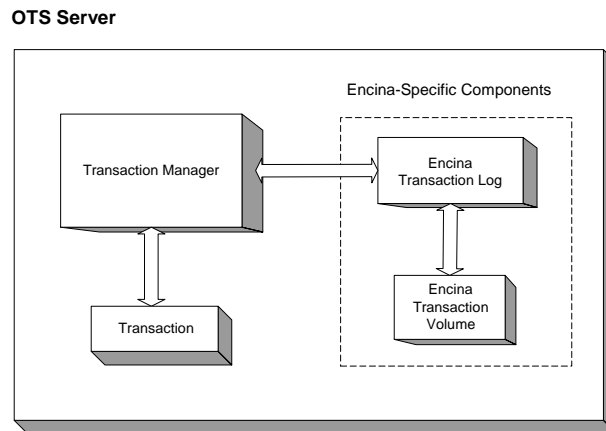


Figure 21: *OTS Management Model*

In Figure 21, the components on the left are common to both OTS Encina and OTS Lite. The components on the right apply to OTS Encina only.

In this model, each OTS server can have multiple Transactions and multiple Encina Transaction Volumes. However, each server can only have one Transaction Manager, and one Encina Transaction Log.

OTS Managed Entities

The following OTS server components have been instrumented for management:

- TransactionManager / Encina TransactionManager
- Transaction / Encina Transaction
- Encina Transaction Log
- Encina Volume

This means that these features can be managed using the Administrator management consoles.

Administrator

The Administrator is a set of tools that enables you to manage and configure server applications at runtime. The Administrator provides a graphical user interface known as the **Administrator Console**. This enables you to manage applications, configuration settings, event logging, and user roles.

The Administrator also provides a web browser interface known as the *Administrator Web Console*. The web console enables you to manage applications and event logging from anywhere, without the need for a lengthy download or installation.

For detailed information about the Orbix management tools, see the *Orbix Management User's Guide*.

Example Managed Entity

[Figure 22](#) shows an OTS **Encina Transaction Manager** running in the Administrator web console. It shows the attributes and operations that are exposed for this entity.

The next sections in this chapter describe the attributes and operations that are displayed for each of the OTS managed entities.




Figure 22: OTS Encina Transaction Manager Entity

TransactionManager Entity

This section describes the managed attributes and operations that are exposed for the TransactionManager and Encina TransactionManager entity. These attributes and operations are displayed in the Administrator Console.

TransactionManager Attributes

The managed attributes for the TransactionManager entity are shown in Table 9. These attributes apply to both OTS Encina and OTS Lite.

Table 9: TransactionManager Attributes (Sheet 1 of 2)

Attribute	Type	Description
Name	string	Name of the transaction manager.

Table 9: *TransactionManager Attributes (Sheet 2 of 2)*

Attribute	Type	Description
Supports 1PC	boolean	Whether the manager supports one-phase commit.
Supports 2PC	boolean	Whether the manager supports two-phase commit.
Active	long	Number of active transactions.
Completed	long	Number of completed transactions (since the server started).
Committed	long	Number of committed transactions (since the server started).
Aborted	long	Number of aborted transactions.
In Doubt	long	Number of transactions that are in doubt.
TPM	long	Number of transactions per minute.
TPM Peak	long	Maximum number of transactions per minute (since the server started).
TPM Peak Time	string	Time when the maximum transactions per minute was reached.
TPM Peak Average	double	Average transactions per minute (since server started).
Timeout	long	Default value for transaction timeout (same as the <code>default_transaction_timeout</code> configuration variable for the <code>ots_lite</code> and <code>ots_encina</code> plug-ins). This attribute is writable.
Transaction Log	hyperlink	Hyperlink to the Transaction Log entity (null for OTS Lite).

Encina TransactionManager Attributes

The additional managed attributes for the Encina TransactionManager entity are shown in [Table 10](#). These attributes apply to OTS Encina only.

Table 10: *Encina TransactionManager Attributes*

Attribute	Type	Description
Trace File	string	The file to which the trace output is written (stderr if the string is empty). This attribute is writable.
Trace On	boolean	Whether Encina tracing is enabled or not. This attribute is writable.
Trace Level bde	space-separated list of strings, where each element is one of the following: GLOBAL, EVENT, PARAM, NONE, INTERNAL_PARAM, INTERNAL_EVENT (for example, "EVENT PARAM")	These attributes specify the trace level for the corresponding Encina module (one of bde, log, restart, tran, util, vol, respectively). These attributes are writable.
Trace Level log		
Trace Level restart		
Trace Level tran		
Trace Level util		
Trace Level vol		

Encina TransactionManager Operations

The managed operations for the Encina TransactionManager entity are shown in [Table 11](#).

Table 11: *Encina TransactionManager Operations*

Operation	Parameters	Type	Description
dump	file name overwrite	string boolean	Writes the contents of the Encina trace buffer to the specified file. Depending on the value of the <code>overwrite</code> parameter, appends to an existing file, or overwrites it.

Transaction Entity

This section describes the managed attributes and operations exposed for the Transaction and Encina Transaction entity. These attributes and operations are displayed in the **Administrator Console**.

Transaction Attributes

The managed attributes for the Transaction entity are shown in [Table 12](#). These attributes apply to both OTS Encina and OTS Lite.

Table 12: *Transaction Attributes*

Attribute	Type	Description
Global TID	string	Global transaction identifier.
Timeout	boolean	Transaction-specific timeout.
Creation Time	boolean	Time when the transaction was created.
Status	long	CostTransactions::Status values.
Resources	long	Available resources for the transaction.

Encina Transaction Attributes

The additional managed attributes for the Encina Transaction entity are shown in [Table 13](#). These attributes apply to OTS Encina only.

Table 13: *Encina Transaction Attributes*

Attribute	Type	Description
Local TID	string	Local Encina-specific transaction identifier.

Transaction Operations

The managed operations for the Transaction entity are shown in [Table 14](#). These operations apply to both OTS Encina and OTS Lite.

Table 14: *Transaction Operations*

Operation	Parameter	Description
Rollback	none	Roll back the transaction.
Mark Rollback	none	Mark the transaction for being rolled back.
Commit	none	Commit the transaction.

Table 14: *Transaction Operations*

Operation	Parameter	Description
Remove Resource	string	Remove (unregister) the resource identified by the stringified object reference from the transaction. For example, this enables a transaction to complete if repeated attempts to deliver an outcome to a resource are failing.

Note: These operations are applicable to all transactions. In practice however, these operations will most likely fail for well-behaved transactions because of their short lifetime. They would only be applied in critical situations (for example, on a transaction with resource failures).

Encina Transaction Log Entity

This section describes the managed attributes and operations exposed for the Encina Transaction Log entity. These attributes and operations are displayed in the **Administrator Console**.

Encina Transaction Log Attributes

The managed attributes for the Encina Transaction Log are shown in [Table 15](#).

Table 15: *Encina Transaction Log Attributes (Sheet 1 of 2)*

Attribute	Type	Description
Name	string	Name of the log (always <code>tranLog</code> for the Encina Transaction Log).
Size	long	Size (in pages of 512 K).
Free	long	Free space (in pages).
Threshold	long	Percentage of used pages versus total pages that (when exceeded) cause a management event to be sent to the management service. This attribute is writable.
Check Interval	long	Interval (in seconds) for checking the amount of free space in the log. This attribute is writable.

Table 15: *Encina Transaction Log Attributes (Sheet 2 of 2)*

Attribute	Type	Description
Growth	long	Difference of free space in the log at beginning and end of the last check interval.
Average Growth	double	Average of the growth rate in the lifetime of the OTS server.
Archive Device	string	File name of the archive device of the log.
Mirrors	list of hyperlinks	List of hyperlinks to Encina Volume entities.

Encina Transaction Log Operations

The managed operations for the Encina Transaction Log are shown in [Table 16](#).

Table 16: *Encina Transaction Log Operations*

Operation	Parameters	Description
Expand	none	Expands the log to maximum possible size. This is necessary to avail of the increased disk space after a mirror has been added.
Add Mirror	string	Creates a new physical volume backed up by the specified disk, and adds it to the list of volumes currently mirroring the transaction log. The raw partition or file specified by the string parameter must exist. You can create files using the <code>itadmin</code> tool.

Encina Volume Entity

This section describes the managed attributes and operations exposed for the Encina (Physical) Volume entity. These attributes and operations are displayed in the **Administrator Console**.

Encina Volume Attributes

The managed attributes for the Encina (Physical) Volume entity are shown in [Table 17](#).

Table 17: *Encina (Physical) Volume Attributes*

Attribute	Type	Description
Name	string	Logical name of the physical volume.
Disks	list of strings	List of fully qualified file or raw partition names for the different disks that backup the volume.

Encina Volume Operations

The managed operations for the Encina (Physical) Volume are shown in [Table 18](#).

Table 18: *Encina (Physical) Volume Operations*

Operation	Parameter	Description
Remove	none	Removes this physical volume.
Add Disk	string	Adds the specified disk to the physical volume. The raw partition or file must exist. You can create files using the <code>itadmin</code> tool.

Management Events

The following OTS events are logged with the Administrator management service:

- The heuristic outcome of a transaction.
This event includes the `otid` and the heuristic outcome type.
- When the used space in the transaction log exceeds the threshold.

This event includes the actual percentage of used versus the total number of pages in the transaction log.

Glossary

A

administration

All aspects of installing, configuring, deploying, monitoring, and managing a system.

C

client

An application (process) that typically runs on a desktop and requests services from other applications that often run on different machines (known as server processes). In CORBA, a client is a program that requests services from CORBA objects.

configuration

A specific arrangement of system elements and settings.

configuration domain

Contains all the configuration information that Orbix ORBs, services and applications use. Defines a set of common configuration settings that specify available services and control ORB behavior. This information consists of configuration variables and their values. Configuration domain data can be implemented and maintained in a centralised Orbix configuration repository or as a set of files distributed among domain hosts. Configuration domains let you organise ORBs into manageable groups, thereby bringing scalability and ease of use to the largest environments. See also [configuration file](#) and [configuration repository](#).

configuration file

A file that contains configuration information for Orbix components within a specific configuration domain. See also [configuration domain](#).

configuration repository

A centralised store of configuration information for all Orbix components within a specific configuration domain. See also [configuration domain](#).

configuration scope

Orbix configuration is divided into scopes. These are typically organized into a root scope and a hierarchy of nested scopes, the fully-qualified names of which map directly to ORB names. By organising configuration properties into scopes, different settings can be provided for individual ORBs, or common settings for groups of ORB. Orbix services have their own configuration scopes.

CORBA

Common Object Request Broker Architecture. An open standard that enables objects to communicate with one another regardless of what programming language they are written in, or what operating system they run on. The CORBA specification is produced and maintained by the OMG. See also [OMG](#).

CORBA objects

Self-contained software entities that consist of both data and the procedures to manipulate that data. Can be implemented in any programming language that CORBA supports, such as C++ and Java.

D

deployment

The process of distributing a configuration or system element into an environment.

E

event

The occurrence of a condition or state change, or the availability of some information that is of interest to one or more modules in a system. Suppliers generate events and consumers subscribe to receive them.

I

IDL

Interface Definition Language. The CORBA standard declarative language that allows a programmer to define interfaces to CORBA objects. An IDL file defines the public API that CORBA objects expose in a server application. Clients use these interfaces to access server objects across a network. IDL interfaces are independent of operating systems and programming languages.

IIOP

Internet Inter-ORB Protocol. The CORBA standard messaging protocol, defined by the OMG, for communications between ORBs and distributed applications. IIOP is defined as a protocol layer above the transport layer, TCP/IP.

installation

The placement of software on a computer. Installation does not include configuration unless a default configuration is supplied.

Interface Definition Language

See [IDL](#).

invocation

A request issued on an already active software component.

IOR

Interoperable Object Reference. See [object reference](#).

M

management

To direct or control the use of a system or component. Sometimes used in a more general way meaning the same as Administration. management console

N

node daemon

Starts, monitors, and manages servers on a host machine. Every machine that runs a server must run a node daemon.

O

object reference

Uniquely identifies a local or remote object instance. Can be stored in a CORBA naming service, in a file or in a URL. The contact details that a client application uses to communicate with a CORBA object. Also known as interoperable object reference (IOR) or proxy.

object transaction service

See [Orbix OTS](#).

OMG

Object Management Group. An open membership, not-for-profit consortium that produces and maintains computer industry specifications for interoperable enterprise applications, including CORBA. See www.omg.com.

ORB

Object Request Broker. Manages the interaction between clients and servers, using the Internet Inter-ORB Protocol (IIOP). Enables clients to make requests and receive replies from servers in a distributed computer environment. Key component in CORBA.

Orbix OTS

Object Transaction Service. An implementation of the OMG Transaction Service Specification. Provides interfaces to manage the demarcation of transactions and the propagation of transaction contexts.

POA

Portable Object Adapter. Maps object references to their concrete implementations in a server. Creates and manages object references to all objects used by an application, manages object state, and provides the infrastructure to support persistent objects and the portability of object implementations between different ORB products. Can be transient or persistent.

protocol

Format for the layout of messages sent over a network.

S

server

A program that provides services to clients. CORBA servers act as containers for CORBA objects, allowing clients to access those objects using IDL interfaces.

T

transaction manager

Manages global transactions on behalf of application programs. A transaction manager coordinates commands from application programs and resource managers to start and complete global transactions. When an application completes a transaction, either with a commit or rollback request, the transaction manager communicates the outcome with each resource manager.

Index

Numerics

- 1PC 5, 98
 - operation 108
 - Orbix 3 OTS 115
 - OTS Lite 120
 - OTS Lite deployment 124
 - resource objects 89
 - successful 99
- 2PC 96
 - ACID properties 4
 - commit() 106
 - operations 108
 - OTS Encina 120
 - OTS plug-in configuration 130
 - otstf transaction manager 118
 - prepare() 105
 - resource objects 89, 92
 - rollback() 107
 - rollbacks 96
 - successful 97
 - transaction management 128
 - transaction manager 115

A

- ADAPTS policy 42
- AUTOMATIC policy 52
 - code example 45
 - default_ots_policy 134
 - InvalidPolicy exception 44
 - Orbix 3 OTS 116
 - ots_v11_policy 133
 - POA policies 19
 - policy mappings 55
 - SERVER_SIDE policy 53
 - Transactional objects 113
 - using 50
- after_completion() 77
- agent_ior_file 139
- allow_registration_after_rollback_only
 - OTS Encina 137
 - OTS Lite 136
- Allows_either TransactionPolicy 55
- Allows_unshared TransactionPolicy 55
- asynchronous XA 62
- AUTOMATIC policy 52
 - policy mappings 55
 - SEVER_SIDE policy 53
- automatic transactions 52

B

- backup_restart_file 138
- before_completion 72
 - after_completion 77

- before_completion() 76
- BeforeCompletionCallback interface 12
- BeforeCompletionCallback objects,
 - registering 72
- begin() 9
 - current interface 28
 - invoking 17
 - JIT transactions 53
 - nested transactions 33
 - new transactions 30
- bindings 25

C

- C API
 - resource manager integration 12
 - XA specification 60
- client_binding_list 25
- client OTS policy 40
- close_string 141
 - Orbix namespace 66
- commit() 9
 - 2PC 96
 - code example 36
 - exceptions 17
 - functions 106
 - heuristic exceptions 32
 - heuristic outcomes 100
 - invoking 17
 - JIT transactions 53
 - new transactions 31
 - resource failure 101
 - resource interface 88
 - resource objects 91
- commit_on_completion_of_next_call() 54
- commit_one_phase() 98
 - invoking 107
- Connector interface 12
- Control interface 10
- Coordinator interface 10
 - identity operations 79
 - relationship operations 83
 - status operations 81
- CosTransactions.hh 16
- create()
 - Control interface 10
 - new top-level transactions 85
 - timeouts 36
- create_POA() 45
 - exceptions 55
- create_policy() 45
- create_resource_manager() 21
 - calling 22
- CurrentConnection interface 12

- CurrentConnection object 22
- Current interface 9, 10
 - commit_on_completion_of_next_call() 54
 - definition 28
 - Transaction Factory 9
- Current object
 - nested transactions 33
 - transaction demarcation 16

D

- database access 22
- default_ots_policy 134
- default_transaction_policy 134
- default_transaction_timeout 133
- direct mode transactions 11
- direct_persistence 137

E

- EITHER policy 44
 - policy mappings 55
- Encina plug-In
 - configuring 126
 - loading 130
- Encina plug-in
 - configuring 129
 - itotstm service 128
- Encina Transaction Manager 126
- exceptions
 - forget() 108
 - heuristic 100, 106
 - HeuristicCommit 107
 - HeuristicMixed and HeuristicHazard 32
 - inactive 94
 - InvalidControl 35
 - InvalidPolicy 44, 55
 - INVALID_TRANSACTION 42, 43
 - NotPrepared 103
 - NoTransaction 32, 36
 - OBJECT_NOT_EXIST 103
 - See Also system exceptions
 - TRANSACTION_MODE 44
 - TRANSACTION_REQUIRED 42
 - TRANSACTION_ROLLBACK 52
 - TRANSACTION_ROLLEDBACK 17, 31, 99
 - user 106, 107
- explicit mode transactions 11
- explicit propagation
 - IDL 57
 - TransactionFactory reference 40

F

- FORBIDS policy 20, 42
 - InvalidPolicy exception 44
- forget() 108

G

- get_control() 35
 - real transactions 53
- get_parent_status() 82

- get_status() 35
 - Current interface return values 81
- get_timeout() 34
- get_top_level_status() 82
- get_transaction_name() 35, 79
 - real transactions 53
- get_txcontext() 80
 - PropagationContext 86
- global_namespace_poa 137

H

- hash_top_level_transaction() 80
- hash_transaction() 79
 - maintaining data 80
 - tracking resource objects 92
- HeuristicCommit exception 100, 107
- heuristic exception 100
- HeuristicMixed and HeuristicHazard exceptions 32
- HeuristicRollbackException 106
- heuristics outcomes 99

I

- implicit propagation policy 40
- Inactive exception 94
- indirect(implicit) mode transactions 11
- indirect mode transactions 11
- initial_disk 139
- initial_disk_size 139
- initial_references:OTSMangement:plugin 132
- initial_references:TransactionFactory:plugin 132
- initial_references:TransactionFactory:reference 132
- interposition_style 133
- InvalidControl exception 35
- InvalidPolicy exception 44
 - create_POA() 55
- INVALID_TRANSACTION exception
 - FORBIDS policy 42
 - PREVENTS policy value 43
- InvocationPolicy 40
 - transaction models 41
 - values 44
- is_ancestor_transaction() 83
- is_descendant_transaction() 84
- is_related_transaction() 83
- is_same_transaction() 79
 - description 83
 - maintaining data 80
 - tracking resource objects 92
- is_top_level_transaction() 84
- itadmin
 - transient POAs 126
- itotstm
 - configuring 129
 - service 126
 - transaction manager service 128
- itotstm service 126

J

JIT transaction creation 53
jit_transactions 134

L

Lite plug-in
 deployment 124
 loading 124
 transaction manager 115
log_check_interval 140
log_threshold 139

M

max_resource_failures 140
Multi-threading 63

N

namespace_poa 137
nested transaction families 82
nested transactions 33
NonTxTargetPolicy 40
 default value 49
 steps for using 47
 values 43
non_tx_target_policy 142
NotPrepared exception 103
NoTransaction exception 32, 36

O

OBJECT_NOT_EXIST exception 103
one-phase-commit (1PC) protocol See 1PC
open_string 141
open-string specification 21
Oracle database example 21
orbix/cos_transactions.hh 54
orbix/xa.hh 21
Orbix 3 OTS applications 115
OrbixOTS.INTEROP variable 117
orb_name 135, 137
orb_plugins configuration variable 129
otid field 92
otid_format_id 136, 138
OTS Application example
 funds transfer 14
OTS application example
 completion steps 15
ots_encina Namespace Variables 137
OTS Encina See Under Encina
OTS Interfaces 10
ots_lite Namespace Variables 135
OTS Lite See Lite
OTS plug-in
 loading 122
OTS plug-ins 120
 deployment scenarios 123
 loading 25
 purpose of 122
OTSPolicies, Orbix specific 52
OTSPolicy 40

 creating objects 45
 values 19, 42
OTS Resource interface 9
otstf
 bypassing 116
 server 116
OTS transaction modes 11
ots_v11_policy 133

P

PERMIT NonTxTargetPolicy 117
PERMIT policy 113
 value 43
PERSISTENT lifespan policy 92
ping_period 141
plugins:ots_encina:orb_name 137
plugins:ots_encina Namespace
 Variables 137
plugins:ots_lite Namespace Variables 135
poa_name 141
POA policies 19
 transaction propagation 40
POAs and Encina plug-in 126
PolicyCurrent object 47
PolicyManager object 47
prepare() 96, 105
PREVENT policy value 43
propagate_separate_tid_optimization 13
 4
PropagationContext structure 85
propagation policies 40

R

RecoveryCoordinator interface 10, 103
recovery coordinator object 94
recreate() 85
register_resource() 24, 93
register_synchronization() 77
replay_completion() 94, 102
 usage model 104
 using 109
REQUIRES policy value 19
resolve_initial_references() 16
 transaction factory object 36
 XAConnector 22
Resource interface 9, 10
resource interface operations 24
Resource interface transaction
 operations 88
ResourceManager interface 12
ResourceManager object 22
resource managers, XA compliant 12
resource objects
 creating 92
 failure/recovery 101
 implementation checklist 108
 implementing servants 91
 protocols supported 95
 registering 93
 tracking 92
 usage model 89

- ResourcePOA class 91
- resource_retry_timeout 138
- restart_file 138
- resume() 34
- rmid 141
- rollback() 96
 - current transactions 33
 - invoking 18
 - occasions when called 107
 - transaction demarcation 9
 - user exceptions 107
- rollback_only() 33, 76
 - real transactions 53
- rollback_only_on_system_ex 134
- rollbacks, reasons for 95

S

- server_binding_list 25
- SERVER_SIDE policy value 52, 134
 - JIT 53
- set_policy_overrides() 47
- set_timeout() 34
- SHARED policy 44
- shared transaction model 41
- StatusActive value 81
- StatusCommitted value 81
- StatusCommitting value 81
- StatusMarkedRollback 81
- StatusMarkedRollback value 81
- StatusNoTransaction value 81
- StatusPrepared value 81
- StatusPreparing value 81
- StatusRolledBack value 81
- StatusRollingBack value 81
- StatusUnknown value 81
- SubtransactionAwareResource
 - interface 10
- superior_ping_timeout 136
- support_ots_v11 133
- supports_async_rollback 141
- suspend() 34
 - real transactions 53
- Synchronization interface 11, 76
- synchronization objects 78
- system exceptions
 - effects of raising 76
 - INVALID_TRANSACTION 43
 - OBJECT_NOT_EXIST 103
 - rollback_only_on_system_ex 134
 - TRANSACTION_MODE 44
 - TRANSACTION_REQUIRED 42
 - TRANSACTION_ROLLEDBACK 17, 31, 52, 101, 108

T

- Terminator interface 11, 36
- thread_model configuration variable 21
- threads 29
- timeouts 34, 96
- trace_comp 139
- trace_file 139

- trace_on 139
- TransactionalObject interface 11, 15
 - Orbix support 113
- transaction coordinator failure 102
- transaction demarcation 9
- TransactionFactory interface 11
 - Current interface 9
 - declaring 85
- transaction_factory_name 133
 - ots_encina 138
 - ots_lite 135
- transaction_factory_ns_name 138
- transaction family 33
- transaction identifier 92
- Transaction interface 8
 - resource manager integration 9
- transaction management
 - OTS interfaces 9
- TransactionManager 4
- TRANSACTION_MODE exception
 - SHARED policy value 44
- transaction modes 11
- TransactionPolicies 112
- TransactionPolicy
 - migrating from 55
- transaction propagation 9
- TRANSACTION_REQUIRED exception 42
- transaction rollbacks, reasons for 95
- TRANSACTION_ROLLEDBACK
 - exception 17, 31, 52, 101, 108
- transactions 2
 - automatic 52
 - creating 30
 - creating new 17
 - database access steps 22
 - example 2
 - maintaining data 80
 - nested 33
 - orbix support 2
 - POA policies 19
 - propagation policies 40
 - properties 3
 - suspending/resuming 34
 - threads 29
- transaction_timeout_period 135, 138
- two-phase-commit (2PC) protocol See 2PC

U

- UNSHARED policy value 44
- unshared transaction model 41
- use_internal_orb 135, 138
- use_raw_disk 139
- user exceptions 106, 107
- USER_ID ID assignment policy 92, 108

V

- VoteCommit value 96
 - using 109
- VoteReadOnly value 96, 105
 - using 109

VoteRollback value 105

X

X/Open XA interface 12

xa_close() 12, 61

xa_commit() 12, 61

xa_complete() 62

XA-compliant database 23

xa_end() 12, 62

xa_forget() 12, 61

XA interfaces 12

xa_open() 12, 61

 open-string 21

xaosw 22, 61

xa_prepare() 12, 61

xa_recover() 12, 62

XA resource manager

 OTS managed transactions

 integration 21

XA Resource Manager Variables 141

xa_rollback() 12, 61

xa_start() 12, 62

xa_switch_t instance 61

XID transaction identifier format 86

