



Orbix 6.3.7

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.

**Migrating from Orbix
3.3 to Orbix 6.3**

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

<http://www.microfocus.com>

Copyright © Micro Focus 2014. 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.

2014-06-09

Contents

Preface	1
Contacting Micro Focus	2
 Part I Overview of Migration	
Introduction	7
Advantages of Orbix 6.3.....	7
Migration Resources	8
Migration Options.....	8
Migrating to Orbix 6.3	8
Mixed Deployment	9
 Part II Migrating to Orbix 6.3	
IDL Migration	13
The opaque Type	13
The Principal Type.....	13
Client Migration	15
Replacing the _bind() Function.....	15
Callback Objects	17
IDL-to-C++ Mapping	18
System Exception Semantics	19
Dynamic Invocation Interface	19
Server Migration	21
Function Signatures.....	21
Object IDs versus Markers	21
CORBA Objects versus Servant Objects	22
BOA to POA Migration	22
Creating an Object Adapter	23
Defining an Implementation Class	23
Creating and Activating a CORBA Object	24
Migrating Proprietary Orbix 3 Features	27
Orbix 3 Locator.....	27
Filters.....	29
Request Logging	29
Piggybacking Data on a Request	29
Multi-Threaded Request Processing	30
Accessing the Client's TCP/IP Details	31
Security Using an Authentication Filter.....	31
Loaders	31
Smart Proxies.....	32
Transformers.....	34

I/O Callbacks	34
Connection Management	35
Session Management	36
CORBA Services	39
Interface Repository	39
Naming Service	39
Notification Service	40
CORBA Specification Updates	40
Quality of Service Properties	42
Configuration/Administration Changes	43
Deprecated Features	44
SSL/TLS Toolkit	45
Changes to the Programming Interfaces	45
Configuration and Administration	47
Migrating Certificate and Private Key Files	50
Administration	53
Orbix Daemons	53
POA Names	53
Command-Line Administration Tools	53
Activation Modes	55
Part III Interoperability	
Configuring for Interoperability	59
Interoperability Overview	59
Launch and Invoke Rights	60
GIOP Versions	62
IDL Issues	65
Using the #pragma Prefix	65
Use of #pragma ID in IDL	67
Fixed Data Type and Interoperability	68
Use of wchar and wstring	70
C++ Keywords as Operation Names	70
Exceptions	73
Orbix 3.3 C++ Edition—System Exceptions	73
New Semantics and Old Semantics	73
The INV_OBJREF and OBJECT_NOT_EXIST Exceptions	75
The TRANSIENT and COMM_FAILURE Exceptions	76
Orbix 3.3 C++ Edition and Orbix 6.3	76
Orbix 3.3 Java Edition—System Exceptions	78
New Semantics and Old Semantics	78
The INV_OBJREF and OBJECT_NOT_EXIST Exceptions	79
The TRANSIENT and COMM_FAILURE Exceptions	79
Orbix 3.3 Java Edition and Orbix 6.3	80
FILTER_SUPPRESS Exception	81
Dynamic Invocation Interface and User Exceptions	82
Dynamic Invocation Interface and LOCATION_FORWARD	83
Services	85

The Orbix 6.3 Interoperable Naming Service.....	85
Interface Repository Interoperability.....	90
SSL/TLS Toolkit Interoperability.....	90
High Availability and Orbix 3.3 Clients.....	91
Connection Management	93
Orbix 6.3 Active Connection Management.....	93
Callbacks and Bidirectional GIOP.....	93
Setting the Listen Queue Size in Orbix 3.3 C++ Edition.....	94
Multiple LOCATION_FORWARD.....	95
Codeset Negotiation	97
Introduction to Codeset Negotiation	97
Configuring Codeset Negotiation	97
Default Codesets.....	98
Configuring Legacy Behavior	101
Index.....	103

Preface

This document explains how to migrate applications from the Orbix and OrbixWeb products, which conform to CORBA 2.1, to Orbix 6.3, which conforms to CORBA 2.6.

Audience

This document is aimed at *C++ or Java programmers* who are already familiar with Orbix or OrbixWeb products and who now want to migrate all or part of a system to use Orbix 6.3.

Parts of this document are relevant also to *administrators* familiar with Orbix and OrbixWeb administration. See [“Administration” on page 53](#) and [“Configuring for Interoperability” on page 59](#).

Organization of this guide

This guide is divided as follows:

Part I “Overview of Migration”

This part briefly discusses the advantages of migrating and the options for your migration strategy.

Part II “Migrating to Orbix 6.3”

This part explains how to migrate client and server source (in C++ or Java) to Orbix 6.3. For each of the features that have been modified or removed from Orbix 6.3, relative to the features supported by Orbix 3 and OrbixWeb 3, this part discusses the replacement features offered by Orbix 6.3.

Part III “Interoperability”

This part discusses the issues that affect a mixed deployment of interoperating Orbix 3, OrbixWeb 3 and Orbix 6.3 applications. With appropriate customization of the ORB configuration, you can obtain an optimum level of compatibility between the various applications in your system.

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 <code>CORBA::Object</code> class. Constant width paragraphs represent code examples or information a system displays on the screen. For example: <pre>#include <stdio.h></pre>
----------------	---

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 or Windows 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 WebSync service, where you can download fixes and documentation updates.
- The Knowledge Base, a large collection of product tips and workarounds.
- Examples and Utilities, including demos and additional product documentation.

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

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.

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>

Part I

Overview of Migration

In this part

This part contains the following chapter:

Introduction	page 7
------------------------------	------------------------

Introduction

The newest generation of Orbix provides significant advances over the previous generation of products.

Advantages of Orbix 6.3

The recommended path for customers upgrading to a new version of Orbix is to move to Orbix 6.3. The extra features offered by Orbix can be divided into the following categories:

- [CORBA 2.6-compliant features](#).
- [Unique features](#).

CORBA 2.6-compliant features

Because Orbix 6.3 contains a CORBA 2.6-compliant ORB, it offers the following advantages over Orbix 2.x (all minor versions of Orbix 2) and Orbix 3.x (all minor versions of Orbix 3):

- Portable interceptor support.
- Codeset negotiation support.
- Value type support.
- Asynchronous method invocation (AMI) support.
- Persistent State Service (PSS) support.
- Dynamic any support.

Unique features

Orbix 6.3 also offers some unique benefits over other commercial ORB implementations, including:

- ORB extensibility using Adaptive Runtime Technology (ART).
Orbix 6.3 has a modular structure built on a micro-kernel architecture. Required ORB modules, ORB plug-ins, are specified in a configuration file and loaded at runtime, as the application starts up. The advantage of this approach is that new ORB functionality can be dynamically loaded into an Orbix application without rebuilding the application.
- Improved performance.
The performance of Orbix 6.3 has been optimized, resulting in performance that is faster than Orbix 3.x and OrbixWeb 3.x in every respect.
- Advanced deployment and configuration.
Orbix 6.3 supports a flexible model for the deployment of distributed applications. Applications can be grouped into configuration domains and organized either as file-based configuration domains or as configuration repository-based configuration domains.
- Rapid application development using the Orbix code generation toolkit.

The code generation toolkit is an extension to the IDL compiler that generates a working application prototype—based on your application IDL—in a matter of seconds.

Migration Resources

Overview of resources

Micro Focus is committed to assisting you with your migration effort, to ensure that it proceeds as easily and rapidly as possible. The following resources are currently available:

- This migration and interoperability guide.
This technical document provides detailed guidance on converting source code to Orbix 6.3. The document aims to provide comprehensive coverage of migration issues, and to demonstrate how features supported in earlier Orbix versions can be mapped to Application Server Platform features.
- Micro Focus Orbix 6.x Upgrade Assessment.
For customers on Orbix 2000, Orbix E2A ASP or versions of Orbix predating Orbix 3.3 one may consider the Orbix 6.x Upgrade Assessment, ensuring the application of best practices and access to the latest updated functionality.:
<http://www.microfocus.com/products/corba/services/orbixupgradeassessment.aspx>.

Migration Options

The basic alternatives for migrating a distributed application to Orbix are to migrate the whole application at once, or to perform the migration gradually, replacing parts of the application piece by piece. For the latter option (gradual migration), you will end up with a mixed deployment consisting of Orbix and older Orbix products.

Migrating to Orbix 6.3

The CORBA 2.6 specification, on which the Orbix 6.3 ORB is based, standardizes almost every aspect of CORBA programming. Migrating your source code to Application Server Platform, therefore, represents a valuable investment because your code will be based on a stable, highly standardized programming interface.

Client side

On the client side, the main issue for migration is that the Orbix `_bind()` function is not supported in Orbix 6.3. The CORBA Naming Service is now the recommended mechanism for establishing contact with CORBA servers.

Server side

On the server side, the basic object adapter (BOA) must be replaced by the portable object adapter (POA). This is one of the major differences between the CORBA 2.1 and the CORBA 2.6 specifications. The POA is much more tightly specified than the old BOA; hence server code based on the POA is well standardized.

Proprietary features

Orbix 3.x and OrbixWeb 3.x support a range of proprietary features not covered by the CORBA standard—for example, the Orbix locator, filters, loaders, smart proxies, transformers and I/O callbacks. When migrating to Orbix 6.3, the proprietary features must be replaced by standard CORBA 2.6 features. This migration guide details how each of the proprietary features can be replaced by equivalent Orbix 6.3 functionality.

Further details

The details of migrating to Orbix 6.3 are discussed in Part II of this guide. See [“Migrating to Orbix 6.3” on page 11](#).

Mixed Deployment

Mixed Deployment is appropriate when a number of CORBA applications are in deployment simultaneously. Some applications might be upgraded to use Orbix 6.3 whilst others continue to use Orbix 3.x and OrbixWeb 3.x. This kind of mixed environment requires on-the-wire compatibility between the generation 3 products and Orbix 6.3. Extensive testing has been done to ensure interoperability with Orbix 6.3.

On-the-wire interoperability

Both Orbix 3.3 and Orbix 6.3 have been modified to achieve an optimum level of on-the-wire compatibility between the two products.

Further details

Interoperability is discussed in Part III of this guide. See [“Interoperability” on page 57](#).

Part II

Migrating to Orbix 6.3

In this part

This part contains the following chapters:

IDL Migration	page 13
Client Migration	page 15
Server Migration	page 21
Migrating Proprietary Orbix 3 Features	page 27
CORBA Services	page 39
Administration	page 53

IDL Migration

This chapter discusses the Orbix 3.x IDL features that are not available in Orbix 6.3.

The opaque Type

Migrating to Orbix 6.3

The object-by-value (OBV) specification, introduced in CORBA 2.3 and supported in Orbix 6.3, replaces opaques.

The Principal Type

Principal IDL type

The CORBA specification deprecates the `Principal` IDL type; therefore the `Principal` IDL type is not supported by Orbix 6.3.

Interoperability

Orbix 6.3 has some limited on-the-wire support for the `Principal` type, to support interoperability with Orbix 3.x applications.

See ["Launch and Invoke Rights" on page 60](#).

Client Migration

Migration of client code from Orbix 3 to Orbix 6.3 is generally straightforward, because relatively few changes have been made to the client-side API.

Replacing the `_bind()` Function

The `_bind()` function is not supported in Orbix 6.3. All calls to `_bind()` must be replaced by one of the following:

- [CORBA Naming Service](#).
- [CORBA Trader Service](#).
- [Object-to-string conversion](#).
- [corbaloc URL](#).
- [ORB::resolve_initial_references\(\)](#).

CORBA Naming Service

The naming service is the recommended replacement for `_bind()` in most applications. Migration to the naming service is straightforward on the client side. The triplet of `markerName`, `serverName`, and `hostName`, used by the `_bind()` function to locate an object, is replaced by a simple name in the naming service.

When using the naming service, an object's name is an abstraction of the object location and the actual location details are stored in the naming service. Object names are resolved using these steps:

1. An initial reference to the naming service is obtained by calling `resolve_initial_references()` with `NameService` as its argument.
2. The client uses the naming service reference to resolve the names of CORBA objects, receiving object references in return.

Orbix 6.3 supports the CORBA Interoperable Naming Service, which is backward-compatible with the old CORBA Naming Service and adds support for stringified names.

CORBA Trader Service

The Orbix 6.3 trader service provides advanced capabilities for object location and discovery. Unlike the Orbix Naming Service where an object is located by name, an object in the Trading Service does not have a name. Rather, a server advertises an object in the Trading Service based on the kind of service provided by the object. A client locates objects of interest by asking the Trading Service to find all objects that provide a particular service. The client can further restrict the search to select only those objects with particular characteristics.

Object-to-string conversion

CORBA offers two CORBA-compliant conversion functions:

```
CORBA::ORB::object_to_string()
```

```
CORBA::ORB::string_to_object()
```

These functions allow you to convert an object reference to and from the stringified interoperable object reference (stringified IOR) format. These functions enable a CORBA object to be located as follows:

1. A server generates a stringified IOR by calling `CORBA::ORB::object_to_string()`.
2. The server passes the stringified IOR to the client (for example, by writing the string to a file).
3. The client reads the stringified IOR from the file and converts it back to an object reference, using `CORBA::ORB::string_to_object()`.

Because they are not scalable, these functions are generally not useful in a large-scale CORBA system. Use them only to build initial prototypes or proof-of-concept applications.

corbaloc URL

A *corbaloc URL* is a form of human-readable stringified object reference. If you are migrating your clients to Orbix 6.3 but leaving your servers as Orbix 3.3 applications, the corbaloc URL offers a convenient replacement for `_bind()`.

To access an object in an Orbix 3.3 server from an Orbix 6.3 client using a corbaloc URL, perform the following steps:

1. Obtain the object key, *ObjectKey*, for the object in question, as follows:
 - i. Get the Orbix 3.3 server to print out the stringified IOR using, for example, the `CORBA::ORB::object_to_string()` operation. The result is a string of the form `IOR:00...`
 - ii. Use the Orbix 6.3 `iordump` utility to parse the stringified IOR. Copy the string that represents the object key field, *ObjectKey*.
2. Construct a corbaloc URL of the following form:

```
corbaloc:iiop:1.0@DaemonHost:DaemonPort/ObjectKey%00
```

Where *DaemonHost* and *DaemonPort* are the Orbix daemon's host and port respectively. A null character, `%00`, is appended to the end of the *ObjectKey* string because Orbix 3.3 applications expect object key strings to be terminated by a null character.
3. In the source code of the Orbix 6.3 client, use the `CORBA::ORB::string_to_object()` operation to convert the corbaloc URL to an object reference.

The general form of a corbaloc URL for this case is as follows:

```
corbaloc:iiop:GIOPVersion@Host:Port/Orbix3ObjectKey%00
```

Where the components of the corbaloc URL are:

- *GIOPVersion*—The maximum GIOP version acceptable to the server. Can be either 1.0 or 1.1.

- *Host* and *Port*—The daemon's (or server's) host and port. The *Host* can either be a DNS host name or an IP address in dotted decimal format.

The *Orbix3ObjectKey* has the following general form:

```
: \Host:SvrName:Marker:IFRSvrName:InterfaceName%00
```

Where the components of the Orbix 3 object key are:

- *Host*—The server host. The *Host* can either be a DNS host name or an IP address in dotted decimal format.
- *SvrName*—The server name of the Orbix 3.3 server.
- *Marker*—The CORBA object's marker.
- *IFRSvrName*—Can be either IR or IFR.
- *InterfaceName*—The object's IDL interface name.

WARNING: Constructing an Orbix 3.3 object key directly based on the preceding format does *not* always work because some versions of Orbix impose extra restrictions on the object key format. Extracting the object key from the server-generated IOR is a more reliable approach.

If you encounter any difficulties with using corbaloc URLs, please contact support@iona.com.

ORB::resolve_initial_references()

The CORBA::ORB::resolve_initial_references() operation provides a mechanism for obtaining references to basic CORBA objects (for example, the naming service, the interface repository, and so on).

Orbix 6.3 allows the resolve_initial_references() mechanism to be extended. For example, to access the BankApplication service using resolve_initial_references(), simply add the following variable to the Orbix 6.3 configuration:

```
# Orbix 6.3 Configuration File
initial_references:BankApplication:reference =
"IOR:010347923849..."
```

Use this mechanism sparingly. The OMG defines the intended behavior of resolve_initial_references() and the arguments that can be passed to it. A name that you choose now might later be reserved by the OMG. It is generally better to use the naming service to obtain initial object references for application-level objects.

Callback Objects

POA policies for callback objects

Callback objects must live in a POA, like any other CORBA object; hence, there are certain similarities between a server and a client with callbacks. The most sensible POA policies for a POA that manages callback objects are shown in [Table 1](#).

Table 1: *POA Policies for Callback Objects*

Policy Type	Policy Value
Lifespan	TRANSIENT ^a
ID Assignment	SYSTEM_ID
Servant Retention	RETAIN
Request Processing	USE_ACTIVE_OBJECT_MAP_ONLY

a. By choosing a TRANSIENT lifespan policy, you remove the need to register the client with an Orbix 6.3 locator daemon.

These policies allow for easy management of callback objects and an easy upgrade path. Callback objects offer one of the few cases where the root POA has reasonable policies, provided the client is multi-threaded (as it normally is for callbacks).

IDL-to-C++ Mapping

The definition of the IDL-to-C++ mapping has changed little going from Orbix 3.x to Orbix 6.3 (apart from some extensions to support valuetypes). Two notable changes are:

- [The CORBA::Any Type](#).
- [The CORBA::Environment Parameter](#).

The CORBA::Any Type

In Orbix 6.3, it is not necessary to use the type-unsafe interface to `Any`. Recent revisions to the CORBA specification have filled the gaps in the IDL-to-C++ mapping that made these functions necessary. That is, the following functions are deprecated in Orbix 6.3:

```
// C++
// CORBA::Any Constructor.
Any(
    CORBA::TypeCode_ptr tc,
    void* value,
    CORBA::Boolean release = 0
);

// CORBA::Any::replace() function.
void replace(
    CORBA::TypeCode_ptr,
    void* value,
    CORBA::Boolean release = 0
);
```

The CORBA::Environment Parameter

The signatures of IDL calls no longer contain the `CORBA::Environment` parameter. This parameter was needed for languages that did not support native exception handling. However, Orbix applications also use it for operation timeouts.

System Exception Semantics

Orbix and OrbixWeb clients that catch specific system exceptions might need to change the exceptions they handle when they are migrated to Orbix.

System exceptions

Orbix 6.3 follows the latest CORBA standards for exception semantics. [Table 2](#) shows the two system exceptions most likely to affect existing code.

Table 2: *Migrated System Exceptions*

When This Happens	Orbix 3 and OrbixWeb Raise	Orbix 6.3 Raises
Server object does not exist	INV_OBJREF	OBJECT_NOT_EXIST
Cannot connect to server	COMM_FAILURE	TRANSIENT

Minor codes

System exception minor codes are completely different between OrbixWeb 3.2 and Orbix 6.3 for Java. Applications that examine minor codes need to be modified to use Orbix 6.3 for Java minor codes.

Dynamic Invocation Interface

Proprietary dynamic invocation interface

Orbix-proprietary dynamic invocation interface (DII) functions are not available in Orbix 6.3. Code that uses `CORBA::Request::operator<<()` operators and overloads must be changed to use CORBA-compliant DII functions.

Note: Orbix 6.3-generated stub code consists of sets of statically generated CORBA-compliant DII calls.

Server Migration

Server code typically requires many more changes than client code. The main issue for server code migration is the changeover from the basic object adapter (BOA) to the portable object adapter (POA).

Function Signatures

Changes to the signature

In Orbix 6.3, two significant changes have been made to C++ function signatures:

- The `CORBA::Environment` parameter has been dropped.
- New types are used for `out` parameters. An `out` parameter of `T` type is now passed as a `T_out` type.

Consequently, when migrating C++ implementation classes you must replace the function signatures that represent IDL operations and attributes.

Object IDs versus Markers

C++ conversion functions

Orbix 6.3 uses a sequence of octets to compose an object's ID, while Orbix 3 uses string markers. CORBA provides the following helper methods to convert between the two types; hence migration from marker dependencies to Object IDs is straightforward.

```
// C++
// Converting string marker -----> ObjectId
PortableServer::ObjectId *
PortableServer::string_to_ObjectId(const char *);

// Converting ObjectId -----> string marker
char *
PortableServer::ObjectId_to_string(
    const PortableServer::ObjectId&
);
```

Java conversion functions

In Java, an object ID is represented as a byte array, `byte[]`. Hence the following native Java methods can be used to convert between string and object ID formats:

```
// Java
// Converting string marker -----> ObjectId
byte[]
java.lang.String.getBytes();

// Converting ObjectId -----> string marker
// String constructor method:
java.lang.String.String(byte[]);
```

CORBA Objects versus Servant Objects

Orbix 3

In Orbix 3 there is no need to distinguish between a CORBA object and a servant object. When you create an instance of an implementation class in Orbix 3, the instance already has a unique identity (represented by a marker) and therefore represents a unique CORBA object.

Orbix 6.3

In Orbix 6.3, a distinction is made between the identity of a CORBA object (its object ID) and its implementation (a servant). When you create an instance of an implementation class in Orbix 6.3, the instance is a servant object, which has no identity. The identity of the CORBA object (represented by an object ID) must be grafted on to the servant at a later stage, in one of the following ways:

- The servant becomes associated with a unique identity. This makes it a CORBA object, in a similar sense to an object in a BOA-based implementation.
- The servant becomes associated with multiple identities. This case has no parallel in a BOA-based implementation.

The mapping between object IDs and servant objects is controlled by the POA and governed by POA policies.

BOA to POA Migration

It is relatively easy to migrate a BOA-based server by putting all objects in a simple POA that uses an active object map; however, this approach is unable to exploit most of the functionality that a POA-based server offers. It is worth while redesigning and rewriting servers so they benefit fully from the POA.

Creating an Object Adapter

Creating a BOA in Orbix 3.x

In Orbix 3, a single BOA instance is used. All CORBA objects in a server are implicitly associated with this single BOA instance.

Creating a POA in Orbix 6.3

In Orbix 6.3, an application can create multiple POA instances (using the `PortableServer::POA::create_POA()` operation in C++ and the `org.omg.PortableServer.create_POA()` operation in Java). Each POA instance can be individually configured, using POA policies, to manage CORBA objects in different ways. When migrating to Orbix 6.3, you should give careful consideration to the choice of POA policies, to obtain the maximum benefit from the POA's flexibility.

Defining an Implementation Class

There are two approaches to defining an implementation class in CORBA:

- [The inheritance approach.](#)
- [The tie approach.](#)

The inheritance approach

The most common approach to implementing an IDL interface in Orbix is to use the inheritance approach. Consider the following IDL fragment:

```
//IDL
module BankSimple {
    Account {
        //...
    };
};
```

The `BankSimple::Account` IDL interface can be implemented by defining a class that inherits from a standard base class. The name of this standard base class for Orbix 3 and Orbix 6.3 is shown in [Table 3](#).

Table 3: *Standard Base Classes for the Inheritance Approach*

Application Type	Implementation Base Class Name
Orbix 3, C++ (BOA)	<code>BankSimple::AccountBOAImpl</code>
Orbix 6.3, C++ (POA)	<code>POA_BankSimple::Account</code>
Orbix 3, Java (BOA)	<code>BankSimple._AccountImplBase</code>

Table 3: *Standard Base Classes for the Inheritance Approach*

Application Type	Implementation Base Class Name
Orbix 6.3, Java (POA)	BankSimple.AccountPOA

Consider a legacy Orbix 3 application that implements `BankSimple::Account` in C++ as the `BankSimple_Account_i` class. The `BankSimple_Account_i` class might be declared as follows:

```
// C++
// Orbix 3 Version
// Inheritance Approach
class BankSimple_Account_i : BankSimple::AccountBOAImpl {
public:
    // Declare IDL operation and attribute functions...
};
```

When this implementation class is migrated to Orbix 6.3, the `BankSimple::AccountBOAImpl` base class is replaced by the `POA_BankSimple::Account` base class, as follows:

```
// C++
// Orbix 6.3 Version
// Inheritance Approach
class BankSimple_Account_i : POA_BankSimple::Account {
public:
    // Declare IDL operation and attribute functions...
};
```

The tie approach

The tie approach is an alternative mechanism for implementing IDL interfaces. It allows you to associate an implementation class with an IDL interface using a delegation approach rather than an inheritance approach.

In Application Server Platform (C++) the tie classes are generated using C++ templates. When migrating from Orbix 3 to Orbix 6.3, all `DEF_TIE` and `TIE` preprocessor macros must be replaced by the equivalent template syntax.

In Orbix 6.3 (Java) the tie approach is essentially the same as in Orbix 3. However, the names of the relevant Java classes and interfaces are different. For example, given an IDL interface, `Foo`, an Orbix 6.3 servant class implements the `FooOperations` Java interface and the associated Java tie class is called `FooPOATie`.

Creating and Activating a CORBA Object

To make a CORBA object available to clients, you should:

1. Create an implementation object. An implementation object is an instance of the class that implements the operations and attributes of an IDL interface. In Orbix 3, an implementation

object is the same thing as a CORBA object. In Orbix 6.3, an implementation object is a servant object, which is not the same thing as a CORBA object.

2. Activate the servant object. Activating a servant object attaches an identity to the object (a marker in Orbix 3 or an object ID in Orbix 6.3) and associates the object with a particular object adapter.

Orbix 3

In Orbix 3, creating and activating an object are rolled into a single step. For example, in C++ you might instantiate a `BankSimple::Account` CORBA object using the following code:

```
// C++
// Orbix 3
// Create and activate a new 'Account' object.
BankSimple_Account_i * acc1 =
    new BankSimple_Account_i("ObjectID");
```

This step creates the CORBA object and attaches the *ObjectID* identity to it (initializing the object's marker). The constructor automatically activates the CORBA object.

Orbix 6.3

In Orbix 6.3, creating and activating an object are performed as separate steps. For example, in C++ you might instantiate a `BankSimple::Account` CORBA object using the following code:

```
// C++
// Orbix 6.3

// Step 1: Create a new 'Account' object.
BankSimple_Account_i * acc1 = new BankSimple_Account_i();

// Step 2: Activate the new 'Account' object.
PortableServer::ObjectId_var oid =

    PortableServer::string_to_ObjectId("ObjectID");
// persistent_poa created previously
persistent_poa->activate_object_with_id(oid, acc1);
```

Activation is performed as an explicit step in Orbix 6.3. The call to `PortableServer::POA::activate_object_with_id()` attaches the *ObjectID* identity to the object and associates the `persistent_poa` object adapter with the object.

Migrating Proprietary Orbix 3 Features

Proprietary Orbix 3 features are replaced by a range of standards-compliant Orbix 6.3 features.

Orbix 3 Locator

The Orbix 3 locator is an Orbix-specific feature that is used in combination with `_bind()` to locate server processes. Because Orbix 6.3 does not support `_bind()`, it cannot use the Orbix 3 style locator.

Note: Orbix 6.3 has a feature called a locator, which is not related in any way to the Orbix 3 locator. The Orbix 6.3 locator is a daemon process, `itlocator`, that locates server processes for clients.

If your legacy code uses the Orbix 3 locator, you must replace it with one of the following Orbix 6.3 features:

- [High availability](#).
- [The CORBA Naming Service](#).
- [The CORBA Initialization Service](#).

High availability

The Orbix 6.3's high availability feature provides fault tolerance—that is, a mechanism that avoids having a single point of failure in a distributed application. With the enterprise edition of Orbix 6.3, you can protect your system from single points of failure through clustered servers.

A clustered server comprises multiple instances, or replicas, of the same server; together, these act as a single logical server. Clients invoke requests on the clustered server and Orbix routes the requests to one of the replicas. The actual routing to any replica is transparent to the client.

The CORBA Naming Service

If your legacy code uses the load-balancing feature of the Orbix 3 locator, you can replace this by the `ObjectGroup` feature of the Orbix 6.3's naming service. Object groups are an Orbix-specific extension to the naming service that allow you to register a number of servers under a single name.

Table 4 shows how the Orbix 3 locator maps to the equivalent naming service functionality.

Table 4: Replacing the Orbix 3 Locator by the Naming Service

Orbix 3-Locator	Orbix 6.3-Naming Service
Entry in the locator file, mapping the server name, <i>SrvName</i> , to a single server host, <i>HostName</i> : <i>SrvName</i> : <i>HostName</i> :	Object binding in the naming service, mapping a <i>name</i> to a single object reference.
Entry in the locator file, mapping the server name, <i>SrvName</i> , to multiple host names: <i>SrvName</i> : <i>Host1</i> , <i>Host2</i> , <i>Host3</i> :	Object group in the naming service, mapping a name to multiple object references.
Overriding functionality of <code>CORBA::LocatorClass</code> .	Custom implementation of the <code>IT_LoadBalancing::ObjectGroup</code> interface.

The naming service is the preferred way to locate objects in Orbix 6.3. It is a standard service and is highly scalable.

The CORBA Initialization Service

The initialization service uses the `CORBA::ORB::resolve_initial_references()` operation to retrieve an object reference from an Orbix 6.3 configuration file, `DomainName.cfg`.

Table 5 shows how the Orbix 3 locator maps to the equivalent initialization service functionality.

Table 5: Replacing the Orbix 3 Locator by the Initialization Service

Orbix 3-Locator	Orbix 6.3-Initialization Service
Entry in the locator file, mapping the server name, <i>SrvName</i> , to a single server host, <i>HostName</i> : <i>SrvName</i> : <i>HostName</i> :	Entry in the <code>DomainName.cfg</code> file, mapping an <i>objectId</i> to a single object reference: <code>initial_references:ObjectId: reference = "IOR:00...";</code>
Entry in the locator file, mapping the server name, <i>SrvName</i> , to multiple host names: <i>SrvName</i> : <i>Host1</i> , <i>Host2</i> , <i>Host3</i> :	<i>No Equivalent</i>
Override functionality of <code>CORBA::LocatorClass</code> .	<i>No Equivalent</i>

The initialization service can only be used as a replacement for the Orbix 3 locator when a simple object lookup is needed.

Filters

Filters are a proprietary Orbix 3 mechanism that allow you to intercept invocation requests on the server and the client side.

Orbix 6.3 does not support the filter mechanism. Instead, a variety of Orbix 6.3 features replace Orbix 3 filter functionality.

Equivalents

[Table 6](#) summarizes the typical uses of Orbix 3 filters alongside the equivalent features supported by Orbix 6.3.

Table 6: Orbix 6.3 Alternatives to Filter Features

Orbix 3 Filter Feature	Orbix 6.3 Equivalent
Request logging	Use portable interceptors.
Piggybacking data on a Request	Use portable interceptors.
Multi-threaded request processing	Use a multi-threaded POA and (optionally) a proprietary <code>WorkQueue</code> POA policy.
Accessing the client's TCP/IP details	<i>Not supported.</i>
Security using an authentication filter	Full security support is provided in the Orbix 6.3 enterprise edition.

Request Logging

Using portable interceptors

In Orbix 6.3, request logging is supported by the new portable interceptor feature. Interceptors allow you to access a CORBA request at any stage of the marshaling process, offering greater flexibility than Orbix filters. You can use them to add and examine service contexts. You can also use them to examine the request arguments.

Piggybacking Data on a Request

Piggybacking in Orbix 3

In Orbix 3, filters support a piggybacking feature that enables you to add and remove extra arguments to a request message.

Piggybacking in Orbix 6.3

In Orbix 6.3, piggybacking is replaced by the CORBA-compliant approach using *service contexts*. A service context is an optional block of data that can be appended to a request message, as specified in the IIOP 1.1 standard. The content of a service context can be arbitrary and multiple service contexts can be added to a request.

Multi-Threaded Request Processing

Orbix 3

In Orbix 3, concurrent request processing is supported using an Orbix thread filter. The mechanism is flexible because it gives the developer control over the assignment of requests to threads.

Orbix 6.3

In Orbix 6.3, request processing conforms to the CORBA 2.6 specification. Each POA can have its own threading policy:

- `SINGLE_THREAD_MODEL` ensures that all servant objects in that POA have their functions called in a serial manner. In Orbix 6.3, servant code is called only by the main thread, therefore no locking or concurrency-protection mechanisms need to be used.
- `ORB_CTRL_MODEL` leaves the ORB free to dispatch CORBA invocations to servants in any order and from any thread it chooses.

Orbix 6.3 request processing extensions

Because the CORBA 2.6 specification does not specify exactly what happens when the `ORB_CTRL_MODEL` policy is chosen, Orbix 6.3 makes some proprietary extensions to the threading model.

The multi-threaded processing of requests is controlled using the Orbix 6.3 work queue feature. Two kinds of work queue are provided by Orbix 6.3:

- *Automatic Work Queue*: A work queue that feeds a thread pool. When a POA uses an automatic work queue, request events are automatically dequeued and processed by threads. The size of the thread pool is configurable.
- *Manual Work Queue*: A work queue that requires the developer to explicitly dequeue and process events.

Manual work queues give developers greater flexibility when it comes to multi-threaded request processing. For example, prioritized processing of requests could be implemented by assigning high-priority CORBA objects to one POA instance and low-priority CORBA objects to a second POA instance. Given that both POAs are associated with manual work

queues, the developer can write threading code that preferentially processes requests from the high-priority POA.

Accessing the Client's TCP/IP Details

Recommendations for Orbix 6.3

Some Orbix 3 applications use Orbix-specific extensions to access socket-level information, such as the caller's IP address, in order to implement proprietary security features. These features are not available in Orbix 6.3, because providing access to low-level sockets would considerably restrict the flexibility of CORBA invocation dispatch.

To provide security for your applications, it is recommended that you use an implementation of the security service provided with the Orbix 6.3 Enterprise Edition instead.

Security Using an Authentication Filter

Recommendations for Orbix 6.3

Some Orbix 3 applications use authentication filters to implement security features. In Orbix 6.3, it is recommended that you use the security service that is made available with the Orbix 6.3 Enterprise Edition.

Loaders

Orbix 3 loader

The Orbix 3 loader provides support for the automatic saving and restoration of persistent objects. The loader provides a mechanism that loads CORBA objects automatically into memory, triggered in response to incoming invocations.

Servant manager

The Orbix 3 loader is replaced by equivalent features of the Portable Object Adapter (POA) in Orbix 6.3. The POA can be combined with a servant manager to provide functionality equivalent to the Orbix 3 loader. There are two different kinds of servant manager:

- *Servant activator*: Triggered only when the target CORBA object cannot be found in memory.
- *Servant locator*: Triggered for every invocation.

Servant activator

Taking the `PortableServer::ServantActivator` class as an example, the member functions of `CORBA::LoaderClass` correspond approximately as shown in [Table 7](#).

Table 7: Comparison of Loader with Servant Activator Class

CORBA::LoaderClass Member Function	ServantActivator Member Function
<code>save()</code>	<code>etherealize()</code>
<code>load()</code>	<code>incarnate()</code>
<code>record()</code>	<i>No equivalent function.</i> An Orbix 6.3 object ID (equivalent to an Orbix 3 marker) can be specified at the time a CORBA object is created. This gives sufficient control over object IDs.
<code>rename()</code>	<i>No equivalent function.</i> An Orbix 6.3 object ID (equivalent to an Orbix 3 marker) cannot be changed after a CORBA object has been created.

Servant locator

A servant locator can also be used to replace the Orbix 3 loader. In general, the servant locator is more flexible than the servant activator and offers greater scope for implementing sophisticated loader algorithms.

Smart Proxies

Orbix 3

The Orbix 3 smart proxies feature is a proprietary mechanism for overriding the default implementation of the proxy class. This allows applications to intercept outbound client invocations and handle them within the local client process address space, rather than using the default proxy behavior of making a remote invocation on the target object. Smart proxies can be used for such purposes as client-side caching, logging, load-balancing, or fault-tolerance.

Orbix 6.3

Orbix 6.3 does not support smart proxies. The primary difficulty is that, in the general case, it is not possible for the client-side ORB to determine if two object references denote the same server object. The CORBA standard restricts the client-side ORB from interpreting the object key or making any assumptions about it. Orbix 3 was able to avoid this limitation by making assumptions about the structure of the object key. This is neither CORBA-compliant nor interoperable with other ORBs.

At best, the ORB can only determine that two object references are equivalent if they have exactly the same server location (host and port in IIOP) and object key. Unfortunately, this can be an unreliable indicator if object references pass through bridges, concentrators, or firewalls that change the server location or object key.

In this case, it is possible for two object references denoting the same CORBA object to appear different to the ORB, and thus have two different smart proxy instances. Since smart proxies are commonly used for caching, having two smart proxy instances for a single CORBA object is unacceptable.

Replacing smart proxies

[Table 8](#) shows how smart proxy tasks can be mapped to equivalent features in Orbix 6.3.

Table 8: Orbix 6.3 Alternatives to Smart Proxy Features

Orbix 3 Smart Proxy Task	Orbix 6.3 Equivalent Feature
Fault tolerance	Orbix 6.3 high availability, based on server clusters.
Logging	Orbix 6.3 built-in logging facility or portable interceptors
Caching	Implement smart proxy-like functionality by hand.

Fault tolerance

Fault tolerance is provided by the high availability feature of the Orbix 6.3's locator. See ["High availability" on page 27](#).

Logging

For logging that requires access to request parameters, portable interceptors can be used in Orbix 6.3. Portable interceptors are similar to Orbix 3 filters, but they are more flexible in that they allow you to read request parameters.

Caching

A smart proxy that implements client-side caching of data cannot be mimicked by a standard Orbix 6.3 feature. In this case, you have no option but to implement smart proxy-like functionality in Orbix 6.3, and this can be done as follows:

1. Create a local implementation of the object to be proxified, by writing a class that derives from the client-side stub class.
2. Every time the client receives an object reference of the appropriate type, wrap the object reference with a corresponding smart proxy object. Before wrapping the object reference, however, you must determine the target object's identity by making an invocation on the remote target object, asking it for a system-wide unique identifying name. This is the key step that avoids the object identity problem described in ["Orbix 6.3" on page 33](#).

Based on the system-wide unique identifying name, the application can then either create a new smart proxy, or reuse the target object's existing smart proxy. The client application should consistently use the smart proxy in place of the regular proxy throughout the application.

Transformers

Orbix 3

Transformers are a deprecated feature of Orbix 3 that allow you to apply customized encryption to CORBA request messages. This could be used to implement a primitive substitute for a security service.

Orbix 6.3

In Orbix 6.3, transformers are not supported. It is recommended, instead, that you use the security service that is made available with the enterprise edition of Orbix 6.3.

I/O Callbacks

Orbix 6.3 does not allow access to TCP/IP sockets or transport-level information. This is incompatible with the Orbix 6.3 architecture, which features a pluggable transport layer. Using Orbix 6.3, you can replace TCP/IP with another transport plug-in such as IP multicast (which is connectionless), simple object access protocol (SOAP), hypertext transfer protocol (HTTP), asynchronous transfer mode (ATM), and so on. For example, the shared memory transport (SHMIOP) does not use file descriptors or sockets.

Purposes for using I/O callbacks

Orbix 3 I/O Callback functionality is generally used for two main purposes:

- *Connection Management*—the number of TCP/IP connections that can be made to a single process is typically subject to an operating system limit. Some form of connection management is required if this limit is likely to be reached in a deployed system.
- *Session Management*—I/O Callback functionality can be used to implement an elementary session-tracking mechanism. The opening of a connection from a client defines the beginning of a session and the closing of the connection defines the end of the session.

Because Orbix 6.3 has no equivalent to the Orbix 3 I/O Callback functionality, you must migrate any code that uses it.

Connection Management

Active connection management

Orbix 6.3 provides an active connection manager (ACM) that allows the ORB to reclaim connections automatically, and thereby increases the number of clients that can use a server beyond the limit of available file descriptors.

ACM configuration variables

IIOP connection management is controlled by four configuration variables:

- `plugins:iiop:incoming_connections:hard_limit` sets the maximum number of incoming (server-side) connections allowed to IIOP. IIOP refuses new connections above this limit.
- `plugins:iiop:incoming_connections:soft_limit` specifies the number of connections at which IIOP begins closing incoming (server-side) connections..
- `plugins:iiop:outgoing_connections:hard_limit` sets the maximum number of outgoing (client-side) connections allowed to IIOP. IIOP refuses new outgoing connections above this limit.
- `plugins:iiop:outgoing_connections:soft_limit` specifies the number of connections at which IIOP begins closing outgoing (client-side) connections.

Closing client connections

The ORB first tries to close idle connections in least-recently-used order. If there are no idle connections, the ORB closes busy connections in least-recently-opened order.

Active connection management effectively remedies file descriptor limits that has constrained past Orbix applications. If a client is idle for a while and the server ORB reaches its connection limit, it sends a GIOP `CloseConnection` message to the client and closes the connection. Later, the same client can transparently reestablish its connection, to send a request without throwing a CORBA exception.

Note: In Orbix 3, Orbix tended to throw a `COMM_FAILURE` on the first attempt at reconnection; server code that anticipates this exception should be reevaluated against current functionality.

Default file descriptor limits

Orbix 6.3 is configured to use the largest upper file descriptor limit on each supported operating system. On UNIX, it is typically possible to rebuild the kernel to obtain a larger number. However, active connection management should make this unnecessary.

Session Management

Because Orbix 6.3 features a pluggable transport layer, it is not appropriate to relate the duration of a client session to the opening and closing of TCP/IP connections from clients. This type of session management, which is typically implemented using I/O callbacks in Orbix 3, has to be migrated to an alternative model.

Session management in Orbix 6.3

Support for session management in Orbix 6.3 is provided by a *lease plug-in*. The lease plug-in implements a scheme for automatically tracking client sessions, based on the idea that a client obtains a lease from the server for the duration of a client session.

Client migration

Client applications can easily be modified to use session management. Just edit the Orbix 6.3 configuration to make the client load the lease plug-in. No changes to the client source code are required.

Server migration

On the server side, the following changes are required to use session management in Orbix 6.3:

- Edit the Orbix 6.3 configuration to make the server load the lease plug-in.
- Modify the server source code so that it uses the lease plug-in to track client sessions.

Further details

See the *CORBA Session Management Guide* for details of how to program and configure the lease plug-in for session management. Demonstration code for the lease plug-in is also provided with the Orbix 6.3 product.

CORBA Services

Orbix includes several CORBA services, such as the interface repository, the naming service, the notification service, and the security service. Because these service are based mainly on the CORBA standard, there are not many changes between Orbix 3 and Orbix 6.3.

Interface Repository

Migration

Migrating source code that uses the Interface Repository (IFR) to Orbix 6.3 is straightforward. Link the migrated application against the stub code derived from the Orbix 6.3 version of the interface repository. No further changes should be necessary.

Naming Service

Backward compatibility

The Orbix 6.3's naming service is backward compatible with Orbix 3.x in two respects:

- *Source code backward compatibility:* source code that is written to use the standard naming service interfaces can be migrated to Orbix 6.3 without modification.
- *On-the-wire backward compatibility:* Orbix 3.x applications can interoperate with the Orbix 6.3 naming service. If you need to interoperate Orbix 3.x applications, it is recommended that you recompile the naming stub code from the Orbix 6.3 IDL files.

New interface

Orbix 6.3 adds a new interface, `CosNaming::NamingContextExt`, which is defined by the CORBA Interoperable Naming Service specification. This interface adds support for using names in stringified format.

Load balancing

The naming service load-balancing extensions provided in Orbix 3 are also present in Orbix 6.3. The Orbix 6.3 load-balancing interfaces are only slightly different from Orbix 3, requiring small modifications to your source code.

Notification Service

The Orbix 6.3 notification service has undergone significant modifications since the OrbixNotification 3 generation of the notification service.

Many of the changes that impact application migration reflect changes in the CORBA standard and require minimal changes to legacy OrbixNotification 3 application code.

CORBA Specification Updates

The Orbix 6.3 notification service complies with both the CORBA 2.6 specification and the OMG's Notification Service Specification, approved in June of 2000. To achieve compliance with these specifications several changes were made to the notification services IDL and APIs.

These changes require that any applications that use generation 3 code need to be recompiled and re-linked, at the very least. Other minor changes might also need to be made to generation 3 code to accommodate the changes in the APIs. Compiler warnings warn you of most changes that need to be made.

`_bind()`

The Orbix 6.3 notification service clients do not use `_bind()` to contact the notification service. Instead, clients should call `resolve_initial_references("NotificationService")` to obtain an object reference to the notification service. See ["Replacing the `_bind\(\)` Function" on page 15](#) for more information.

Subscription and publication notification

Orbix 6.3 provides notification service clients greater flexibility over how they receive subscription and publication details from the notification channel. To accomplish this, an input parameter has been added to `obtain_offered_types()` and `obtain_subscription_types()`.

The Orbix 6.3 operation signatures are:

```
// IDL
CosNotification::EventTypeSeq obtain_subscription_types(
    in ObtainInfoMode mode);
CosNotification::EventTypeSeq obtain_offered_types(
    in ObtainInfoMode mode);
```

The new parameter is of type `ObtainInfoMode` which is an enum defined in `CosNotifyChannelAdmin` as:

```
// IDL
enum ObtainInfoMode
{
    ALL_NOW_UPDATES_OFF,
    ALL_NOW_UPDATES_ON,
    NONE_NOW_UPDATES_OFF,
    NONE_NOW_UPDATES_ON
};
```

Any generation 3 clients that call `obtain_offered_types()` or `obtain_subscription_types()` need to add the parameter. `ALL_NOW_UPDATES_OFF` mimics generation 3 functionality. For more information on the other values, see the *CORBA Notification Service Guide*.

Unstructured event clients

Orbix 6.3 introduces unstructured event, any-style, client interfaces into the `CosNotifyComm` module. This allows any-style clients to support the enhanced subscription features and it standardizes notification service client development. Any-style clients developed for OrbixNotification 3 use the interfaces from `CosEventComm`.

In addition, the Orbix 6.3 any-style proxy interfaces, defined in `CosNotifyChannelAdmin`, inherit their client interfaces directly from `CosNotifyComm`. In OrbixNotification 3 any-style proxies inherit client interfaces from `CosNotifyComm:NotifyPublish` and `CosEventComm::PushConsumer`.

Note: The `connect()` operation's parameter is still an interface defined in `CosEventComm`.

Not updating legacy code does not generate any compiler errors. However, at runtime any-style clients using legacy code are not able to contact the notification service.

TimeBase::TimeT

Orbix 6.3 supports the new OMG standard definition of `TimeBase::TimeT`. In OrbixNotification 3 `TimeBase::TimeT` is defined as a structure containing two unsigned longs. In Orbix 6.3 it is defined as a `CORBA::ULongLong`.

Any generation 3 clients that use the timing features of the service need to be updated to support the new definition of `TimeBase::TimeT`. If they are not, the Orbix 6.3 notification service generates marshalling errors at runtime.

Quality of Service Properties

Orbix 6.3 notification service uses new several new Quality-of-Service (QoS) properties and has reimplemented others.

PacingInterval

`PacingInterval` is re-implemented as a `TimeBase::TimeT` in Orbix 6.3 and is specified in units of 10^{-7} seconds. In Orbix 3 it is a `TimeBase:UtcT` and is specified in milliseconds.

Orbix 6.3 QoS properties

[Table 9](#) lists the new Orbix 6.3 QoS properties. For more detailed information on Orbix 6.3 QoS properties, see the *CORBA Notification Service Guide*.

Table 9: *Orbix 6.3 QoS Properties (Sheet 1 of 2)*

QoS Property	Description
<code>MaxEventsPerConsumer</code>	Specifies the maximum number of undelivered events that a channel will queue for a consumer. It is set with a <code>long</code> and is valid for supplier proxies, consumer admins, and notification channels.
<code>MaxRetries</code>	Specifies the maximum number of times a proxy push supplier calls <code>push()</code> on its consumer before giving up, or the maximum number of times a proxy pull consumer calls <code>pull()</code> or <code>try_pull()</code> on its supplier before giving up. It is set with a <code>CORBA::Ulong</code> and is valid for consumer admins and notification channels.
<code>RetryTimeout</code>	Specifies the amount of time that elapses between attempts by a proxy push supplier to call <code>push()</code> on its consumer. It is set with a <code>TimeBase::TimeT</code> and defaults to 1 second.
<code>MaxRetryTimeout</code>	Sets the ceiling for the calculated value of <code>RetryTimeout</code> . It is set with a <code>TimeBase::TimeT</code> and defaults to 60 seconds.
<code>RequestTimeout</code>	Specifies the amount of time a channel object has to perform an operation on a client. It is set using a <code>TimeBase::TimeT</code> .

Table 9: *Orbix 6.3 QoS Properties (Sheet 2 of 2)*

QoS Property	Description
PullInterval	Specifies the amount of time that elapses between attempts by a proxy pull consumer to call <code>pull()</code> or <code>try_pull()</code> on its consumer. It is specified with a <code>long</code> and defaults to 1 second.
RetryMultiplier	Specifies the number used to calculate the amount of time between attempts by a proxy push supplier to call <code>push()</code> on its consumer. It is set with a <code>CORBA::double</code> and defaults to 1.0.

Channel administration properties

Orbix 6.3 has introduced two properties to control the administration of a notification channel. These properties can only be set on a notification channel. For more information, see the *CORBA Notification Service Guide*.

[Table 10](#) describes the new properties.

Table 10: *Orbix 6.3 Administration Properties*

Property	Description
MaxConsumers	Specifies the maximum number of consumers that can be connected to a channel at a given time. It is set using a <code>long</code> and defaults to 0 (unlimited).
MaxSuppliers	Specifies the maximum number of suppliers that can be connected to a channel at a given time. It is set using a <code>long</code> and defaults to 0 (unlimited).

Configuration/Administration Changes

Centralized configuration

Orbix 6.3 has a centralized configuration mechanism. This means that the notification service is configured using the standard Orbix 6.3 configuration tools and the information is stored in the common Orbix 6.3 database.

Starting the notification service

The Orbix 6.3 notification service can be configured to start on system boot, on demand, or from the command line.

To start the notification service from the command line use:

```
itnotify run [-background]
```

The `-background` flag is optional and starts the notification service to run as a background process.

Managing the notification service

The Orbix 6.3 notification service can be managed in one of two ways.

- The Orbix 6.3 `itadmin` tool. For more information, see the *CORBA Administrator's Guide*.
- The Orbix 6.3 notification console, `itnotifyconsole`. For more information on using the console, see the *CORBA Notification Service Guide*.

Configuration variables

The Orbix 6.3 notification service uses a new set of configuration variables. See the *CORBA Administrator's Guide* for a detailed listing of the new configuration variables.

Deprecated Features

Orbix 6.3 has deprecated some proprietary features from OrbixNotification 3. Any notification clients that make use of these features need to be updated.

HealthCheck

The OrbixNotification 3 HealthCheck feature allows notification channels, and optionally notification clients, to monitor their connections. In Orbix 6.3 this feature is no longer supported.

Code Modification

To find code using the HealthCheck feature search for the following strings:

- `DO_HEALTHCHECK`
- `DO_GL_HEALTHCHECK`
- `initializeHealthCheck`
- `startHealthCheck`
- `stopHealthCheck`
- `HealthCheck.h`

This code must be removed before the clients can be compiled using the Orbix 6.3 libraries.

Simulating HealthCheck in Orbix 6.3

HealthCheck-like functionality is implemented in Orbix 6.3, using the `MaxRetries` QoS property. If a `ProxyPushSupplier` or a `ProxyPullConsumer` fails to communicate with its associated client in `MaxRetries` attempts, the notification channel forces a disconnect and destroys all of the resources used to support the client.

String events

Orbix 6.3 no longer supports string events. All generation 3 clients using string events must be rewritten to use a valid event type.

SSL/TLS Toolkit

This section describes how to migrate from OrbixSSL or Orbix 3.3 security to the Orbix 6.3 SSL/TLS security service. Orbix 6.3 SSL/TLS has a very similar set of features to Orbix 3.3 security and it supports interoperability with legacy Orbix applications (see [“SSL/TLS Toolkit Interoperability” on page 90](#)).

The programming interfaces and administration of security have, however, changed significantly between Orbix 3.3 and Orbix 6.3. This section provides an overview of these changes.

Changes to the Programming Interfaces

Support for security level 2

The APIs for Orbix 6.3 SSL/TLS are based on the CORBA security level 2 interfaces. The programming interface is, therefore, based on the following standard IDL modules:

- `Security`
- `SecurityLevel1`
- `SecurityLevel2`

Note: Orbix 6.3 SSL/TLS does not implement every interface in the `SecurityLevel1` and `SecurityLevel2` modules. The CORBA security API is a mechanism-neutral API that can be layered over a variety of security toolkits. Some of the standard interfaces are more appropriately implemented by a higher level security layer.

CORBA policy-based API

In contrast to OrbixSSL 3.x, the Orbix 6.3 SSL/TLS product supports a *CORBA policy-based* approach to setting security properties. This represents a significant enhancement over OrbixSSL 3.x, because the policy-based approach lets you set properties at a finer granularity than before.

For example, client policies can be set at the following levels:

- ORB
- Thread
- Object reference

Server policies can be set at the following levels:

- ORB
- POA

No support for certificate revocation lists

Orbix 6.3 SSL/TLS has no support for certificate revocation lists (CRL). Therefore, the following OrbixSSL 3.x interfaces have no Orbix 6.3 equivalent:

```
IT_CRL_List
IT_X509_CRL_Info
IT_X509_Revoked
IT_X509_RevokedList
```

If you require certificate revocation in Orbix 6.3, you can programmatically implement any required revocation checks by registering a certificate validator policy,

```
IT_TLS_API::CertValidatorPolicy.
```

Mechanism-specific API

Orbix 6.3 SSL/TLS provides a number of value-added APIs that deal with the mechanism-specific aspects of the SSL/TLS toolkit. The extra IDL interfaces provide the facility to parse X.509 certificates and set Orbix-specific security policies.

The mechanism-specific API is defined by the following IDL modules:

- IT_Certificate
- IT_TLS
- IT_TLS_API

Migrating OrbixSSL 3.x classes and data types

When migrating to Orbix 6.3, most of the old C++ and Java classes from OrbixSSL 3.x are replaced by equivalent IDL interfaces. [Table 11](#) shows which OrbixSSL classes and data types to replace by the equivalent Orbix 6.3 SSL/TLS types.

Table 11: *Mapping OrbixSSL 3.x Types to Orbix 6.3 SSL/TLS (Sheet 1 of 2)*

OrbixSSL 3.x Type	Orbix 6.3 SSL/TLS Equivalent
IT_AVA	IT_Certificate::AVA
IT_AVAList	IT_Certificate::AVAList
IT_CertError	IT_Certificate::CertError

Table 11: Mapping OrbixSSL 3.x Types to Orbix 6.3 SSL/TLS (Sheet 2 of 2)

OrbixSSL 3.x Type	Orbix 6.3 SSL/TLS Equivalent
IT_CRL_List	<i>No equivalent</i>
IT_Extension	IT_Certificate::Extension
IT_ExtensionList	IT_Certificate::ExtensionList
IT_OID	IT_Certificate::ASN_OID
IT_OIDTag	IT_Certificate::OIDTag
IT_SSL	Equivalent functionality provided by the Security, SecurityLevel1, SecurityLevel2, and IT_TLS_API IDL modules.
IT_UTCTime	IT_Certificate::UTCTime
IT_ValidateX509CertCB	Use a combination of the IT_TLS::CertValidator interface and the IT_TLS_API::CertValidatorPolicy interface.
IT_X509_CRL_Info	<i>No equivalent</i>
IT_X509_Revoked	<i>No equivalent</i>
IT_X509_RevokedList	<i>No equivalent</i>
IT_X509Cert	IT_Certificate::X509Cert
IT_X509CertChain	IT_Certificate::X509CertChain

Configuration and Administration

Enabling security in Orbix 6.3

Security in Orbix 6.3 is enabled by configuring an application to load the security plug-in, `iiop_tls`. This is a relatively simple procedure involving just a few changes in the Orbix 6.3 configuration file; although advanced applications might also need to use security APIs.

Because application security is controlled by editing the configuration file, you must ensure that access to the configuration file is restricted.

External configuration granularity

The external configuration granularity refers to the effective scope of security configuration settings that are made in a configuration file. The external configuration granularity is mapped as follows:

- In OrbixSSL 3.x, it is identified with a process.
- In Orbix 6.3 SSL/TLS, it is identified with a single ORB instance.

KDM support

The key distribution management (KDM) is a framework that enables automatic activation of secure servers. Both OrbixSSL 3.x and Orbix 6.3 SSL/TLS provide a KDM and the functionality is similar in each.

There is one significant difference between the OrbixSSL 3.x KDM and the Orbix 6.3 KDM. Protection against server impostors implemented differently in the two products:

- In OrbixSSL 3.x, a binary checksum is calculated from the contents of the server executable file. The server is launched only if the calculated checksum matches the cached value.
- In Orbix 6.3 SSL/TLS, the node daemon relies on the server executables being stored in a secured directory to prevent tampering. A different sort of checksum is calculated (based on the contents of the server activation record) to ensure that the node daemon cannot be fooled into launching a server from an insecure directory.

No CRL support

Orbix 6.3 SSL/TLS does not support certificate revocation lists. Hence, there are no equivalents for the corresponding OrbixSSL 3.x configuration variables. See also [“No support for certificate revocation lists” on page 46](#).

Migrating OrbixSSL 3.x configuration

Most of the OrbixSSL 3.x configuration variables have direct equivalents in Orbix 6.3, as shown in [Table 12](#). In addition, many of the properties listed in [Table 12](#) can also be set programmatically in Orbix 6.3.

Table 12: *Mapping OrbixSSL 3.x Configuration Variables to Orbix 6.3 (Sheet 1 of 2)*

OrbixSSL 3.x Configuration Variable	Orbix 6.3 SSL/TLS Equivalent
IT_CA_LIST_FILE	policies:trusted_ca_list_policy
IT_AUTHENTICATE_CLIENTS	policies:target_secure_invocation_policy
IT_SERVERS_MUST_AUTHENTICATE_CLIENTS.	policies:target_secure_invocation_policy
IT_INVOCATION_POLICY	policies:target_secure_invocation_policy policies:client_secure_invocation_policy
IT_SECURE_REMOTE_INTERFACES IT_SECURE_SERVERS IT_INSECURE_REMOTE_INTERFACES IT_INSECURE_SERVERS	These properties cannot currently be specified in the Orbix 6.3 configuration file. You can, however, set the properties programmatically using the following interfaces: SecurityLevel2::EstablishTrustPolicy SecurityLevel2::QOPPolicy
IT_CIPHERSUITES	policies:mechanism_policy

Table 12: Mapping OrbixSSL 3.x Configuration Variables to Orbix 6.3 (Sheet 2 of 2)

OrbixSSL 3.x Configuration Variable	Orbix 6.3 SSL/TLS Equivalent
IT_ALLOWED_CIPHERSUITES	<i>No equivalent in Orbix 6.3.</i>
IT_CERTIFICATE_FILE IT_CERTIFICATE_PATH	Equivalent functionality provided by: principal_sponsor:auth_method_data
IT_BIDIRECTIONAL_IOP_BY_DEFAULT	
IT_CACHE_OPTIONS	policies:session_caching_policy plugins:atli_tls_tcp:session_cache_validity _period plugins:atli_tls_tcp:session_cache_size
IT_DEFAULT_MAX_CHAIN_DEPTH	policies:max_chain_length
IT_MAX_ALLOWED_CHAIN_DEPTH.	<i>No equivalent in Orbix 6.3.</i>
IT_DAEMON_POLICY IT_DAEMON_UNRESTRICTED_METHODS IT_DAEMON_AUTHENTICATES_CLIENTS IT_ORBIX_BIN_SERVER_POLICY	In Orbix 6.3, the services are configured using standard Orbix 6.3 configuration variables such as the secure invocation policies.
IT_DAEMON_UNRESTRICTED_METHODS	<i>No equivalent in Orbix 6.3.</i> There is currently no concept of service authorization in Orbix 6.3.
IT_FILTER_BAD_CONNECTS_BY_DEFAULT	Not needed in Orbix 6.3.
IT_ENABLE_DEFAULT_CERT	Not needed in Orbix 6.3. There is no need for this option because Orbix 6.3 supports security unaware applications.
IT_DISABLE_SSL	Not needed in Orbix 6.3. Configure your application not to load the security plug-in.
IT_KDM_CLIENT_COMMON_NAMES IT_KDM_ENABLED IT_KDM_PIPES_ENABLED IT_KDM_REPOSITORY IT_KDM_SERVER_PORT	Equivalent functionality is provided by the KDM in Orbix 6.3. See the <i>CORBA SSL/TLS Guide</i> .
IT_CHECKSUMS_ENABLED IT_CHECKSUM_REPOSITORY	<i>No equivalent in Orbix 6.3.</i> There is no binary checksum functionality in Orbix 6.3. Orbix 6.3 SSL/TLS relies on storing server executables in secured directories.
IT_CRL_ENABLED IT_CRL_REPOSITORY IT_CRL_UPDATE_INTERVAL	<i>No equivalent in Orbix 6.3.</i> There is no CRL functionality in Orbix 6.3.

Migrating Certificate and Private Key Files

In OrbixSSL 3.x, a variety of certificate and private key formats are used in different parts of the product. Orbix 6.3 SSL/TLS is based on a unified certificate file format, the industry standard PKCS#12 format, and the PEM format for storing trusted CA certificates. This subsection describes how to convert each of the legacy formats to PKCS#12.

Certificate file formats

The following certificate file formats are used by OrbixSSL 3.x and Orbix 6.3 SSL/TLS:

- *Privacy enhanced mail (PEM) format*—A PEM file typically contains a single certificate. OrbixSSL 3.x can use this format to hold peer certificates. Orbix 6.3 SSL/TLS *cannot* use this format for peer certificates.
- *PKCS#12 format*—A PKCS#12 file contains a peer certificate chain, concatenated with a private key at the end. Both OrbixSSL 3.x and Orbix 6.3 SSL/TLS can use this format for peer certificates.

Migrating certificate files

You can migrate OrbixSSL 3.x certificate files to Orbix 6.3 SSL/TLS as shown in [Table 13](#).

Table 13: *Converting Certificate Files*

Source OrbixSSL 3.x File Format	Target Orbix 6.3 File SSL/TLS Format	How to Convert
PEM format	PKCS#12 format	Use the <code>openssl pkcs12</code> utility, specifying the complete peer cert chain, private key and pass phrase.
PKCS#12 format	PKCS#12 format	<i>No conversion needed.</i>

Private key file formats

The following private key file formats are used by OrbixSSL 3.x and Orbix 6.3 SSL/TLS:

- *PKCS#1 format*—An unencrypted private key format. Orbix 6.3 SSL/TLS only supports this format programmatically.
- *PKCS#8 format*—An encrypted private key format. Orbix 6.3 SSL/TLS only supports this format programmatically.
- *OpenSSL proprietary private key format*—A proprietary encrypted format generated by the OpenSSL toolkit utilities.

- *Proprietary KEYENC format (deprecated)*—An encrypted private key format generated by the OrbixSSL 3.x `keyenc` utility. This format was formerly used by OrbixSSL 3.x Java applications and is now deprecated.

Migrating key files

You can migrate OrbixSSL 3.x private key files to Orbix 6.3 SSL/TLS as shown in [Table 14](#).

Table 14: *Converting Private Key Files*

Source OrbixSSL 3.x File Format	Target Orbix 6.3 SSL/TLS File Format	How to Convert
PKCS#1 format	PKCS#12 format	Use the <code>openssl pkcs12</code> utility, specifying the complete peer cert chain, private key, and pass phrase.
OpenSSL proprietary encrypted private key format	PKCS#12 format	Convert as follows: <ol style="list-style-type: none"> 1. Decrypt using the <code>openssl rsa</code> command. 2. Encrypt as PKCS#12 using the <code>openssl pkcs12</code> utility, specifying the complete peer cert chain, private key, and pass phrase.
Proprietary <code>keyenc</code> format	PKCS#12 format	Convert as follows: <ol style="list-style-type: none"> 1. Decrypt using the <code>keyenc -d</code> command: 2. Encrypt as PKCS#12 using the <code>openssl pkcs12</code> utility, specifying the complete peer cert chain, private key, and pass phrase.

Trusted CA certificate lists

In both OrbixSSL 3.x and Orbix 6.3 SSL/TLS, a trusted CA certificate list file consists of a concatenated list of PEM certificates.

Note: The Orbix 6.3 SSL/TLS Java Edition product currently does not accept any extraneous text (comments and so on) in a trusted CA list file. The extra text must therefore be removed if you are using Orbix 6.3 SSL/TLS Java Edition.

Interoperability

In a mixed system containing Orbix 3.3 Java Edition and Orbix 6.3 SSL/TLS, the PKCS#12 format can be used for peer certificates because Orbix 3.3 Java Edition also accepts the PKCS#12 format.

Administration

The administration of Orbix 6.3 has changed significantly from Orbix 3. This chapter provides a brief overview of the main changes in Orbix administration.

Orbix Daemons

Orbix 6.3 daemons

To provide greater flexibility and scaling, Orbix 6.3 replaces the Orbix 3 daemon, `orbixd`, with two daemons:

- The locator daemon, `itlocator`, helps clients to find Orbix 6.3 servers.
- The node daemon, `itnode_daemon`, launches dormant Orbix 6.3 servers in response to a client's request for service.

POA Names

Administering POA Names

In Orbix 3, CORBA objects were associated with a named server. In Orbix 6.3, CORBA objects are associated with named POAs. This means that Orbix 6.3 object references include an embedded POA name instead of a server name.

The Orbix 6.3 locator daemon locates the CORBA object using the object reference's embedded POA name. Hence, POA names play a major role in configuring the Orbix 6.3 locator daemon.

Command-Line Administration Tools

Orbix 6.3 unifies many of Orbix 3's command-line tools under a single utility, `itadmin`. Also, some of the Orbix 3 command line-tools have been deprecated.

General command-line tools

[Table 15](#) compares the Orbix 3 general purpose command-line tools with the Orbix 6.3's tools.

Table 15: Comparison of Orbix 3 and Orbix 6.3 General Command-Line Tools (Sheet 1 of 2)

Description	Orbix 3	Orbix 6.3
Show implementation repository (IMR) entry.	<code>catit</code>	<code>itadmin process show</code>
Security commands.	<code>chowit</code> , <code>chmodit</code>	<i>No equivalent</i>
Show configuration.	<code>dumpconfig</code>	<code>itadmin config dump</code>

Table 15: Comparison of Orbix 3 and Orbix 6.3 General Command-Line Tools (Sheet 2 of 2)

Description	Orbix 3	Orbix 6.3
Associate hosts into groups.	grouphosts	<i>No equivalent</i>
C++ IDL compiler.	idl	idl
CodeGen toolkit.	idlgcn	idlgcn
Java IDL compiler.	idlj	idl
Interface Repository (IFR).	ifr	itifr
Kill a server process.	killit	itadmin process stop
List server.	lsit	itadmin process list
Create a sub-directory in the IMR.	mkdirit	<i>No equivalent</i>
Orbix daemon.	orbixd	itlocator and itnode_daemon
Ping the Orbix daemon.	pingit	<i>No equivalent</i>
List active servers.	psit	itadmin process list -active
Add a definition to the IFR.	putidl	idl -R
Register a server in the IMR.	putit	itadmin process create
Show an IFR definition.	readifr	itadmin ifr show
Remove a sub-directory from the IMR.	rmdirit	<i>No equivalent</i>
Unregister a server from the IMR.	rmit	itadmin process remove
Remove a definition from the IFR.	rmidl	itadmin ifr remove
Associate servers with groups.	servergroups	<i>No equivalent</i>
Associate hosts with servers.	serverhosts	<i>No equivalent</i>

Naming Service Command Line Tools

Table 16 compares the Orbix 3 naming service command-line tools with the Orbix 6.3 tools.

Table 16: Comparison of Orbix 3 and Orbix 6.3 Naming Service Command-Line Tools (Sheet 1 of 2)

Description	Orbix 3	Orbix 6.3
Add a member to an object group.	add_member	itadmin nsog add_member
Print the IOR of an object group.	cat_group	<i>No equivalent</i>

Table 16: Comparison of Orbix 3 and Orbix 6.3 Naming Service Command-Line Tools (Sheet 2 of 2)

Description	Orbix 3	Orbix 6.3
Print the IOR of an object group's member.	cat_member	itadmin nsog show_member
Print the IOR of a given name.	catns	itadmin ns resolve
Remove an object group.	del_group	itadmin nsog remove
Remove a member from an object group.	del_member	itadmin nsog remove_member
List all object groups.	list_groups	itadmin nsog list
List the members of an object group.	list_members	itadmin nsog list_member
List the bindings in a context.	lsns	itadmin ns list
Create an object group.	new_group	itadmin nsog create
Create an unbound context.	newncns	itadmin ns newnc
Select a member of an object group.	pick_member	<i>No equivalent</i>
Bind a name to a context.	putncns	itadmin ns bind -context
Create a bound context.	putnewncns	itadmin ns newnc
Bind a name to an object.	putns	itadmin ns bind -object
Rebind a name to a context.	reputncns	itadmin ns bind -context
Rebind a name to an object.	reputns	itadmin ns bind -object
Remove a binding.	rmns	itadmin ns remove

Activation Modes

Orbix 3

Orbix 3 process activation modes, *shared*, *unshared*, *per-method*, *per-client-pid*, and *persistent* are used for a variety of reasons. For example, they are used to achieve multi-threaded behavior in a single-threaded environment, to increase server reliability, and so on. The two most popular modes are:

- *Shared mode*—which enables all clients to communicate with the same server process.
- *Per-client-pid mode*—which enforces a 1-1 relationship between client process and server process, is sometimes used to maximize server availability.

Orbix 6.3

Orbix 6.3 provides the following activation modes:

- `on_demand`—the process only activates when required.
- `per_client`—a new process is activated for each client.

Orbix 6.3 moved CORBA object association from the server to the POA. Because of this, all Orbix 6.3 processes are shared.

Migration

Migration of source code should be straightforward, because the choice of activation mode has almost no impact on BOA or POA-based server code.

Load balancing

The additional activation modes provided by Orbix 3 are typically used to achieve some form of load-balancing that is transparent to the client. The Enterprise Edition of Orbix 6.3 includes transparent locator-based load balancing over a group of replica POAs. This answers the needs currently addressed by Orbix 3 activation modes.

Part III

Interoperability

In this part

This part contains the following chapters:

Configuring for Interoperability	page 59
IDL Issues	page 65
Exceptions	page 73
Services	page 85
Connection Management	page 93
Codeset Negotiation	page 97

Configuring for Interoperability

This chapter describes the main configuration changes that must be made to facilitate interoperability between Orbix 3.x and Orbix 6.3 applications.

Interoperability Overview

This Interoperability Guide describes how to configure applications that use a mixture of Orbix products and any feature limitations that apply to such interoperating systems.

Orbix 6.3 interoperability

Because Orbix 6.3 is binary-compatible with Orbix E2A ASP v6.0, Orbix 6.3 has the same interoperability characteristics as ASP 6.0.

Orbix E2A ASP v6.0 interoperability

The following product releases have been tested for interoperability with Orbix E2A ASP v6.0:

- Orbix 3.3.4 C++ Edition
- Orbix 3.3.4 Java Edition

Orbix E2A ASP v5.1 interoperability

The following product releases have been tested for interoperability with Orbix E2A ASP v5.1:

- Orbix 3.0.1-82
- OrbixWeb 3.2-15
- Orbix 3.3.2 C++ Edition
- Orbix 3.3.2 Java Edition

The `_bind()` function

Orbix 6.3 does not support the `_bind()` function for establishing connections between clients and servers. Neither Orbix 3.0.1-82, OrbixWeb 3.2-15, nor Orbix 3.3 clients can use the `_bind()` function to establish a connection to an Orbix 6.3 server. You must use a CORBA Naming Service instead. For example, you could use either the Orbix 3.3 naming service or the Orbix 6.3 naming service.

IDL feature support

Orbix 6.3 supports a larger set of IDL data types and features than Orbix 3.3. When developing IDL interfaces for use with Orbix 6.3 and other products you need to restrict your IDL to a subset that is supported by all of the interoperating products.

In particular, the following describe IDL features that are subject to limitations or require special configuration:

- [“Using the #pragma Prefix” on page 65](#)
- [“Use of #pragma ID in IDL” on page 67](#)
- [“Fixed Data Type and Interoperability” on page 68](#)
- [“Use of wchar and wstring” on page 70](#)
- [“C++ Keywords as Operation Names” on page 70](#)

Changed exception semantics

The semantics of some CORBA system exceptions are different in Orbix 6.3, as compared with Orbix 3.0.1-82, OrbixWeb 3.2-15, or Orbix 3.3. If you have existing code written for Orbix 3.0.1-82, OrbixWeb 3.2-15, or Orbix 3.3, you should read the following:

- [“Orbix 3.3 C++ Edition—System Exceptions” on page 73](#)
- [“Orbix 3.3 Java Edition—System Exceptions” on page 78](#)

These sections describe how to configure your legacy application so that it is insulated from any differences in exception semantics.

Bidirectional GIOP

Orbix 6.3 introduces support for bidirectional GIOP, based on an OMG standard. Previously (Orbix E2A ASP v5.x and v6.0), bidirectional GIOP was not supported, or was not based on an OMG standard (Orbix 3.x and earlier).

See [“Callbacks and Bidirectional GIOP” on page 93](#) for details.

Other affected features

If you want to use the Orbix 6.3 interoperable naming service as the common naming service for your interoperating system, see [“The Orbix 6.3 Interoperable Naming Service” on page 85](#).

The rest of this guide describe miscellaneous issues that might affect interoperability in a mixed product environment.

Launch and Invoke Rights

When an Orbix 6.3 client attempts to open a connection to an Orbix 3.0.1-82, OrbixWeb 3.2-15, or Orbix 3.3 server you must make sure that the system is configured such that the Orbix 6.3 client has launch and invoke rights.

Role of launch and invoke rights

In Orbix 3.3 the `orbixd` daemon process is responsible both for launching servers and for redirecting client requests to servers. These two functions are governed by *launch rights* and *invoke rights*, respectively.

Launch and invoke rights on Orbix 3.3 servers are based on the idea that the client *userID* is transmitted along with request messages. The field of the request message that contains the user ID is known as the Principal of the invocation.

If launch and invoke rights are not configured correctly, the Orbix 6.3 client raises a `CORBA::OBJECT_NOT_EXIST` system exception.

Setting launch rights

The launch rights associated with an Orbix 3.3 server specify which users are allowed to cause automatic launching of the server. Launch rights in Orbix 3.3 are granted with the following form of `chmodit`:

```
chmodit 1+userID ServerName
```

Setting invoke rights

The invoke rights associated with an Orbix 3.3 server are used to determine which users are allowed to invoke on the server. Invoke rights are granted using:

```
chmodit i+userID ServerName
```

Orbix 6.3 and Orbix 3.3

The configuration must be altered for an Orbix 6.3 client invoking on an Orbix 3.3 server. There are two possible approaches to fix the launch and invoke rights:

- [Alter the configuration of the Orbix 6.3 Client.](#)
- [Relax the security on the orbixd daemon.](#)

Alter the configuration of the Orbix 6.3 Client

Three configuration variables must be made (or changed) in the Orbix 6.3 configuration file:

```
# Orbix 6.3 Configuration File
policies:giop:interop_policy:send_locate_request = "false";
policies:giop:interop_policy:send_principal      = "true";
policies:giop:interop_policy:enable_principal_service_context =
"true";
```

The `policies:giop:interop_policy:send_locate_request` option controls whether Orbix 6.3 sends `LocateRequest` messages before sending initial `Request` messages. This option must be set to false because `LocateRequest` messages do not contain a `Principal` field.

The `policies:giop:interop_policy:send_principal` option controls whether Orbix 6.3 sends `Principal` information containing the current user name in GIOP 1.0 and GIOP 1.1 requests. The user name is matched against the launch and invoke rights listed in the `orbixd` daemon, to determine the permissions of the Orbix 6.3 client.

Relax the security on the orbixd daemon

Alternatively, you can relax the security on the `orbixd` daemon so that all clients have launch and invoke rights. For example, use the `chmodit` command line utility to change the launch and invoke rights:

```
chmodit l+a11 ServerName
```

```
chmodit i+a11 ServerName
```

These commands give permission for any client to invoke or launch the server *ServerName*. Permissions are granted even if the `Principal` value is left blank in the incoming requests.

GIOP Versions

GIOP version of a connection

The GIOP version used by a client-server connection is determined by the client. When a client is about to open a connection to a CORBA object, the client examines the version information in the object's IOR:

- If the GIOP version in the IOR is greater than or equal to the default GIOP version of the client, the client initiates a connection using the client's default GIOP version.
- Otherwise, the client initiates a connection using the GIOP version in the IOR.

Effect of GIOP version

The GIOP version of a connection is important, because some CORBA features are not supported in early GIOP versions. [Table 17](#) shows the minimum GIOP version required for some CORBA features, according to the CORBA specification.

Table 17: *CORBA-Specified Minimum GIOP Versions*

CORBA Feature	CORBA-Specified Minimum GIOP Version
fixed type	1.1
wchar and wstring types	1.1
codeset negotiation (Orbix 6.3 only)	1.1

Orbix-specific minimum GIOP versions

Notwithstanding the CORBA-specified minimum GIOP versions, Orbix allows some features to be used at a lower GIOP version (in some cases requiring specific configuration variables to be set). [Table 18](#) shows the Orbix-specific minimum GIOP versions.

Table 18: *Orbix-Specific Minimum GIOP Versions*

CORBA Feature	Orbix-Specific Minimum GIOP Version
fixed type	1.0
wchar and wstring types	1.0
codeset negotiation (Orbix 6.3 only)	1.1

For more details on these CORBA features, see these sections:

- [“Fixed Data Type and Interoperability” on page 68.](#)
- [“Use of wchar and wstring” on page 70.](#)
- [“Introduction to Codeset Negotiation” on page 97.](#)

Table of default GIOP versions

Table 19 shows the default GIOP versions for different Orbix clients when opening a connection to a server.

Table 19: *Default GIOP Version Used by Orbix Clients*

Client Version	Default GIOP Version
Orbix 3.0.1-82	1.0
OrbixWeb 3.2-15	1.0
Orbix 3.3 C++ Edition	1.1
Orbix 3.3 Java Edition	1.0
Orbix 6.3	1.1

IDL Issues

This chapter describes those features of IDL that affect interoperability between Orbix 3.x and Orbix 6.3 applications.

Using the #pragma Prefix

Using the `#pragma prefix` preprocessor directive in your IDL affects the semantics of the `_narrow()` function. When an Orbix 3.0.1-82 or Orbix 3.3 C++ client attempts to `_narrow()` an object reference originating from an Orbix 6.3 server, a remote `_is_a()` call is implicitly made.

The `#pragma prefix` preprocessor directive is not fully supported in OrbixWeb 3.2-15 and Orbix 3.3 Java Edition. An OrbixWeb 3.2-15 or Orbix 3.3 Java application can, however, interoperate with Orbix 6.3, with an implicit `_is_a()` call being made by the Orbix runtime.

Effect of #pragma prefix

The `#pragma prefix` directive is used to add a prefix to the `RepositoryId` of all the IDL declarations that follow. For example:

```
//IDL
#pragma prefix "mydomain.com"

interface Foo {
    //Various operations and attributes (not shown)
    ...
};
```

The default `RepositoryId` of the `Foo` interface would be `IDL:Foo:1.0`. When used as above, the `#pragma prefix` causes the `RepositoryId` of the interface `Foo` to change to `IDL:mydomain.com/Foo:1.0`.

C++ code example

Consider, a `Foo` object reference that is generated by an Orbix 6.3 server. The Orbix 6.3 server stringifies the object reference, using the `CORBA::ORB::object_to_string()` operation and writes it to a temporary file.

An Orbix 3.3 C++ client then reads the stringified object reference from the temporary file and converts it back to a `Foo` object reference, as follows:

```
//C++
...
//-----
-
// The following variables are assumed to be initialized
// already:
// 'stringObj' - A stringified object reference of char * type
// 'orbV'      - A reference to an ORB object,
//              of CORBA::ORB_var type
//
try {
    CORBA::Object_var objV = orbV->string_to_object(stringObj);
    // Attempt to 'narrow' the object reference to type
    'Foo_ptr'
    Foo_var myFooV = Foo::_narrow(objV);
    if (CORBA::is_nil(myFooV) ) {
        cerr << "error: narrow to Foo failed" << endl;
        exit(1);
    }
}
catch (CORBA::SystemException& sysEx) {
    ... // deal with exceptions
}
```

Semantics of the `_narrow()` function

When `Foo::_narrow(objV)` is invoked, the object's `RepositoryId` is checked to make sure that it really is of type `Foo`. There are two ways a client can check the type of an object when it performs a `_narrow()`:

- Check the type locally, using the information in the client stub code.
- Check the type remotely, by calling back to the Orbix 6.3 server. The `_is_a()` function is invoked on the remote `Foo` object.

Because the `Foo` object reference originates from an Orbix 6.3 server, the Orbix 3.3 C++ client is unable to check the `RepositoryId` using its local stub code. It must call back to the server instead. The implementation of `_narrow()` calls the remote operation `CORBA::Object::_is_a()` on the object reference `objV`. The `_is_a()` function returns `TRUE` if the object is really of type `Foo`, otherwise it returns `FALSE`.

Effect on the CORBA Naming Service

The naming service is affected because it uses a `#pragma prefix` directive:

```
//IDL for the CORBA Naming Service
#pragma prefix "omg.org"

module CosNaming {
    ...
    interface NamingContext {
        ...
    };
};
```

When used as above, `#pragma prefix` causes the `RepositoryId` of the interface `NamingContext` to change to `IDL:omg.org/CosNaming/NamingContext:1.0`. An Orbix 3.3 C++ client that uses the Orbix 6.3 naming service, therefore, implicitly makes a remote `_is_a()` invocation whenever it invokes `_narrow()` on a naming service object.

Orbix 3.3 C++ Edition and Orbix 6.3

When Orbix 3.3 C++ Edition and Orbix 6.3 applications are mixed in the same system, you can use IDL that has a `#pragma prefix` directive, but the semantic behavior of `_narrow()` is affected.

Orbix 3.3 Java Edition and Orbix 6.3

If a `#pragma prefix` preprocessor directive appears in your IDL, it is ignored by the Orbix 3.3 IDL-to-Java compiler. The Java stub and skeleton code is generated as if the `#pragma prefix` was not there.

When Orbix 3.3 Java Edition and Orbix 6.3 applications are mixed in the same system, you can use IDL that has a `#pragma prefix` directive, but implicit `is_a()` calls are made by the Orbix runtime.

Use of `#pragma ID` in IDL

The `#pragma ID` directive is supported in Orbix 6.3, but is not supported in Orbix 3.3.

Syntax of `#pragma ID`

The `#pragma ID` directive is used to associate an arbitrary repository ID with a given IDL type name. It has the following syntax:

```
#pragma ID TypeName "RepositoryID"
```

The *RepositoryId* must be of the form *Format:String* where no colon can appear in *Format*. For example, if the *Format* of the repository ID is IDL:

```
//IDL
module Example {
    interface Foo {};
#pragma ID Foo "IDL:ArbitraryFooId:1.1"
};
```

The default repository ID that would normally be associated with Foo is IDL:Example/Foo:1.0. By including the #pragma ID directive the repository ID becomes IDL:ArbitraryFooId:1.1 instead.

Orbix 3.3 C++ Edition and Orbix 6.3

IDL that makes use of the #pragma ID directive cannot be used interoperably between Orbix 3.3 C++ Edition and Orbix 6.3 applications.

Orbix 3.3 Java Edition and Orbix 6.3

IDL that makes use of the #pragma ID directive cannot be used interoperably between Orbix 3.3 Java Edition and Orbix 6.3 applications.

Fixed Data Type and Interoperability

When interoperating between an Orbix 3.0.1-82/OrbixWeb 3.2-15 application and an Orbix 6.3 C++ application, it is necessary to change the configuration of Orbix 6.3 in order to be able to use the fixed-point IDL type.

C++ applications

To enable the fixed-point type to be sent between an Orbix 3.0.1-82 application and an Orbix 6.3 application, the following configuration entry must be made (or changed) in the Orbix 6.3 configuration file:

```
# Orbix 6.3 Configuration File
policies:giop:interop_policy:allow_fixed_types_in_1_0 = "true";
```

If set to true, Orbix 6.3 permits fixed-point types to be sent over GIOP 1.0. Defaults to false.

Java applications

Orbix 6.3 accepts fixed-point types through GIOP 1.0 and GIOP 1.1 connections. No special configuration is needed, therefore, when sending fixed-point types between Orbix 6.3 and legacy products such as Orbix 3.0.1-82 or Orbix 3.3.

Orbix 3.0.1-82 and Orbix 6.3

Orbix 3.0.1-82 uses GIOP 1.0 by default and Orbix 6.3 does not permit fixed-point types to be sent over GIOP 1.0. It is necessary, therefore, to reconfigure Orbix 6.3 in this case by setting the `allow_fixed_types_in_1_0` variable to `true`.

Orbix 3.3 C++ Edition and Orbix 6.3

Orbix 6.3 uses GIOP 1.1 by default and Orbix 6.3 permits fixed-point types to be sent over GIOP 1.1. There is, therefore, no need to reconfigure Orbix 6.3 in this case.

Orbix 3.3 Java Edition and Orbix 6.3

To enable the fixed-point type to be sent between Orbix 3.3 Java Edition and Orbix 6.3 applications, two alternative configurations can be used:

- Make, or change, the following configuration entry in the Orbix 6.3 configuration file:

```
# Orbix 6.3 Configuration File
policies:giop:interop_policy:allow_fixed_types_in_1_0 =
  "true";
```

If set to `true`, Orbix 6.3 permits fixed-point types to be sent over GIOP 1.0. Defaults to `false`.

- Alternatively, you can configure Orbix 3.3 Java Edition to use GIOP 1.1, using the `IT_DEFAULT_IIOPI_VERSION` configuration variable. This configuration variable can be set in any of the ways described in the Orbix 3.3 Administrator's Guide. For example, you can set it in the `orbixweb3.cfg` file, as follows:

```
#File: 'orbixweb3.cfg'
OrbixWeb {
  # Other options not shown
  # ...
  IT_DEFAULT_IIOPI_VERSION = "11";
};
```

By setting the `IT_DEFAULT_IIOPI_VERSION` configuration variable to `11` you ensure that Orbix 3.3 Java Edition uses GIOP 1.1 by default on connections to servers. Because GIOP 1.1 officially supports marshalling of fixed-point data, this enables you to use fixed-point data interoperably.

Note: Orbix 3.3 C++ Edition has a similarly named environment variable, `IT_IIOPI_VERSION`. However, setting `IT_IIOPI_VERSION` in Orbix 3.3 C++ Edition does not have the same effect as setting `IT_DEFAULT_IIOPI_VERSION` in Orbix 3.3 Java Edition. The `IT_IIOPI_VERSION` environment variable cannot be used to enable use of the fixed point type between Orbix 3.3 C++ Edition and Orbix 6.3.

Use of wchar and wstring

Table 20 summarizes the support for the `wchar` and `wstring` IDL types in the Orbix 3.3 and Orbix 6.3 products.

Table 20: *Support for the `wchar` and `wstring` Types by Product*

Product	Supports <code>wchar</code>	Supports <code>wstring</code>
Orbix 6.3 (C++)	Yes	Yes
Orbix 6.3 (Java)	Yes	Yes
Orbix 3.3 C++ Edition	No	No
Orbix 3.3 Java Edition	Yes	Yes

All of the products that support `wchar` and `wstring` types can interoperate with each other.

C++ Keywords as Operation Names

Previously, if your IDL contained operation names that were the same as C++ keywords, Orbix 3.0.1-82 and Orbix 3.3 C++ Edition could not interoperate with Orbix 6.3.

This problem is now fixed. Orbix 3.3 applications can now interoperate with Orbix 6.3 even when your IDL contains C++ keywords as operation names.

IDL example

Consider the following IDL:

```
//IDL
interface CplusplusKeywords {
    void    for();
    boolean class();
};
```

C++ stub code

The Orbix 3.3 IDL-to-C++ compiler maps this interface to the following proxy class:

```
//C++
class CplusplusKeywords: public virtual CORBA::Object {
    ...
public:
    ...
    virtual void _for (...);
    virtual CORBA::Boolean _class (...);
    ...
};
```

The names of the functions in C++ have a leading underscore character (for example, `_for` and `_class`) to avoid clashing with the `for` and `class` C++ keywords.

On-the-wire format for operation names

When an Orbix 3.3 C++ or Java client makes a remote invocation using the `_for()` and `_class()` functions, the operation names are marshaled as "for" and "class" respectively. This behavior complies with CORBA 2.6 and is compatible with Orbix 6.3 CORBA servers.

Exceptions

This chapter discusses the differences in the handling of CORBA exceptions between Orbix 3.x and Orbix 6.3.

Orbix 3.3 C++ Edition—System Exceptions

The semantics of system exceptions in Orbix prior to Orbix 3.0.1-20 are different from the semantics in Orbix 6.3. In Orbix 3.0.1-20 and later Orbix 3.x versions, however, exception semantics have been altered to make them compatible with Orbix 6.3. An environment variable, `IT_USE_ORBIX3_STYLE_SYS_EXC`, is introduced that enables you to insulate legacy code from the change.

New Semantics and Old Semantics

Some system exceptions in Orbix 6.3 have different semantics to the corresponding exceptions in Orbix prior to Orbix 3.0.1-20. The exception semantics used by Orbix 6.3 are referred to here as *new semantics*. The exception semantics used by Orbix prior to Orbix 3.0.1-20 are referred to here as *old semantics*.

The `IT_USE_ORBIX3_STYLE_SYS_EXC` Variable

The `IT_USE_ORBIX3_STYLE_SYS_EXC` variable affects three different aspects of Orbix 3.0.1-82 and Orbix 3.3 applications:

- System exceptions raised by the server.
- System exceptions raised by the client.
- Transformation of exceptions arriving at the client.

System exceptions are not only raised by servers, they can also be raised on the client side. If a client encounters an error before it sends a `Request` message to a server, or after it receives a `Reply` message from a server, the client raises a system exception. The `IT_USE_ORBIX3_STYLE_SYS_EXC` variable therefore affects both client and server applications.

System exceptions raised by the server

[Table 21](#) shows how system exceptions raised by an Orbix 3.0.1-82 and an Orbix 3.3 server are influenced by `IT_USE_ORBIX3_STYLE_SYS_EXC`.

Table 21: *Effect of `IT_USE_ORBIX3_STYLE_SYS_EXC` on a Server*

<code>IT_USE_ORBIX3_STYLE_SYS_EXC</code>	Orbix 3.0.1-82 and Orbix 3.3 Server - Exception Raising
<i>Not defined</i>	Old semantics

Table 21: *Effect of IT_USE_ORBIX3_STYLE_SYS_EXC on a Server*

IT_USE_ORBIX3_STYLE_SYS_EXC	Orbix 3.0.1-82 and Orbix 3.3 Server - Exception Raising
YES	Old semantics
NO	New semantics

System exceptions raised by the client

Table 22 shows how system exceptions raised by an Orbix 3.0.1-82 and an Orbix 3.3 client are influenced by IT_USE_ORBIX3_STYLE_SYS_EXC.

Table 22: *Effect of IT_USE_ORBIX3_STYLE_SYS_EXC on a Client*

IT_USE_ORBIX3_STYLE_SYS_EXC	Orbix 3.0.1-82 and Orbix 3.3 Client - Exception Raising
<i>Not defined</i>	Old semantics
YES	Old semantics
NO	New semantics

Transformation of exceptions arriving at the client

Table 23 shows how transformation of exceptions arriving at an Orbix 3.0.1-82 and an Orbix 3.3 client are influenced by IT_USE_ORBIX3_STYLE_SYS_EXC.

Table 23: *Transformation of Exceptions at the Client Side*

IT_USE_ORBIX3_STYLE_SYS_EXC	Orbix 3.0.1-82 and Orbix 3.3 Client - Exception Transformation
<i>Not defined</i>	Transform to old semantics
YES	Transform to old semantics
NO	Transform to new semantics

Transformation is applied to system exceptions incoming from the network. This feature dynamically intercepts system exceptions arriving at the client and, if necessary, converts them to the type of system exception expected by the client (consistent with either new or old semantics). This is essential to ensure that the client can apply a consistent style of exception handling, irrespective of the type of server it is talking to.

Difference between Orbix Prior to Orbix 3.0.1-82 and Orbix 3.3

The presence of the transformation feature means that there is a significant difference between Orbix clients prior to Orbix 3.0.1-20 and Orbix 3.0.1-82/Orbix 3.3 clients even when the `IT_USE_ORBIX3_STYLE_SYS_EXC` variable is not set (or set equal to YES). An Orbix 3.0.1-82 or Orbix 3.3 client that uses old semantics actively transforms incoming system exceptions to old semantics. A pre-Orbix 3.0.1-20 client does not.

The `INV_OBJREF` and `OBJECT_NOT_EXIST` Exceptions

Orbix 6.3 semantics

In Orbix 6.3 the `INV_OBJREF` and `OBJECT_NOT_EXIST` system exceptions are raised under the following circumstances:

- The `INV_OBJREF` system exception is raised by `CORBA::ORB::string_to_object()` to indicate that the stringified object reference is malformed in some way.
- The `OBJECT_NOT_EXIST` system exception is raised by a server to indicate that a CORBA object does not exist.

Orbix 3.3 (new semantics)

In Orbix 3.0.1-82 and Orbix 3.3 (new semantics) the `INV_OBJREF` and `OBJECT_NOT_EXIST` system exceptions are raised under the following circumstances:

- The `INV_OBJREF` system exception is raised for a variety of reasons. However, it is not raised to indicate that a CORBA object does not exist.
- The `OBJECT_NOT_EXIST` system exception is raised by a server to indicate that a CORBA object does not exist.

Pre-Orbix 3.0.1-20 (old semantics)

Prior to Orbix 3.0.1-20 (old semantics) the `INV_OBJREF` and `OBJECT_NOT_EXIST` system exceptions are raised under the following circumstances:

- The `INV_OBJREF` system exception is raised for a variety of reasons. When raised by a server, with minor code 10101, it indicates that a CORBA object does not exist.
- The `OBJECT_NOT_EXIST` system exception is never raised by pre-Orbix 3.0.1-20 applications.

The TRANSIENT and COMM_FAILURE Exceptions

Orbix 6.3 Semantics and Orbix 3.3 (new semantics)

In Orbix 6.3 and in Orbix 3.0.1-82/Orbix 3.3 (new semantics) the `TRANSIENT` and `COMM_FAILURE` system exceptions are raised under the following circumstances:

- The `TRANSIENT` exception is raised if a client tries to send a message to a server, but is unable to do so. In terms of the TCP/IP transport layer, this means an error occurred before or during an attempt to write to or connect to a socket.
- The `COMM_FAILURE` exception is raised if a client has already sent a message to a server, but is unable to receive the associated reply. In terms of the TCP/IP transport layer, this means either the connection went down or an error occurred during an attempt to read from a socket.

Pre-Orbix 3.0.1-20 (old semantics)

Prior to Orbix 3.0.1-20 (old semantics) the `TRANSIENT` and `COMM_FAILURE` system exceptions are raised under the following circumstances:

- The `TRANSIENT` exception is never raised in pre-Orbix 3.0.1-20 applications.
- The `COMM_FAILURE` exception is raised in pre-Orbix 3.0.1-20 applications if an error occurs while writing to, reading from, or connecting to a TCP/IP socket.

Orbix 3.3 C++ Edition and Orbix 6.3

There are three different ways of setting the `IT_USE_ORBIX3_STYLE_SYS_EXC` configuration value:

- Setting an environment variable.
- Setting a configuration variable.
- Using the `SetConfigValue()` function.

Setting an environment variable

Set the environment variable, `IT_USE_ORBIX3_STYLE_SYS_EXC`, as follows:

Windows

```
set IT_USE_ORBIX3_STYLE_SYS_EXC=yes_or_no
```

UNIX

```
export IT_USE_ORBIX3_STYLE_SYS_EXC=yes_or_no
```

Where `yes_or_no` can be the string `YES` or the string `NO`.

Setting a configuration variable

Set the configuration variable, `IT_USE_ORBIX3_STYLE_SYS_EXC`, by editing the Orbix 3.3 configuration file:

```
# Orbix 3.3 Configuration File
Orbix {
    IT_USE_ORBIX3_STYLE_SYS_EXC = "yes_or_no";
};
```

Using the SetConfigValue() function

Use the `CORBA::ORB::SetConfigValue()` function:

```
// C++
orb_p->SetConfigValue(
    "Orbix.IT_USE_ORBIX3_STYLE_SYS_EXC",
    "yes_or_no"
);
```

Where `orb_p` is a pointer to a `CORBA::ORB` instance.

Compatibility matrix

Table 24 shows the compatibility matrix between Orbix 3.0.1-82/Orbix 3.3 and Orbix 6.3.

Table 24: System Exception Handling Compatibility between Orbix 3.0.1-82/Orbix 3.3 and Orbix 6.3

Client Application	Orbix 3.0.1-82/Orbix 3.3 Server (Old Semantics)	Orbix 3.0.1-82/Orbix 3.3 Server (New Semantics)	Orbix 6.3 CORBA Server
Orbix 3.0.1-82/Orbix 3.3 Client (Old Semantics)	Yes	Yes	Yes
Orbix 3.0.1-82/Orbix 3.3 Client (New Semantics)	Yes	Yes	Yes
Orbix 6.3 CORBA Client	No	Yes	Yes

A *Yes* entry in the above table indicates compatible exception semantics for that combination.

An Orbix 3.0.1-82/Orbix 3.3 application described in the table as old semantics has its `IT_USE_ORBIX3_STYLE_SYS_EXC` variable set equal to `YES`, or unset. An Orbix 3.0.1-82/Orbix 3.3 application described in the table as new semantics has its `IT_USE_ORBIX3_STYLE_SYS_EXC` variable set equal to `NO`.

Orbix 3.3 Java Edition—System Exceptions

The semantics of system exceptions in OrbixWeb prior to OrbixWeb 3.2-05 are different from the semantics in Orbix 6.3. In OrbixWeb 3.2-15 and Orbix 3.3 Java Edition, however, exception semantics have been altered to make them compatible with Orbix 6.3. An environment variable, `IT_USE_ORBIX3_STYLE_SYS_EXC`, is introduced that enables you to insulate legacy code from the change.

New Semantics and Old Semantics

Some system exceptions in Orbix 6.3 have different semantics to the corresponding exceptions in OrbixWeb prior to OrbixWeb 3.2-05. The exception semantics used by Orbix 6.3 are referred to here as new semantics. The exception semantics used by OrbixWeb prior to OrbixWeb 3.2-05 are referred to here as old semantics.

The `IT_USE_ORBIX3_STYLE_SYS_EXC` variable

The `IT_USE_ORBIX3_STYLE_SYS_EXC` variable affects two aspects of OrbixWeb 3.2-15 and Orbix 3.3 Java Edition applications:

- System exceptions raised by the server.
- System exceptions raised by the client.

The `IT_USE_ORBIX3_STYLE_SYS_EXC` variable therefore affects both client and server applications.

Note: OrbixWeb 3.2-15 and Orbix 3.3 Java applications do not perform transformations on incoming system exceptions.

System exceptions raised by the server

System exceptions raised by an OrbixWeb 3.2-15/Orbix 3.3 Java server are influenced in the following way by

`IT_USE_ORBIX3_STYLE_SYS_EXC`.

Table 25: *Effect of `IT_USE_ORBIX3_STYLE_SYS_EXC` on a Server*

<code>IT_USE_ORBIX3_STYLE_SYS_EXC</code>	Orbix 3.3 Java Server - Exception Raising
<i>Not defined</i>	Old semantics
TRUE	Old semantics
FALSE	New semantics

System exceptions raised by the client

System exceptions raised by an OrbixWeb 3.2-15/Orbix 3.3 Java client are influenced in the following way by `IT_USE_ORBIX3_STYLE_SYS_EXC`.

Table 26: *Effect of `IT_USE_ORBIX3_STYLE_SYS_EXC` on a Client*

<code>IT_USE_ORBIX3_STYLE_SYS_EXC</code>	Orbix 3.3 Java Client - Exception Raising
<i>Not defined</i>	Old semantics
TRUE	Old semantics
FALSE	New semantics

The `INV_OBJREF` and `OBJECT_NOT_EXIST` Exceptions

Orbix 6.3 Semantics and Orbix 3.3 Java Edition (new semantics)

In Orbix 6.3 and OrbixWeb 3.2-15/Orbix 3.3 Java Edition (new semantics) the `INV_OBJREF` and `OBJECT_NOT_EXIST` system exceptions are raised under the following circumstances:

- The `INV_OBJREF` system exception is raised by `CORBA::ORB::string_to_object()` to indicate that the stringified object reference is malformed in some way.
- The `OBJECT_NOT_EXIST` system exception is raised by a server to indicate that a CORBA object does not exist.

Orbix 3.3 Java Edition (old semantics)

In OrbixWeb 3.2-15/Orbix 3.3 Java Edition (old semantics) the `INV_OBJREF` and `OBJECT_NOT_EXIST` system exceptions are raised under the following circumstances:

- The `INV_OBJREF` system exception, with minor code 10100, is raised by a server to indicate that a CORBA object does not exist.
- The `OBJECT_NOT_EXIST` system exception is never raised in OrbixWeb 3.2-15/Orbix 3.3 Java Edition.

The `TRANSIENT` and `COMM_FAILURE` Exceptions

Orbix 6.3 Semantics and Orbix 3.3 Java Edition (new semantics)

In Orbix 6.3 and OrbixWeb 3.2-15/Orbix 3.3 Java Edition (new semantics) the `TRANSIENT` and `COMM_FAILURE` system exceptions are raised under the following circumstances:

- The `TRANSIENT` exception is raised if a client tries to send a message to a server but is unable to do so. In terms of the TCP/IP transport layer, this means an error occurred before or during an attempt to write to or connect to a socket.
- The `COMM_FAILURE` exception is raised if a client has already sent a message to a server but is unable to receive the associated reply. In terms of the TCP/IP transport layer, this means either the connection went down or an error occurred during an attempt to read from a socket.

Orbix 3.3 Java Edition (old semantics)

In OrbixWeb 3.2-15/Orbix 3.3 Java Edition (old semantics) the `TRANSIENT` and `COMM_FAILURE` system exceptions are raised under the following circumstances:

- The `TRANSIENT` exception can be raised in an OrbixWeb 3.2-15/Orbix 3.3 Java client when attempting to make a connection through Orbix Wonderwall, or when attempting to deal with a `LOCATION_FORWARD` Reply message.
- The `COMM_FAILURE` exception is raised in OrbixWeb 3.2-15/Orbix 3.3 Java Edition if an error occurs while writing to, reading from, or connecting to a TCP/IP socket.

Orbix 3.3 Java Edition and Orbix 6.3

Setting the `IT_USE_ORBIX3_STYLE_SYS_EXC` variable

The `IT_USE_ORBIX3_STYLE_SYS_EXC` variable can be set in any of the ways described in the OrbixWeb Administrator's Guide.

For example, to switch on new semantics you can make the following entry in the `OrbixWeb3.cfg` configuration file:

```
# Orbix 3.3 Configuration File
OrbixWeb.IT_USE_ORBIX3_STYLE_SYS_EXC = "FALSE";
```

Compatibility matrix

Table 27 shows the compatibility matrix between OrbixWeb 3.2-15/Orbix 3.3 Java Edition and Orbix 6.3.

Table 27: *System Exception Handling Compatibility between OrbixWeb 3.2-15/Orbix 3.3 Java Edition and Orbix 6.3*

Client Application	OrbixWeb 3.2-15/Orbix 3.3 Java Server (Old Semantics)	OrbixWeb 3.2-15/Orbix 3.3 Java Server (New Semantics)	Orbix 6.3 CORBA Server
OrbixWeb 3.2-15/Orbix 3.3 Java Client (Old Semantics)	Yes	No	No
OrbixWeb 3.2-15/Orbix 3.3 Java Client (New Semantics)	No	Yes	Yes
Orbix 6.3 CORBA Client	No	Yes	Yes

A Yes entry in the above table indicates compatible exception semantics for that combination.

An OrbixWeb 3.2-15/Orbix 3.3 Java application described in the table as old semantics has its `IT_USE_ORBIX3_STYLE_SYS_EXC` variable set equal to `TRUE`, or unset. An OrbixWeb 3.2-15/Orbix 3.3 Java application described in the table as new semantics has its `IT_USE_ORBIX3_STYLE_SYS_EXC` variable set equal to `FALSE`.

FILTER_SUPPRESS Exception

The `FILTER_SUPPRESS` exception is a system exception specific to Orbix and OrbixWeb. If an Orbix 3.3 C++ server or an Orbix 3.3 Java server sends the `FILTER_SUPPRESS` exception to an Orbix 6.3 CORBA client, it is converted to the standard system exception `CORBA::UNKNOWN`.

Purpose of the FILTER_SUPPRESS exception

Filters are a proprietary feature of Orbix 3.3 that enable you to read and manipulate all incoming and outgoing messages. Prior to the availability of a standard CORBA Security Service, some applications used filters to implement a rudimentary security mechanism. These legacy applications could block the execution of an operation on the server side, by raising the `FILTER_SUPPRESS` exception in a filter.

How Orbix 6.3 handles a `FILTER_SUPPRESS` exception

When a `FILTER_SUPPRESS` exception is sent back to an Orbix 6.3 CORBA client, the Orbix 6.3 CORBA client does not recognize the exception. A `CORBA::UNKNOWN` system exception is raised instead by the Orbix 6.3 CORBA client.

Dynamic Invocation Interface and User Exceptions

The dynamic invocation interface (DII) in Orbix 3.3 cannot handle CORBA user exceptions.

Orbix 3.3 and user exceptions

If a user exception is received by an Orbix 3.3 invocation, the Orbix 3.3 runtime converts the exception into a `CORBA::UNKNOWN` system exception, which is then thrown by the `CORBA::Request::invoke()` operation.

Handling user exceptions in Orbix 3.3 C++ Edition

Given an initialized request object, `req`, the following example shows an outline of how to deal with user exceptions in the DII:

```
// C++ - Orbix 3.3
// Initialize DII Request object, req.
...
// Make the invocation
try {
    req.invoke();
}
catch (...) {
    // You will reach this point if a user exception is thrown.
    ...
}
```


Handling user exceptions in Orbix 3.3 Java Edition

Given an initialized request object, `req`, the following example shows an outline of how to deal with user exceptions in the DII:

```
// Java - Orbix 3.3
// Initialize DII Request object, req.
...
// Make the invocation
try {
    req.invoke();
}
catch (java.lang.Exception) {
    // You will reach this point if a user exception is thrown.
    ...
}
```

Orbix 6.3 and user exceptions

In the Orbix 6.3 DII, however, user exceptions are supported in the DII. The `CORBA::UnknownUserException` standard exception class holds a `CORBA::Any` which can then be parsed with the aid of the dynamic any module to obtain the contents of the user exception.

Dynamic Invocation Interface and LOCATION_FORWARD

The dynamic invocation interface (DII) in Orbix 3.3 C++ Edition is now able to handle reply messages that have the `LOCATION_FORWARD` status. Previously, `LOCATION_FORWARD` replies were not supported in Orbix C++ applications.

The DII in Orbix 3.3 Java Edition has always been able to handle reply messages that have the `LOCATION_FORWARD` status.

See also [“Multiple LOCATION_FORWARD” on page 95](#).

Location forwarding mechanisms

The IIOP protocol features support for location forwarding. It is used to dynamically discover the location of CORBA objects. There are two distinct kinds of message exchange that form the basis of location forwarding:

- The client ORB can deliberately probe the location of a CORBA object, by sending a `LocateRequest` message to the server (or agent). The server (or agent) responds with a `LocateReply` message containing details of the object's location.
- When a client sends a regular `Request` message, the server (or agent) might respond with a special type of `Reply` message that has a reply status of `LOCATION_FORWARD`. This reply has details of the object's location.

Support for location forwarding

The location forward mechanism is used by the Orbix 3.3 daemon and the Orbix 6.3 locator service to direct clients to the true location of a CORBA server:

- The first type of message exchange is a `LocateRequest` followed by `LocateReply`.
- The second type of message exchange is a `Request` followed by a `Reply` with status `LOCATION_FORWARD`.

Both kinds of message exchange are supported in Orbix 3.3.

Services

In a mixed system with Orbix 3.x and Orbix 6.3 applications, you generally have a choice between an Orbix 3.x or an Orbix 6.3 implementation of a CORBA service. This chapter discusses the viable configurations of CORBA services in a mixed system.

The Orbix 6.3 Interoperable Naming Service

The naming service provided with Orbix 6.3 is an implementation of the CORBA Interoperable Naming Service (INS) specification. This section explains how to set up Orbix 3.3 applications to use the Orbix 6.3 INS.

Old and new naming services

In an environment that mixes Orbix 3.3 and Orbix 6.3 applications, you have a choice between using the old CORBA Naming Service (NS), provided with Orbix 3.3, or the new CORBA Interoperable Naming Service (INS), provided with Orbix 6.3.

The NamingContextExt interface

The main difference between the old and new naming services is that the INS adds a new IDL `CosNaming::NamingContextExt` interface:

```
// File: CosNaming.idl
#pragma prefix "omg.org"

module CosNaming {
    ...
    interface NamingContextExt : NamingContext {
        typedef string StringName;
        typedef string Address;
        typedef string URLString;

        StringName to_string (in Name n)
            raises (InvalidName);
        Name to_name (in StringName sn)
            raises (InvalidName);

        exception InvalidAddress {};

        URLString to_url (in Address addr, in StringName sn)
            raises (InvalidAddress, InvalidName);
        Object resolve_str (in StringName sn)
            raises (NotFound, CannotProceed, InvalidName);
    };
};
```

Stub code

Applications that use the INS should preferably be built against the new naming stub (generated from the INS IDL). This makes the new `NamingContextExt` interface accessible. However, the old naming stubs (generated from the old NS IDL) can also be used.

Narrowing and remote `_is_a()` operation

When an Orbix 3.3 application invokes `CosNaming::NamingContext::_narrow()` on an Orbix 6.3 `NamingContext`, it makes a remote `_is_a()` invocation on the INS. The `_is_a()` invocation is used to confirm the type of the `NamingContext` object reference. See ["Using the #pragma Prefix" on page 65](#).

Orbix 3.3 and Orbix 6.3

You can configure Orbix 3.3 to use both the Orbix 3.3 NS and the Orbix 6.3 INS. This section describes how to configure the CORBA Initialization Service to obtain a reference to either naming service using the `CORBA::ORB::resolve_initial_references()` function.

Configuring Orbix 3.3 to use the Orbix 6.3 INS

To connect to both the Orbix 3.3 NS and the Orbix 6.3 INS from an Orbix 3.3 application you must first configure the initialization service. Edit the `common.cfg` configuration file and make the following entries in the `Common.Services` scope:

```
# Orbix 3.3 Configuration File
Common {
  Services {
    # This is the stringified IOR for the root 'NamingContext'
    # of the 'Orbix 3' naming service.
    # You can obtain this IOR by running the naming service
    # as follows:
    #   ns -I <iorfile>
    NameService = "IOR:1234.....";

    # This is the stringified IOR for the root 'NamingContext'
    # of the 'Orbix 6.3' Interoperable Naming Service.
    # You can obtain this IOR using the Orbix 6.3 admin
    # utility as follows:
    #   itadmin ns resolve
    INS = "IOR:4567.....";
  };
};
```

Orbix 3.3 configuration variables

The following configuration variables are set in the `Common.Services` scope:

- The `Common.Services.NameService` configuration variable is set to a stringified IOR for a `NamingContext` in the Orbix 3 NS.
- The `Common.Services.INS` configuration variable is set to a stringified IOR for a `NamingContext` in the Orbix 6.3 INS.

Setting the `Common.Services.INS` variable

For example, consider the following IOR string:

```
IOR:010000002f00000049444c3a696f6e612e636f6d2f49545f4e616d696e672f49545f4e616d696e67436f6e746578744578743a312e3000001000000000000006e00000010102000b00000031302e322e312e31313300008a1300003f0000003a3e0232311744656661756c74204c6f636174696f6e20446f6d61696e185f64656661756c745f69745f6e635f6578745f706f615f000800000000000020000100000600000006000000010000003500
```

You can assign this IOR string to `Common.Services.INS` as follows:

```
# Orbix 3.3 Configuration File
Common {
    Services {
        INS =
        "IOR:010000002f00000049444c3a696f6e612e636f6d2f49545f4e616d696e672f49545f4e616d696e67436f6e746578744578743a312e3000001000000000000006e00000010102000b00000031302e322e312e31313300008a1300003f0000003a3e0232311744656661756c74204c6f636174696f6e20446f6d61696e185f64656661756c745f69745f6e635f6578745f706f615f000800000000000020000100000600000006000000010000003500";
    };
};
```

Orbix 3.3 client code for using both naming services

The following C++ code extract shows how an Orbix 3.0.1-20 application can make an initial connection to both naming services:

```
// C++ - Orbix 3 Client Code
int
main (int argc, char *argv[])
{
    CORBA::ORB_var orbV;

    try
    {
        cout << "Initializing the ORB." << endl;
        orbV = CORBA::ORB_init(argc, argv, "Orbix");

        CosNaming::NamingContext_var orbix3RootContextV;
        CosNaming::NamingContext_var orbix2000RootContextV;
        CORBA::Object_var objV;

        try
        {
            objV = orbV->resolve_initial_references("NameService");
            orbix3RootContextV = CosNaming::NamingContext::_narrow(objV);

            objV = orbV->resolve_initial_references("INS");
            orbix2000RootContextV = CosNaming::NamingContext::_narrow(objV);
        }
        catch (CORBA::SystemException &sysEx)
        {
            cerr << &sysEx << endl;
            return 1;
        }
        ...
    }
}
```

After this code runs, `orbix3RootContextV` holds a reference to an Orbix 3 `NamingContext` and `orbix2000RootContextV` holds a reference to an Orbix 6.3 `NamingContext`.

Orbix 3.3 Java Edition and Orbix 6.3

The following steps describe how to configure Orbix 3.3 Java Edition to connect to both the Orbix 3.3 NS and the Orbix 6.3 INS:

Step	Action
1	Obtain the IOR for the root naming context of the naming service.
2	Connect to the Orbix 3.3 NS and the Orbix INS.

Step 1—obtain the IOR

Obtain the IOR for the root naming context of the NS.

Start the Orbix 6.3 INS and enter the following command:

```
itadmin ns resolve > Naming.ref
```

The output of this command is an IOR string that looks similar to the following:

```
IOR:010000002f00000049444c3a696f6e612e636f6d2f49545f4e616
d696e672f49545f4e616d696e67436f6e746578744578743a312e3
0000001000000000000006e00000010102000b00000031302e322
e312e31313300008a1300003f0000003a3e0232311744656661756
c74204c6f636174696f6e20446f6d61696e185f64656661756c745
f69745f6e635f6578745f706f615f000800000000000020000010
00000600000006000000010000003500
```

This is the IOR string for the root naming context of the Orbix 6.3 INS.

Step 2—connect to the naming services

Connect to the Orbix 3.3 NS and the Orbix INS.

The following Java code shows how an Orbix 3.3 Java client connects to both the Orbix 3.3 NS and the Orbix 6.3 INS:

```
//Java
NamingContext OWrootContext = null;

try {
    org.omg.CORBA.Object ncOWeb =
        orb_wrapper.get_orb().resolve_initial_references(
            "NameService"
        );
    OW32rootContext = NamingContextHelper.narrow(ncOWeb);

    // read the ART Naming IOR from the file:
    String objRef = null;
    BufferedReader br = null;

    try {
        br = new BufferedReader( new FileReader("Naming.ref") );
        objRef = br.readLine();
    } catch (IOException e) {
        System.err.println(
            "IOException caught: " + e.toString()
        );
        ioe = new IOException();
    } finally {
        try {
            br.close();
        } catch (IOException ignore) { }
    }
}
```

```

    org.omg.CORBA.Object objNaming =
    orb.string_to_object(objRef);
    O2KRootContext = NamingContextHelper.narrow(objNaming);
} catch (SystemException ex) {
    System.err.println ("Exception caught during bind : " +
    ex.toString());
    System.exit (1);
} catch (org.omg.CORBA.ORBPackage.InvalidName in) {
    System.err.println ("Exception during narrow of initial
    reference : " + in.toString());
    System.exit (1);
}

```

This code reads the stringified IOR for the Orbix 6.3 NS from the Naming.ref file. The stringified IOR is converted to an object reference, O2KRootContext, using the org.omg.CORBA.ORB.string_to_object() function. The O2KRootContext object reference is used to access the root NamingContext of the Orbix 6.3 INS.

Interface Repository Interoperability

Significant changes were made to the IDL definition of the Interface Repository (IFR) between CORBA 2.2 and CORBA 2.3. The Orbix 6.3 IFR is written to conform to the CORBA 2.6 specification and it has many advantages over the Orbix 3.3 IFR. If you have both Orbix 3.3 and Orbix 6.3 applications that use the IFR, it is recommended that you change the Orbix 3.3 applications to use the Orbix 6.3 IFR.

Modifying Orbix 3.3 applications to use the Orbix 6.3 IFR

To change an Orbix 3.3 C++ application to use the Orbix 6.3 IFR, perform the following steps:

1. Take the IDL for the Orbix 6.3 IFR and generate stub code from it using the Orbix 3.3 IDL compiler.
2. Modify the source code of your Orbix 3.3 application to be consistent with the IDL for the Orbix 6.3 IFR.
3. Link your Orbix 3.3 application with the IFR stub code generated in step 1.

SSL/TLS Toolkit Interoperability

Orbix 3.3 to Orbix 6.3 interoperability

Orbix version 3.3 or later is recommended for secure interoperability with Orbix 6.3 SSL/TLS. Both C++ and Java editions of Orbix 3.3 have been tested with Orbix 6.3 SSL/TLS. There are no known SSL-related interoperability problems affecting this product combination.

Orbix 6.3 Interoperability with Orbix 2000

Orbix 6.3 SSL/TLS (both C++ and Java) has been tested for secure interoperability with Orbix 2000 versions 1.2 and 2.0. There are no known SSL-related interoperability problems.

High Availability and Orbix 3.3 Clients

High availability is a feature of Orbix 6.3 that provides fault tolerance by grouping servers into *server clusters*. Orbix 3.3 clients (C++ and Java Editions) are now able to interoperate with Orbix 6.3 server clusters.

Support for multi-profile IORs

In Orbix 3.3.2 the client ORB iterates over a multi-profiled IOR until it is able to establish a connection to a server. It always starts at the first profile, when connecting or reconnecting to a server.

Connection Management

There are some differences in connection management between Orbix 3.x and Orbix 6.3 applications. In most cases these differences are unimportant, but a minority of applications might be affected.

Orbix 6.3 Active Connection Management

Orbix 6.3 has a feature called active connection management (ACM) that is used to limit the number of open connections on an Orbix 6.3 application.

The Orbix 6.3 ACM feature has been interoperably tested with Orbix 3.3 and found to be fully compatible.

Configuring the ACM

To configure ACM in Orbix 6.3, edit the configuration file, making the following additional entries:

```
# Orbix 6.3 Configuration File
plugins:iiop:incoming_connections:hard_limit = "InHardLimit";
plugins:iiop:incoming_connections:soft_limit = "InSoftLimit";
plugins:iiop:outgoing_connections:hard_limit = "OutHardLimit";
plugins:iiop:outgoing_connections:soft_limit = "OutSoftLimit";
```

A value of -1 indicates that there is no limit on the number of connections.

Callbacks and Bidirectional GIOP

Orbix 6.3 supports bidirectional GIOP. This is a new feature introduced since Orbix E2A ASP v 6.0.

Motivation for bi-directional IIOP

Bidirectional GIOP was introduced in Orbix in order to overcome the limitations of standard GIOP in relation to using callback objects through a firewall.

Features

Micro Focus's implementation of bidirectional GIOP has the following features:

1. Compliant with the modified bidirectional GIOP approach described in the firewall submission.
2. Compatible with GIOP 1.2 (that is, not dependent on GIOP 1.4 NegotiateSession messages).

3. Decoupled from IIOP, so that it can be used over arbitrary connection-oriented transports (for example, SHMIOP).
4. Supports weak `BiDirIds` initially.
5. Supports bidirectional invocations on legacy Orbix 3.x callback object references in order to facilitate phased migration to Orbix 6.3.

References

For more details about the bidirectional GIOP support in Orbix 6.3, see the following references:

- *CORBA Programmer's Guide*
- *Administrator's Guide*

Setting the Listen Queue Size in Orbix 3.3 C++ Edition

A new configuration variable, `IT_LISTEN_QUEUE_SIZE`, is defined in Orbix 3.3 C++ Edition. It allows you to set the size of the queue associated with listening ports on an Orbix 3.3 C++ server. This is a useful optimization for a heavily loaded server that might receive many connection attempts in a short time.

Listen queue size

When an Orbix server wants to receive connections from clients, it needs to call the `listen(int, int)` socket function. The second parameter of `listen()` sets the listen queue size associated with the socket. The listen queue size determines the maximum length that the queue of pending connections can grow to. In Orbix 3.3, the queue length is 5, by default.

The `IT_LISTEN_QUEUE_SIZE` configuration variable

Orbix 3.3 C++ Edition supports a new `IT_LISTEN_QUEUE_SIZE` configuration variable that enables you to configure the listen queue size. It can be set subject to the following constraints:

- The value should lie between 5 and 2000 (inclusive).
- If it is set to a value less than 5, the value 5 is used instead.
- If it is set to a value greater than 2000, the value 2000 is used instead.

Queue size hard limit

The maximum queue size is subject to a hard limit that varies between platforms:

- *Solaris*—there is currently no limit.
- *HPUX*—the limit is 20.

- *Windows*—the limit is 5.

Setting the listen queue size

There are three different ways to set the `IT_LISTEN_QUEUE_SIZE` configuration value:

- Set the `IT_LISTEN_QUEUE_SIZE` environment variable:

Windows

```
set IT_LISTEN_QUEUE_SIZE=QueueSize
```

UNIX

```
export IT_LISTEN_QUEUE_SIZE=QueueSize
```

- Set the `IT_LISTEN_QUEUE_SIZE` configuration variable by editing the Orbix 3.3 configuration file, as follows:

```
# Orbix 3.3 Configuration File
Orbix {
    IT_LISTEN_QUEUE_SIZE = "QueueSize";
};
```

- Use the `CORBA::ORB::SetConfigValue()` function:

```
// C++
orb_p->SetConfigValue(
    "Orbix.IT_LISTEN_QUEUE_SIZE",
    "QueueSize"
);
```

Where `orb_p` is a pointer to a `CORBA::ORB` instance.

Querying the listen queue size

An application can query the value of `IT_LISTEN_QUEUE_SIZE`, using the following code:

```
// C++
char* value = 0;
CORBA::Orbix.GetConfigValue("Orbix.IT_LISTEN_QUEUE_SIZE", value);
cout << endl << "Listen Queue size is " << value << endl;
// Caller is responsible for memory allocated
// in out parameter to GetConfigValue
//
delete[] value;
value = 0;
```

Multiple LOCATION_FORWARD

When an Orbix 3.3 C++ client attempts to connect to a server, it can deal with at most one `LOCATION_FORWARD` reply on a single request. In some cases, this limit might be exceeded when an Orbix 3.3 client attempts to connect to an Orbix 6.3 CORBA server.

An Orbix 3.3 Java client can deal with an infinite number of `LOCATION_FORWARD` replies on a single request.

Description

In a pure Orbix 3.3 environment, the only time a `LOCATION_FORWARD` reply can be generated is when an Orbix 3.3 client contacts the Orbix daemon. In Orbix 6.3, any CORBA server can generate a `LOCATION_FORWARD` reply. It is, therefore, possible that the limit of a single `LOCATION_FORWARD` could be exceeded when an Orbix 3.3 client attempts to connect to an Orbix 6.3 CORBA server.

Summary

[Table 28](#) summarizes the handling of multiple `LOCATION_FORWARD` reply messages.

Table 28: *Number of `LOCATION_FORWARD` Replies that Can Be Handled by Orbix Products*

Product	Maximum Number of <code>LOCATION_FORWARD</code> Replies
Orbix 3.3 C++ Edition	1
Orbix 3.3 Java Edition	<i>Infinity</i>
Orbix 6.3	<i>Infinity</i>

Codeset Negotiation

Codeset negotiation enables CORBA applications to agree on a common character set for transmission of narrow and wide characters.

Introduction to Codeset Negotiation

The CORBA codeset conversion framework enables applications to ensure that they communicate using compatible character formats for both narrow characters, `char`, and wide characters, `wchar`.

Support for codeset negotiation

Orbix 2000 (version 1.1 and later) and Orbix 6.3 support codeset negotiation, as defined by the CORBA 2.6 specification.

Neither Orbix 3.3 nor Orbix 2000 version 1.0 support codeset negotiation.

Servers and codeset negotiation

A server that supports codeset negotiation appends a list of supported codesets (character formats) to the interoperable object references (IORs) it generates. The codesets are placed in standard `IOP::TAG_CODE_SETS` components in the IOR.

Clients and codeset negotiation

A client that supports codeset negotiation examines an IOR to check the list of codesets supported by the server. The client compares this list with its own list of supported codesets and, if a match is found, the client chooses the pair of transmission codesets (narrow character format and wide character format) to use for that particular connection.

When sending a `Request` message, the client appends an `IOP::CodeSets` service context that tells the server which codesets are used. The client continues to include an `IOP::CodeSets` service context in `Request` messages until the first `Reply` message is received from the server. Receipt of the first server `Reply` message implicitly indicates that codeset negotiation is complete. The same characters formats are used for subsequent communication on the connection.

Configuring Codeset Negotiation

Orbix 6.3 features greatly enhanced support for internationalization and codeset negotiation. In particular, it is now possible to specify explicitly the codesets that a server exports in an IOR.

CORBA configuration variables

Table 29 gives the configuration variables that are used to specify the codesets for an Orbix 6.3 CORBA application.

Table 29: CORBA Codeset Configuration Variables (Orbix 6.3)

Configuration Variable	Description
<code>plugins:codeset:char:ncs = "<codeset>";</code>	Specifies the native narrow character codeset.
<code>plugins:codeset:char:ccs = ["<codeset1>", "<codeset2>", ...];</code>	Specifies the list of conversion narrow character codesets supported.
<code>plugins:codeset:wchar:ncs = "<codeset>";</code>	Specifies the native wide character codeset.
<code>plugins:codeset:wchar:ccs = ["<codeset1>", "<codeset2>", ...];</code>	Specifies the list of conversion wide character codesets supported.
<code>plugins:codeset:always_use_default = <boolean>;</code>	Specifies that hardcoded default values are used and the preceding variables are ignored, if set in the same configuration scope or higher.

Default Codesets

This section describes the default codesets used by the Orbix 6.3 product. The following default codesets are defined:

- [CORBA C++ codesets for non-MVS platforms.](#)
- [CORBA C++ codesets for MVS platform.](#)
- [CORBA Java codesets for US-ASCII locale.](#)
- [CORBA Java codesets for Shift_JIS locale.](#)
- [CORBA Java codesets for EUC-JP locale.](#)
- [CORBA Java codesets for other locales.](#)

Native and conversion codesets

Native codesets are used by the application to pass `char` and `wchar` data to the ORB.

Conversion codesets are used, where necessary, to facilitate interoperability with other ORBs or platforms.

CORBA C++ codesets for non-MVS platforms

Table 30 shows the default codesets for Orbix 6.3 C++ applications on non-MVS platforms (Latin-1 locale).

Table 30: *CORBA C++ Codesets (Non-MVS Platforms)*

Codeset Type	Codeset
Native codeset for char (NCS-C)	ISO-8859-1
Conversion codesets for char (CCS-C)	<i>none</i>
Native codeset for wchar (NCS-W)	UCS-2 or UCS-4
Conversion codesets for wchar (CCS-W)	UTF-16

In Orbix 6.3, the choice of native wide character codeset, UCS-2 or UCS-4, is based on the size of `CORBA::WChar` (either 2 or 4 bytes). On Windows, UCS-2 is used and on most UNIX platforms, UCS-4 is used.

CORBA C++ codesets for MVS platform

Table 31 shows the default codesets for Orbix 6.3 C++ applications on the MVS platform.

Table 31: *CORBA C++ Codesets (Non-MVS Platforms)*

Codeset Type	Codeset
Native codeset for char (NCS-C)	EBCDIC
Conversion codesets for char (CCS-C)	ISO-8859-1
Native codeset for wchar (NCS-W)	UCS-2 or UCS-4
Conversion codesets for wchar (CCS-W)	UTF-16

CORBA Java codesets for US-ASCII locale

Table 32 shows the codesets supported by Orbix 6.3 Java applications in a US-ASCII locale.

Table 32: *CORBA Java Codesets (ISO-8859-1/Cp-1292/US-ASCII locale)*

Codeset Type	Codeset
Native codeset for char (NCS-C)	ISO-8859-1
Conversion codesets for char (CCS-C)	UTF-8
Native codeset for wchar (NCS-W)	UTF-16

Table 32: *CORBA Java Codesets (ISO-8859-1/Cp-1292/US-ASCII locale)*

Codeset Type	Codeset
Conversion codesets for wchar (CCS-W)	UCS-2

CORBA Java codesets for Shift_JIS locale

Table 33 shows the codesets supported by Orbix 6.3 Java applications in a Shift_JIS locale.

Table 33: *CORBA Java Codesets (Shift_JIS locale)*

Codeset Type	Codeset
Native codeset for char (NCS-C)	UTF-8
Conversion codesets for char (CCS-C)	ISO-8859-1 or Shift_JIS or euc_JP
Native codeset for wchar (NCS-W)	UTF-16
Conversion codesets for wchar (CCS-W)	UCS-2 or Shift_JIS or euc_JP

CORBA Java codesets for EUC-JP locale

Table 34 shows the codesets supported by Orbix 6.3 Java applications in a EUC-JP locale.

Table 34: *CORBA Java Codesets (EUC-JP locale)*

Codeset Type	Codeset
Native codeset for char (NCS-C)	UTF-8
Conversion codesets for char (CCS-C)	ISO-8859-1 or Shift_JIS or euc_JP
Native codeset for wchar (NCS-W)	UTF-16
Conversion codesets for wchar (CCS-W)	UCS-2 or Shift_JIS or euc_JP

CORBA Java codesets for other locales

Table 35 shows the codesets supported by Orbix 6.3 Java applications in other locales.

Table 35: *CORBA Java Codesets (other locale)*

Codeset Type	Codeset
Native codeset for char (NCS-C)	UTF-8

Table 35: CORBA Java Codesets (other locale)

Codeset Type	Codeset
Conversion codesets for char (CCS-C)	ISO-8859-1 or <i>file encoding</i>
Native codeset for wchar (NCS-W)	UTF-16
Conversion codesets for wchar (CCS-W)	UCS-2 or <i>file encoding</i>

Configuring Legacy Behavior

Default behavior

By default, the `IOP::TAG_CODE_SETS` tagged component is included in generated IORs and the transmission codesets are negotiated by clients and transmitted through an `IOP::CodeSets` service context. This is the CORBA-defined behavior.

Legacy behavior

Orbix 6.3 (all versions) also provides legacy behavior, to support the scenario where wide character data is communicated between Orbix 6.3 and Orbix 3.3 Java Edition.

Disabling codeset negotiation

The following configuration variable can be used to explicitly disable the codeset negotiation mechanism:

```
# Orbix 6.3 Configuration File
policies:giop:interop_policy:negotiate_transmission_codeset =
"false";
```

The default is `true`.

This is a proprietary setting provided for interoperability with legacy implementations, such as Orbix 3.3 Java Edition. The native codeset for character data, ISO-8859-1 (Latin-1), is used and the overhead of full negotiation is avoided. If wide character data is used, Orbix 6.3 reverts to the UTF-16 transmission codeset.

Enabling wchar transmission on a GIOP 1.0 connections

Passing `wchar` data over GIOP 1.0 can be enabled using the following configuration variable:

```
# Orbix 6.3 Configuration File
policies:giop:interop_policy:allow_wchar_types_in_1_0 = "true";
```

The default is `false`.

The transmission of `wchar` data is not legal in GIOP 1.0, by default.

Index

Symbols

#pragma ID 67
#pragma prefix
 and naming service 67
#pragma prefix, using 65

A

ACM 35
activate_object_with_id() operation 25
activating CORBA objects
 in Application Server Platform 25
 in Orbix 3 25
activation, and the KDM 48
activation modes 55
active connection management
 and interoperability 93
 in Application Server Platform 35
add_member command 54
administration properties
 MaxConsumers 43
 MaxSuppliers 43
allow_fixed_types_in_1_0 variable 68
allow_wchar_types_in_1_0 variable 102
Any constructor 18
any-style clients 41
Any type
 migrating 18
 type-unsafe functions 18
ASN_OID structure 47
authentication filters 31
auth_method_data variable 49
automatic activation 48, 55
automatic work queues 30
AVA interface 46
AVAList interface 46

B

bi-directional IIOP 93
binary checksums 48
_bind() function 59
 and corbaloc URLs 16
 and notification service 40
 and the Orbix 3 locator 27
BOA, replacing with the POA 22

C

C++ function signatures 21
C++ keywords
 in IDL 70
 on-the-wire format 71
C++ mapping changes 18
caching

 and smart proxies 33
 of data using smart proxies 34
callbacks
 POA policies for 17
cat_group command 54
catit command 53
cat_member command 55
catns command 55
CertError interface 46
certificate authorities
 trusted CA list 51
certificate revocation lists
 configuration of 48
 no support for 46
certificates
 interoperability 51
 migrating 50
 PKCS#12 format 50
CertValidator interface 47
CertValidatorPolicy interface 46, 47
char type 97
checksums, in the KDM 48
chmodit command 53
chmodit utility 61
chownit command 53
client_secure_invocation_policy
 variable 48
clustered servers 27
codeset negotiation
 interoperability 101
 support for 97
codesets
 Application Server Platform (C++) 99
 transmission 97
codesetsApplication Server Platform
 (Java) 99, 100
CodeSets service context 97, 101
command-line tools 53
COMM_FAILURE system exception 19, 36
 new semantics 76, 80
 old semantics 76, 80
common.cfg file 86
Common.Services scope 87
compatibility matrix
 for Java applications 81
 for system exceptions 77
concurrent request processing 30
configuration
 active connection management 93
 allow_fixed_types_in_1_0 variable 68
 itadmin utility 53
 IT_IIOP_VERSION variable 69
 IT_LISTEN_QUEUE_SIZE variable 94

- IT_USE_ORBIX3_STYLE_SYS_EXC
 - variable 73
 - security variables 48
 - send_locate_request variable 61
 - send_principal variable 61
- connection management
 - and ACM 35
 - and I/O callbacks 35
- CORBA::Environment parameter
 - migrating 19
- corbaloc URL 16
- CORBA objects, creating and activating 24
- CORBA Security Level 2 45
- CosNotifyComm module 41
- CRL 46

D

- daemons
 - locator 53
 - node daemon 53
 - orbixd 53
- DEF_TIE macro 24
- del_group command 55
- del_member command 55
- deprecated IDL types
 - Principal 13
- DII, See dynamic invocation interface
- documentation
 - .pdf format 4
 - updates on the web 4
- DO_GL_HEALTHCHECK 44
- DO_HEALTHCHECK 44
- dumpconfig command 53
- dynamic any module 83
- dynamic invocation interface
 - and LOCATION_FORWARD reply
 - status 83
 - and user exceptions 82

E

- enable_principal_service_context
 - variable 61
- Environment parameter 19
 - and C++ function signatures 21
 - migrating 19
- EstablishTrustPolicy interface 48
- etherealize() function 32
- Extension interface 47
- ExtensionList interface 47
- external configuration granularity 47

F

- fault tolerance 27, 33
- file descriptor limits
 - and active connection management 36
 - extending 36
- filters
 - and FILTER_SUPPRESS exception 81
 - migrating to Application Server
 - Platform 29
 - typical uses 29

- FILTER_SUPPRESS system exception 81
- firewalls, and bi-directional IIOP 93
- fixed type, interoperating 68

G

- GIOP
 - default version 69
- groupshosts command 54

H

- HealthCheck
 - overview 44
 - simulating in Application Server
 - Platform 45
- high availability 27, 33

I

- I/O Callbacks 34
- IDL
 - C++ keywords appearing in 70
 - wchar type 70
 - wstring type 70
- idl command 54
- idlgcn command 54
- idlj command 54
- IDL migration 60
- IDL-to-C++ mapping
 - and C++ keywords in IDL 70
 - changes 18
- IFR 90
- ifr command 54
- IIOP
 - bi-directional 93
 - IT_DEFAULT_IIOP_VERSION
 - variable 69
- iiop_tls plug-in 47
- implementing CORBA objects
 - inheritance approach 23
 - tie approach 24
- incarnate() function 32
- incoming_connections:hard_limit
 - variable 35
- incoming_connections:soft_limit
 - variable 35
- inheritance approach 23
- initialization service
 - and the Orbix 3 locator 28
 - configuring for naming service 86
- initializeHealthCheck() function 44
- initial references
 - NotificationService object ID 40
- interface repository 90
- internationalization 70, 97
- interoperable naming service
 - interoperability 85
 - new interface 39
- INV_OBJREF system exception
 - migration 19
 - new semantics 75, 79
 - old semantics 75
- invoke rights 60

- IOR, and supported codesets 97
- `_is_a()` function 65, 67
- itadmin utility 53
- IT_ALLOWED_CIPHERSUITES variable 49
- IT_AUTHENTICATE_CLIENTS variable 48
- IT_AVA interface 46
- IT_AVAList interface 46
- IT_BIDIRECTIONAL_IIOB_BY_DEFAULT variable 49
- IT_CACHE_OPTIONS variable 49
- IT_CA_LIST_FILE variable 48
- IT_CertError structure 46
- IT_CERTIFICATE_FILE variable 49
- IT_Certificate interface 46
- IT_CERTIFICATE_PATH variable 49
- IT_CHECKSUM_REPOSITORY variable 49
- IT_CHECKSUMS_ENABLED variable 49
- IT_CIPHERSUITES variable 48
- IT_CRL_ENABLED variable 49
- IT_CRL_List interface 46, 47
- IT_CRL_REPOSITORY variable 49
- IT_CRL_UPDATE_INTERVAL variable 49
- IT_DAEMON_AUTHENTICATES_CLIENTS variable 49
- IT_DAEMON_POLICY variable 49
- IT_DAEMON_UNRESTRICTED_METHODS variable 49
- IT_DEFAULT_IIOB_VERSION variable 69
- IT_DEFAULT_MAX_CHAIN_DEPTH variable 49
- IT_DISABLE_SSL variable 49
- IT_ENABLE_DEFAULT_CERT variable 49
- IT_Extension interface 47
- IT_ExtensionList interface 47
- IT_FILTER_BAD_CONNECTS_BY_DEFAULT variable 49
- IT_IIOB_VERSION variable 69
- IT_INSECURE_REMOTE_INTERFACES variable 48
- IT_INSECURE_SERVERS variable 48
- IT_INVOCATION_POLICY variable 48
- IT_KDM_CLIENT_COMMON_NAMES variable 49
- IT_KDM_ENABLED variable 49
- IT_KDM_PIPES_ENABLED variable 49
- IT_KDM_REPOSITORY variable 49
- IT_KDM_SERVER_PORT variable 49
- IT_LISTEN_QUEUE_SIZE variable 94, 95
- itlocator daemon 53
- IT_MAX_ALLOWED_CHAIN_DEPTH variable 49
- itnode_daemon daemon 53
- itnotifyconsole utility 44
- IT_OID structure 47
- IT_OIDTag type 47
- IT_ORBIX_BIN_SERVER_POLICY variable 49
- IT_SECURE_REMOTE_INTERFACES variable 48
- IT_SECURE_SERVERS variable 48
- IT_SERVERS_MUST_AUTHENTICATE_CLIENTS variable 48

- IT_SSL interface 47
- IT_TLS_API interface 46
- IT_TLS interface 46
- IT_USE_ORBIX3_STYLE_SYS_EXC variable
 - in C++ 73
 - in Java 78
 - setting for Java applications 80
 - setting in C++ 76
- IT_UTCTime interface 47
- IT_ValidateX509CertCB interface 47
- IT_X509CertChain interface 47
- IT_X509Cert interface 47
- IT_X509_CRL_Info interface 46, 47
- IT_X509_Revoked interface 46, 47
- IT_X509_RevokedList interface 46, 47

K

- KDM 48
- key distribution management 48
- killit command 54

L

- launch rights 60
- lease plug-in, and session management 36
- level 2, security 45
- listen queue size
 - range 94
 - setting for C++ applications 95
- list_groups command 55
- list_members command 55
- load() function 32
- load balancing
 - and activation modes 56
 - and the CORBA Naming Service 27
- loader 31
- LoaderClass class 32
- loading persistent objects 31
- LocateReply messages 83
- LocateRequest messages 62, 83
- LOCATION_FORWARD reply status 83, 95
- locator, Orbix 3 migrating to Application Server Platform 27
- LocatorClass class 28
- locator daemon
 - administering POA names 53
 - in Application Server Platform 53
- logging
 - and portable interceptors 29
 - and smart proxies 33
- lsit command 54
- lsns command 55

M

- manual work queues 30
- markers, converting to object ID 21
- max_chain_length variable 49
- MaxConsumers administration
 - properties 43
- MaxEventsPerConsumer QoS property 42

- MaxRetries QoS property 42, 45
- MaxRetryTimeout QoS property 42
- MaxSuppliers administration property 43
- mechanism_policy variable 48
- minor codes, for system exceptions 19
- mkdir command 54
- multiple location forward 95
- multi-threaded request processing 29

N

- NamingContextExt interface 39, 85
- naming service
 - and #pragma prefix 67
 - and NamingContextExt interface 85
 - C++ code sample 88
 - extensions 39
 - interoperability 39, 85
 - Java code sample 88
 - load-balancing extensions 39
 - source code compatibility 39
 - stub code 86
- _narrow() function 65
 - and NamingContext 86
 - semantics 66
- narrow characters, and codeset negotiation 97
- negotiate_transmission_codeset variable setting 101
- new_group command 55
- newncns command 55
- node daemon 48, 53
- notification console 44
- notification service
 - administration properties 43
 - any-style clients 41
 - CORBA compliance 40
 - deprecated features 44
 - management 44
 - migrating 40
 - overview 40
 - PacingInterval type, migrating 42
 - Quality-of-Service properties 42
 - starting 44
 - subscribing and publishing, updates 40
 - TimeBase::TimeT, migrating 41
 - unstructured events 41
- NotificationService object ID 40

O

- object-by-value 13
- ObjectGroup interface 28
- object groups, and load balancing 27
- object IDs, converting to marker 21
- OBJECT_NOT_EXIST system exception
 - and new semantics 75
 - Application Server Platform semantics 75
 - launch and invoke rights 61
 - migration 19
 - new semantics 79
 - old semantics 75

- object_to_string() function 65
- ObtainInfoMode enumeration 41
- obtain_offered_types() operation 40
- obtain_subscription_types() operation 40
- OIDTag type 47
- on_demand 56
- opaque type 13
- OpenSSL proprietary private key format 50
- operation signatures
 - Environment parameter 19
- ORB_CTRL_MODEL policy 30
- orbixd daemon 53, 54
 - chmodit utility 61
 - invoke rights 60
 - launch rights 60
- OrbixNotification 3 40
- OrbixSSL 3.x configuration, migrating 48
- OrbixWeb3.cfg configuration file 80
- outgoing_connections:hard_limit 35
- outgoing_connections:soft_limit 35
- out parameters, and C++ function signatures 21

P

- PacingInterval type 42
- PEM format 50
- per_client 56
- pick_member command 55
- piggybacking
 - in filters 29
 - migrating to Application Server Platform 30
- pingit command 54
- PKCS#12 format 50
- PKCS#1 format 50
- PKCS#8 format 50
- plug-ins
 - iiop_tls 47
 - lease 36
- POA
 - and object identities 22
 - creating 23
 - names, administering 53
 - replacing the BOA 22
- POA policies
 - and POA creation 23
 - for callback objects 17
- policies
 - allow_wchar_types_in_1_0 102
 - negotiate_transmission_codeset 101
 - threading policies 30
- policy-based API 45
- portable interceptors
 - and logging 29
 - replacement for filters 29
- principal
 - enabling 61
- Principal type 13, 61
 - interoperability 13
- prioritized request processing 30

- privacy enhanced mail format 50
- private keys
 - migrating 51
 - OpenSSL proprietary format 50
 - PKCS#1 format 50
 - PKCS#8 format 50
 - proprietary KEYENC format 51
- Proprietary KEYENC format 51
- psit command 54
- publication, to notification channel 40
- PullInterval QoS property 43
- putidl command 54
- putit command 54
- putncns command 55
- putnewncns command 55
- putns command 55

Q

- QOFPolicy interface 48
- QoS properties 42
 - MaxEventsPerConsumer 42
 - MaxRetries 42, 45
 - MaxRetryTimeout 42
 - PullInterval 43
 - RequestTimeout 42
 - RetryMultiplier 43
 - RetryTimeout 42
- Quality-of-Service properties 42

R

- readifr command 54
- record() function 32
- rename() function 32
- replace() function 18
- replies, LOCATION_FORWARD status 83
- repository IDs 65
 - and #pragma ID 67
- reputncns command 55
- reputns command 55
- request processing, prioritized 30
- RequestTimeout QoS property 42
- RetryMultiplier QoS property 43
- RetryTimeout QoS property 42
- rmdir command 54
- rmdirl command 54
- rmit command 54
- rmns command 55

S

- save() function 32
- saving persistent objects 31
- security 29
 - and filters 81
 - and transformers 34
 - ASN_OID structure 47
 - AVA interface 46
 - AVAList interface 46
 - CertError interface 46
 - CertValidator interface 47
 - CertValidatorPolicy interface 47
 - configuration variables 48

- enabling 47
- EstablishTrustPolicy interface 48
- Extension interface 47
- ExtensionList interface 47
- IT_AVA interface 46
- IT_AVAList interface 46
- IT_CertError structure 46
- IT_Certificate interface 46
- IT_CRL_List interface 47
- IT_Extension interface 47
- IT_ExtensionList interface 47
- IT_OID structure 47
- IT_OIDTag type 47
- IT_SSL interface 47
- IT_TLS_API interface 46
- IT_TLS interface 46
- IT_UTCTime interface 47
- IT_ValidateX509CertCB 47
- IT_X509CertChain interface 47
- IT_X509Cert interface 47
- IT_X509_CRL_Info interface 47
- IT_X509_Revoked interface 47
- IT_X509_RevokedList interface 47
- OIDTag type 47
- QOFPolicy interface 48
- UTCTime type 47
- X509CertChain interface 47
- X509Cert interface 47
- SecurityLevel1 module 45
- SecurityLevel2 module 45
- Security module 45
- security service 31
- send_locate_request variable 61
- send_principal variable 61
- servant activator 31
- ServantActivator class 32
- servant locator 31
- servant manager 31
- servant objects 22
- server clusters 27
- servergroups command 54
- serverhosts command 54
- service contexts
 - CodeSets 97, 101
 - replacement for piggybacking filters 30
- session_cache_size variable 49
- session_cache_validity_period variable 49
- session_caching_policy variable 49
- session management
 - and I/O callbacks 35
 - client migration 36
 - overview 36
 - server migration 36
- SetConfigValue() function
 - and listen queue size 95
 - using 76
- SINGLE_THREAD_MODEL policy 30
- smart proxies
 - caching 34
 - definition 32
 - migrating to Application Server Platform 33

- socket-level information 31
- startHealthCheck() function 44
- stopHealthCheck() function 44
- string events 45
- subscription
 - to notification channel 40
- system exceptions
 - and
 - IT_USE_ORBIX3_STYLE_SYS_EXC 73
 - changes in semantics 19
 - COMM_FAILURE 80
 - compatibility matrix 77, 81
 - FILTER_SUPPRESS 81
 - INV_OBJREF 75, 79
 - IT_USE_ORBIX3_STYLE_SYS_EXC
 - variable 78
 - minor code differences 19
 - new semantics 73, 78
 - OBJECT_NOT_EXIST 75, 79
 - old semantics 73, 78
 - semantics 73
 - TRANSIENT 80
 - UNKNOWN 81, 82

T

- TAG_CODE_SETS IOR component 97, 101
- target_secure_invocation_policy
 - variable 48
- TCP/IP
 - accessing details 31
 - accessing from application 34
 - and session management 36
- templates, and tie approach 24
- thread filter, migrating to Application Server Platform 30
- tie approach 24
- TIE macro 24
- TimeBase::TimeT
 - and notification service 41
 - replacing PacingInterval type 42
- TLS, policy-based API 45
- transformation of exceptions, and
 - IT_USE_ORBIX3_STYLE_SYS_EXC 74
- transformers 34
- TRANSIENT system exception
 - new semantics 76, 80
 - old semantics 76, 80
 - when raised 19
- transmission codesets 97
- transports, accessing TCP/IP layer 34
- trusted_CA_certificate_list 51
- trusted_ca_list variable 48

U

- UNKNOWN system exception 81, 82
- UnknownUserException user exception
 - class 83
- unstructured events 41
- URL, corbaloc format 16
- user exceptions
 - and DII 82

- parsing with dynamic any 83
- UnknownUserException 83
- UTCTime type 47

W

- wchar type
 - and codeset negotiation 97
 - interoperating 70
 - over GIOP 1.0 connections 102
- wide characters
 - and codeset negotiation 97
- Wonderwall 93
- WorkQueue policy 29
- work queues
 - automatic 30
 - manual 30
- wstring type, interoperating 70

X

- X509CertChain interface 47
- X509Cert interface 47