

Strategies for integrating OSS/J systems with CORBA

Written by: Brenton Camac, Borland Software Corporation
Ke Jin, Borland Software Corporation
Dave Stringer, Borland Software Corporation
Date: January 7, 2004
Filename: COM-ARCH-CORBAV2-1.0.doc
Version: 1.0

Copyright © 2004

BEA Systems Inc., Borland Software Corporation, Digital Fairway Corporation, Eftia OSS Solutions Inc., Highdeal, IBM United Kingdom Limited, ILOG SA, IP VALUE GmbH, Mahindra British Telecom Limited, MetaSolv Software Inc., Motorola, Inc., NEC Corporation, Nokia Networks Oy, Nortel Networks Limited, ObjectStore, Solarmetric Inc., Sonic Software Corporation, Sun Microsystems Inc., Telcordia Technologies Inc., TopCoder Inc., Watchmark Corporation.

All rights reserved. Use is subject to license terms.

1. Executive Summary

OSS/J systems are Operational Support System (OSS) components and sub-systems designed for the telecommunications domain and developed for the J2EE (Java) platform. By definition, an OSS needs to interface with other systems of an operational telecommunications environment in order to obtain information and to effect changes. For many of these existing environments, CORBA technology is used as the means to achieve interoperability between subsystems. Consequently, for such environments, the OSS/J system will require interoperability between its J2EE host platform and the CORBA platform of the operational telecommunications environment.

This document discusses how interoperability between CORBA and J2EE platforms can be achieved for the purposes of OSS/J systems. The possible integration points between the two platforms are several, resulting from the many ways in which systems from both platforms can potentially interact with each another. This document enumerates the primary integration points and presents approaches for achieving interworking between the platforms - utilizing standardized capabilities of the platforms where available.

One of the more complicated integration points is between CORBA's Notification service and J2EE's EJBs. One approach at interoperability is to use a CosNotification-JMS bridge, which is technically complicated. A more recent proposal is to directly subscribe EJBs to the Notification Channel. The details of each are presented and compared.

Finally, to illustrate the programming concepts introduced, functional examples are included. The examples are developed with Borland Enterprise Server J2EE and CORBA technology, but are vendor neutral and should work with any other J2EE-compliant application server and CORBA environment.

Contents

1. Executive Summary	i
2. Introduction.....	1
3. Invoking CORBA Objects from J2EE	2
3.1 <i>Selecting a Naming Service.....</i>	2
3.1.1 Using the CosNaming service.....	2
3.1.2 Using the JNDI service.....	3
3.2 <i>Invoking a CORBA object's operation.....</i>	4
3.3 <i>Example Implementation.....</i>	4
4. Invoking EJB operations from CORBA	7
4.1 <i>Synchronous communications.....</i>	7
4.1.1 Case 1 – A Modifiable CORBA-based System	7
4.1.2 Case 2 – A Sealed CORBA-based System	9
4.2 <i>Asynchronous communications.....</i>	11
4.2.1 Approach 1: Interworking via a JMS Bridge.....	12
4.2.1.1 Translating events into messages.....	12
4.2.1.2 Considerations with this approach	15
4.2.2 Approach 2: Interworking directly with CosNotification service.....	16
4.2.2.1 Structured Event Channel Interworking.....	16
4.2.2.2 Typed Event Channel Interworking	18
4.2.3 Comparison of Approaches	20
5. Appendix 1 - A Simplified IDL Mapping.....	21
6. Appendix 2 - The OMG Notification Service	23
7. References.....	26

2. Introduction

The OSS through Java Initiative is a working group of industry leaders who have joined resources to speed the development of innovative OSS/BSS solutions and to facilitate faster deployment and integration of OSS components. Building on the success of J2EE technology in enterprise applications, the Initiative defines and implements an open, standard set of Java technology-based APIs that help jumpstart the implementation of end-to-end services on next-generation wireless networks and leverage the convergence of telecommunications and Internet-based solutions.

To-date the OSS/J initiative has defined many API specifications for OSS components and systems. Deployments of these OSS/J based systems will typically be into pre-existing OSS environments, where they will converse with existing operational OSS systems in order to obtain information and to effect changes. As many existing OSS environments use CORBA as the basis for such interworking, it is important to consider how J2EE-based OSS/J application components may interwork with CORBA-based systems.

This document identifies the primary points of interoperability between the J2EE and CORBA platforms for OSS/J systems. For each point, one or more approaches for interworking are presented along with their advantages, disadvantages and limitations. Since this information is intended to provide the OSS/J developer with practical information on the subject, several functional code examples are also included.

The goal of this information is to provide the OSS/J developer with specific practical information on J2EE and CORBA interworking in the context of OSS/J systems. The use of these approaches is not required by the OSS/J Initiative. Indeed, the Initiative does not even mandate how OSS/J APIs are to be implemented. This document is provided as a resource for OSS/J developers so they are aware of the issues involved in this area in order to help them implement OSS/J systems which must communicate with pre-existing CORBA-based systems.

This document is arranged into two parts. The first part concerns how J2EE components (primarily EJBs) can invoke CORBA systems. The second part concerns the inverse: how CORBA based systems can propagate information to J2EE components (specifically EJBs).

3. Invoking CORBA Objects from J2EE

This section examines how a J2EE application component may invoke a remote CORBA object's operations. Such a scenario might arise, for instance, when a OSS/J system must interact with an existing CORBA-based OSS system via pre-defined CORBA IDL interfaces.

3.1 Selecting a Naming Service

Before invoking the operations of a CORBA object, a caller must first obtain a reference to the object. The reference is then used to locate and contact the object. In CORBA, objects are located by their IORs (Interoperable Object References). CORBA systems will generally bind IORs of their objects to designated names in a Naming Service. The standardized naming service of the CORBA platform is the CosNaming service. Therefore a standardized and hence portable way to locate a CORBA object is to obtain its object reference from the CosNaming service.

A J2EE component could use the CosNaming service to obtain a reference to a remote CORBA object. Since the CosNaming service is required to be supported by J2EE compliant application servers¹, this would be a portable approach. And even though the naming service instance supplied by the J2EE server may not be the same instance used by the CORBA system to bind its object reference, the J2EE component will be able to use the same CosNaming interfaces to access that other instance. Locating object references using this approach is examined in more detail in Section 3.1.1 below.

Another approach available to J2EE components is to use JNDI (the native naming service of the J2EE platform) instead of CosNaming. The advantage of this approach for J2EE components is that they will be using the same type of naming service and therefore same procedure to retrieve object references irrespective of whether those objects are hosted on the J2EE platform or CORBA platform. Interacting with the CosNaming service via JNDI is possible because the JNDI specification defines Service Provider Interfaces (SPIs) to allow JNDI to interwork with other popular naming and directory services (including CosNaming). The J2EE specification requires access to the CosNaming service be available via JNDI (SPIs)². Locating object references via this approach is examined in more detail in Section 3.1.2 below.

Regardless of how each J2EE server vendor chooses to resolve these requirements in their product offering, the J2EE component can expect there to be two naming service interfaces - JNDI and CosNaming. And since the CORBA object is likely to be registered in CosNaming (since it is a CORBA-based system), then the EJB will have the choice of using either naming service interface. Thus, both approaches are discussed here and shown in the accompanying source code example.

3.1.1 Using the CosNaming service

Traversing the CosNaming service's namespace to locate IORs begins with obtaining a reference to the service's root name-context. The portable way to achieve this is using the two steps as follows:

1. Obtain a reference to an ORB in the J2EE environment. The portable way to achieve this is through the component's standard environment.

¹ "All J2EE products must provide a COSNaming name service to meet the EJB interoperability requirements." (J2EE 1.3 Specification. Section 6.2.4.7).

² "..., a COSNaming JNDI service provider must be available through the web, EJB, and application client containers. A COSNaming JNDI service provider is a part of the J2SE 1.3 SDK and JRE from Sun, but is not a required component of the J2SE specification." (J2EE 1.3 Specification. Section 6.2.4.7).

```
// File: ServiceLocator from accompanying code / example-1
javax.naming.Context jndiRootCtx = new InitialContext();
org.omg.CORBA.ORB orb = (ORB)jndiRootCtx.lookup("java:comp/ORB");
```

Figure 1. Retrieving a handle to an ORB within the EJB Container

2. Query the ORB for its registered Naming Service. CORBA defines a portable way to retrieve references to its various object services (e.g. Naming Service, Event Service, Transaction Service, etc) with the ORB's `resolve_initial_references()` method. How an ORB is configured with such initial references is vendor specific and a server administration task.

```
// File: ServiceLocator from accompanying code / example-1
org.omg.CORBA.Object ns = orb.resolve_initial_references("NameService");
org.omg.CosNaming.NamingContext cosNamingRootCtx =
    NamingContextHelper.narrow(ns);
```

Figure 2. Obtaining an object reference to root context of the CosNaming service

Having obtained the root naming context of the desired name service, the component can now navigate the name graph to locate the desired object reference.

In the CosNaming service, the name of the binding is represented as a structured data type - the NameComponent. Converting between this structured form and a string form is facilitated with convenience methods. This example uses a convenience method `to_name()` to achieve this. The componentized name is then supplied to the resolve operation of a Naming Context object - typically the root context.

```
// File: ServiceLocator from accompanying code / example-1
org.omg.CosNaming.NameComponent[] compoundName =
    com.inprise.vbroker.naming.NamingUtil.to_name(nameString);
java.lang.Object objRef = cosNamingRootCtx.resolve(compoundName);
```

Figure 3. Resolving a name binding in CosNaming

This returns the object reference bound to that name in the naming service.

3.1.2 Using the JNDI service

If the component is to use the JNDI naming service to lookup the CORBA object reference, then the root naming context of that service must first be obtained. The portable way to achieve this in the J2EE environment is as follows:

```
// File: ServiceLocator from accompanying code / example-1
javax.naming.Context jndiRootCtx = new InitialContext();
```

Figure 4. Obtaining an object reference to the root context of JNDI service

Having obtained the root naming context to the JNDI naming service, the component can now navigate the name graph to locate the desired object reference. With the JNDI naming service the name is resolved against the naming context using the lookup method with a String argument.

```
// File: ServiceLocator from accompanying code / example-1
java.lang.Object objRef = jndiRootCtx.lookup(nameString);
```

Figure 5. Resolving a name binding in CosNaming

This returns the object reference bound to that name in the naming service.

3.2 Invoking a CORBA object's operation

Having obtained a reference to the remote CORBA object (represented by the variable objRef in the code), it must then be narrowed to the appropriate type. Narrowing the data type does not depend on which method was used to obtain the object reference. Narrowing the object reference is performed by the narrow() function of the Helper class generated by the IDL to Java compiler.

```
// File: EJB1Bean from accompanying code / example-1
// narrow object for use
ServiceAInterface1 corbaObjRef = ServiceAInterface1Helper.narrow(
    (org.omg.CORBA.Object) objRef);
// invoke operation of CORBA object
corbaObjRef.echo("test data");
```

Figure 6. Invoking operations of a CORBA object

Once narrowed, the object reference can be used as any other local object, and its methods invoked. This is shown in the code fragment of Figure 6 which is based on the testJNDI() and testCosNaming() methods of EJB1 from example 1. The echo() method is an operation implemented by the remote CORBA object.

3.3 Example Implementation

The preceding concepts are combined and illustrated in example1 project [1]. This example demonstrates that a CORBA object can be registered in CORBA's CosNaming service, and subsequently located by an EJB either via the JNDI or CosNaming services. Once located, the EJB is able to invoke operations of the CORBA object. The supplied test case exercises the EJB to use both JNDI and CosNaming services to retrieve the CORBA object reference.

The example consists of a CORBA server (ServiceAInterface1) which offers an operation 'echo' that returns the string data it is supplied. When running the supplied test case, the following steps will be performed.

Step 1. When the CORBA service is started it first locates the CosNaming service. There are several mechanisms for discovering a CosNaming service, both standard and proprietary. Typical of the standard mechanisms is to supply the CORBA server with a URL indicating the CosNaming service's location on

the network: e.g. corbaloc::localhost:683/NamingService. Once the service has been located, the CORBA server registers a CORBA object reference with it.

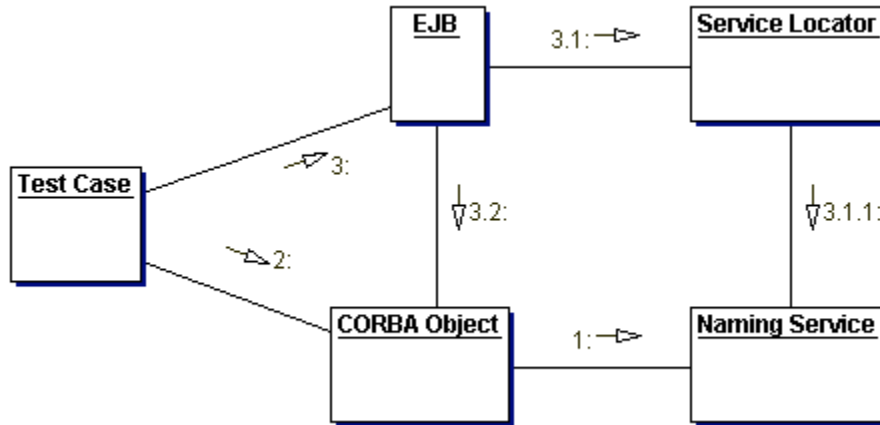


Figure 7 Steps taken when exercising Example 1's test

Step 2. The test 'testCORBASvcExists' is run by the JUnit conformant test case (TestCase1). This is a pre-condition test for the other two which follow (see Step 3) and ensures that the CORBA object is available for them.

Step 3. The tests 'testJNDI' and 'testCORBA' both follow the same execution path which is: first the EJB's testJNDI() or testCORBA() method is invoked depending upon the test case.

Step 3.1. The EJB contacts a utility class (ServiceLocator) to lookup the object reference either via JNDI or CosNaming as directed.

Step 3.1.1. The Service Locator contacts the naming service (either via JNDI or CosNaming protocols) to retrieve the requested object binding.

Step 3.2. The EJB narrows the returned object reference and then invokes the echo operation of the CORBA object. If the data returned from the echo operation is the same as what was supplied, then the EJB's method returns true, otherwise false.

Summary

This section examined how a J2EE component could locate and invoke the operations of a CORBA object. Since the J2EE component could use either naming service of the J2EE platform (JNDI or CosNaming), both approaches were examined. Although each approach used different techniques, both achieved the same result - obtaining an object reference to a remote corba object. To illustrate these concepts, fragments of source code taken from complete examples (see[1]) were included.

The choice of which approach to use will depend on various factors. For a pure-EJB environment, the JNDI approach might be preferred since it uses techniques that are already commonly used in that programming environment to locate other objects (e.g. other EJBs, etc). However, if existing client-side CORBA code is available and is to be reused in the EJB, then the CosNaming approach may be favored since this would require less change to that code. Another consideration which might influence the choice of approach is which Naming Service instance hosts the sought after object reference. If an existing CORBA application is using its own Naming Service to store object references, and if this naming service is not available to EJBs via the JNDI service of the EJB container (ie. the container is not configurable to

work with foreign naming services), then it may be necessary for the EJB to use the CosNaming approach and explicitly connect to the foreign Naming Service.

The ability for J2EE components to use standard means to interoperate with CORBA objects and services this way is possible because the J2EE platform explicitly requires support for CORBA's CosNaming service and RMI-IIOP protocol in J2EE-compliant application servers.

4. Invoking EJB operations from CORBA

This section examines how an EJB's operations may be invoked from a CORBA-based system. Two modes of interaction are considered. The first mode is direct (synchronous) communications where a CORBA-based system directly invokes an operation of an EJB's remote interface. The second mode is indirect (or more precisely - asynchronous) communications where a CORBA-based system propagates events to the EJB container via a messaging system. For both modes there is more than one approach that can be taken; these are presented and compared.

4.1 Synchronous communications

This section examines various approaches whereby a CORBA-based system can directly invoke the methods of an EJB. In the context of OSS/J systems, this arrangement may arise, for instance, under the following situations:

1. An existing CORBA-based system is to be extended to interact with an OSS/J system's EJBs. In this case the interface being used is defined by the EJB and therefore involves RMI semantics. Since many existing CORBA systems are implemented in C++ as well as Java, both Java and non-Java scenarios will be examined.
2. An OSS/J system must register a callback interface with an existing CORBA-based system. In this case the callback interface is defined by the CORBA-based system in IDL.

4.1.1 Case 1 – A Modifiable CORBA-based System

This case examines how a CORBA-based system can be extended to interact with an EJB. Such might occur, for instance, when an existing CORBA-based system is modified to invoke the operations of an OSS/J system. Since the interaction between the two systems is according to the EJB's interface, that interface must be mapped to the local implementation environment of the CORBA-based system.

RMI Clients

If the CORBA system is implemented using Java, then it is sufficient to create RMI-IIOP stubs from the EJB interface. This can be done using tools such as Java2 SDK's "rmic" (with `-iiop` option) or Borland's Visibroker "java2iiop" utility. An illustration of this is shown as Example-2 in the accompanying source code, which demonstrates a CORBA client invoking the methods of an EJB's Remote interface. The CORBA system can then use the generated stubs to invoke operations of the EJB.

Non-rmi clients

If the CORBA system uses C++, then the interaction becomes more complex as the EJB's (Java) RMI semantics and its native language objects (e.g. Collections, etc) will need to be expressed in C++. One approach to achieve this is to map the EJB interface first to language-neutral IDL and then from IDL to C++. Both mappings are defined by the OMG and tools exist to perform this automatically (e.g. Borland's Visibroker product has both `java2idl` and `idl2cpp` compilers).

An illustration of this is found in Example-2. The remote interface of the EJB (EJB2) is shown below.

```
// File: EJB2 from accompanying code / example-2
package ossj.example2.myEJBapp;

public interface EJB2 extends javax.ejb.EJBObject {

    public String echoString(String s) throws RemoteException;

    public String[] echoStrings(String[] ss) throws RemoteException;

    public Date echoDate(Date d) throws RemoteException;
}

```

Figure 8. Remote interface of EJB2

Using the JavaIDL mapping ([2]) to translate this Java interface to IDL produces the following IDL definition. (Some elements have been removed for brevity).

```
...
module org {
    module omg {
        module boxedRMI {
            module java {
                module lang {
                    valuetype seq1_String sequence< ::CORBA::WStringValue >;
                    ...
                };
            };
        };
    };
};

module java {
    module lang {
        custom valuetype _Exception : ::java::lang::Throwable {
        };
        ...
    };
    ...
    module util {
        custom valuetype Date : ::java::lang::Comparable
        supports ::java::lang::Cloneable {
        };
        ...
    };
};

module ossj {
    module example2 {
        module myEJBapp {
            interface EJB2 : ::javax::ejb::EJBObject {
                ::CORBA::WStringValue echoString (in ::CORBA::WStringValue arg0);
                ::org::omg::boxedRMI::java::lang::seq1_String echoStrings (in
                    ::org::omg::boxedRMI::java::lang::seq1_String arg0);
                ::java::util::Date echoDate (in ::java::util::Date arg0);
            };
        };
    };
};

```

Figure 9. Some of the IDL derived from EJB2's Remote Interface

The non-Java client of this interface would need to provide implementations for the custom valuetypes `Date` and `_Exception` in this particular instance. Other valuetypes used in the interface such as `seq1_String` are defined in terms of existing IDL constructs (e.g. `sequence< ::CORBA::WStringValue >`).

Consequently, this approach of mapping an EJB interface to IDL and then to an implementation language typically requires the client provide implementations for some of the ValueTypes. In addition, it also results in complex interfaces even for EJBs with relatively simple interfaces. The complexity arises because the OMG defined mappings attempt to accommodate as many features of the EJB interface as possible.

For a discussion of a non-standardized approach to the problem of mapping EJB interfaces to IDL see Appendix 1 - A Simplified IDL Mapping.

Summary

For CORBA systems implemented in Java, integration with EJBs is achieved using commonly available RMI-IIOP compilers such as `rmic` and `java2iiop`. For CORBA systems implemented in C++, mapping the EJB interface to the target language requires more effort. A 'brute-force' approach of transforming the Remote interface to IDL and then to C++ is possible, but results in complex C++ constructs which must be implemented by the client and manipulated by the client at runtime. Typically, such mapped interfaces are quite complex as they attempt to accommodate all of the semantics of the original EJB (RMI) interface. If the EJB's interface is relatively simple, then such a comprehensive mapping may not be required. In such cases, a 'simplified' Java to IDL mapping might be considered. Such a scheme is non-standard and is offered by at least one vendor (See Appendix 1 - A Simplified IDL mapping).

4.1.2 Case 2 – A Sealed CORBA-based System

This section examines the case where an existing interface (defined in IDL) is to be implemented by an EJB. Such might occur, for instance, if an OSS/J system is to receive callbacks from existing CORBA-based system. In this case, the callback interface is defined by the CORBA-based system (in IDL) and the EJB must implement it in order to be contacted by the CORBA-based system.

The task then is for the EJB to provide an implementation of the IDL interface. The IDL can be mapped to Java using standard language mapping [4]. However, this can not be implemented directly as an EJB Remote interface is given RMI semantics [5]. Since the caller using the original IDL would not be using RMI semantics the invocations would fail.

This issue can be resolved by using a proxy object to add necessary RMI semantics. The proxy's implementation is based on the IDL to Java language mapping [4]. Invocations it receives via that interface are propagated to the target EJB via stubs compiled from the EJB's interface.

An illustration of this is shown in Example-3. There, an IDL interface (with name `Callback`) is used by a CORBA-based service to notify an interested party of events. The CORBA system has another interface (`CORBASvc`) which exposes two operations: `registerCallback()` which allows an object implementing the `Callback` interface to be registered in order to receive future callbacks; and `eventNotification()` which contacts the registered callback object via its `eventNotification()` method for testing purposes. These are shown in Figure 10.

```

module MyCORBASystem3 {

    interface Callback {
        long eventNotification(in long param1);
    };

    interface CORBAsvc {
        void registerCallback(in Callback c);
        long eventNotification(in long param1);
    };
};

```

Figure 10. IDL Interface of CORBAsvc

Once the configuration has been initialized, invocations on the CORBAsvc's eventNotification() method will propagate to the registered Callback object which in turn propagates to the EJB. The EJB3's implementation of echoInt() is to return the value supplied it. In this way, the test driver (caller) can confirm that the call to the CORBAsvc has successfully completed its round-trip. This is depicted below.

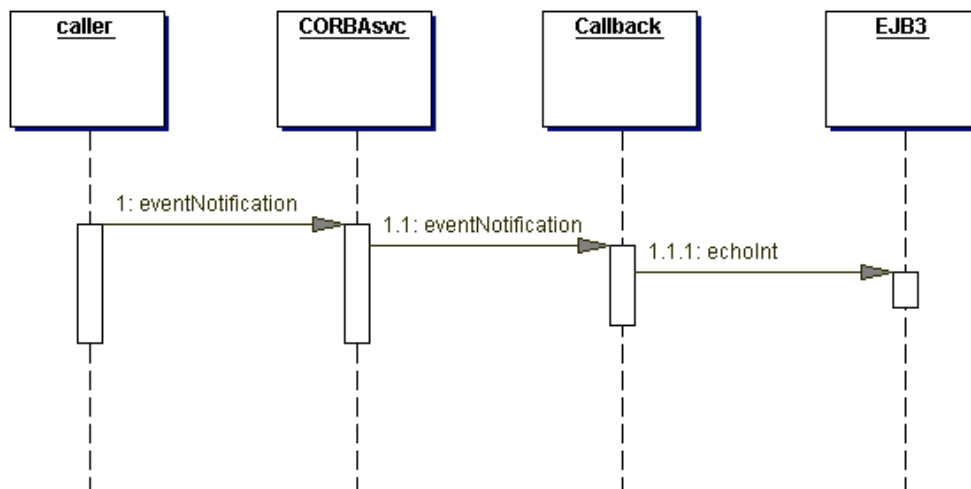


Figure 11. Testing the configuration of Example-3

To produce this configuration, the following steps are taken in the example. (For convenience, the proxy object is collocated with the EJB in the EJB container). When the EJB (EJB3) is initialized, it creates a POA to host the callback object. Then when EJB3's registerForCallback() method is invoked, it creates a local instance of the Callback object, configures it to delegate its calls to this EJB3 instance, registers the object with the POA so it can be contacted by remote CORBA clients, and then contacts the CORBAsvc to register it as a Callback object. The methods used are shown in the sequence diagram below. This sequence of activities is initiated by the test client TestCase3 during its setup phase.

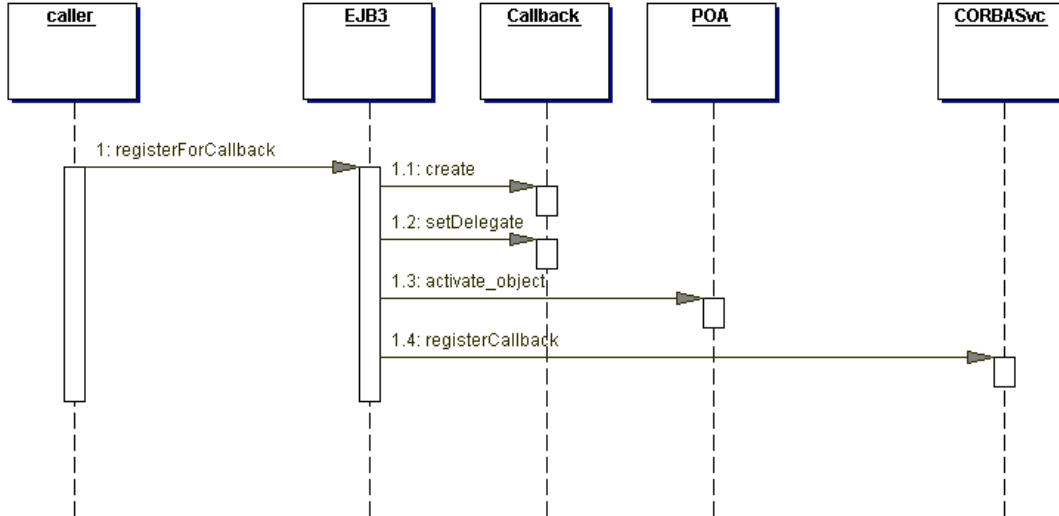


Figure 12. Configuring EJB3 to receive callbacks from CORBASvc

Summary

In summary, EJBs can be integrated with existing CORBA services through pre-existing IDL interfaces. EJBs require callers to use RMI semantics which typically won't be satisfied by existing pure-CORBA systems. One approach to resolve this is to use a proxy that implements an IDL interface and delegates its invocations to a configured EJB using the required RMI semantics.

4.2 Asynchronous communications

Many OSS/J services require interaction with existing underlying operational systems. For some OSS/J systems, the interaction will be the receiving of events produced within those existing systems. For instance, the OSS/J Quality of Service (QoS) system is interested in alarm events and threshold triggers occurring in underlying OSS systems in order to fulfill its service to its clients.

The majority of these existing operational systems use CORBA's CosNotification service³ to propagate events. (The CosNotification service is CORBA's service that provides asynchronous event propagation and is explained in more detail in Appendix 1.) CosNotification is used widely in this domain with many telecom standards (e.g. recommendations ITU-T X.780 and TMF 814A) defining how telecom-domain specific events are to be expressed as CosNotification events. Since there are no equivalent standards, recommendations or tools devised for alternative messaging technologies (e.g. JMS) it is reasonable to assume that OSS will continue to use the CosNotification service and that OSS/J systems will therefore need to integrate with CosNotification in order to interact with underlying OSS systems.

The problem confronting OSS/J systems then is how their EJB components may receive events from a CosNotification Channel. Since the J2EE platform possess its own asynchronous messaging service – JMS – one approach is to map CosNotification events into JMS messages and then deliver these to an EJB via a Message Driven Bean (MDB). A proposal for this approach is discussed below. Another approach is for an EJB to be directly subscribed to the notification channel as an event consumer. This also is discussed below as well as a comparison of the relative merits and limitations of each approach.

³ The use of the CosNotification service in the telecom domain should not come as a surprise considering it was the telecom membership of the OMG who were the main force behind the proposal and definition of an event service for CORBA.

4.2.1 Approach 1: Interworking via a JMS Bridge

This approach considers a bridge that links the CosNotification service with JMS and which translates events from one domain to the other. The bridge thus acts as an event consumer to the Notification service and a message supplier to JMS. Such an interworking service is currently being discussed within the OMG as part of a general purpose JMS / CosNotification Interworking Service [6].

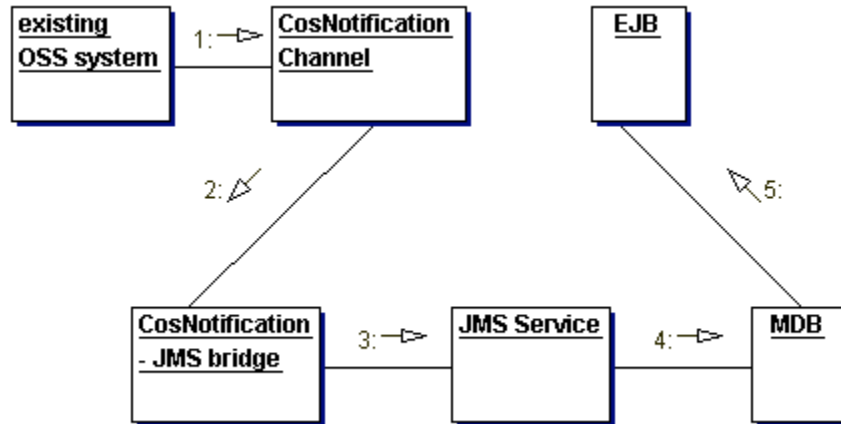


Figure 13. Interworking via the CosNotification-JMS Bridge

4.2.1.1 Translating events into messages

According to the proposal [6], information in a CosNotification StructuredEvent would be mapped into a JMS message using the following general formula.

1. If the event's variable_header standard optional part is present (i.e. Event_Reliability, Timeout and Priority fields), then these are mapped to corresponding fields of the JMS message header. If the event supplier does not define these fields, then the JMS header fields are populated using default values as specified in the JMS specification.
2. The event's fixed_header, the rest of the optional header fields as well as its filterable body are placed in JMS properties fields. These will therefore be seen by the JMS receiver application as user-defined fields.
3. Fields in the JMS message which have no counterpart in the event domain (such as JMSDestination and JMSMessageID) are populated by the translator.
4. The event's remainder_of_body section is mapped to a JMS message body type depending on the type and complexity of the data.

Conceptually, this translation is depicted as follows.

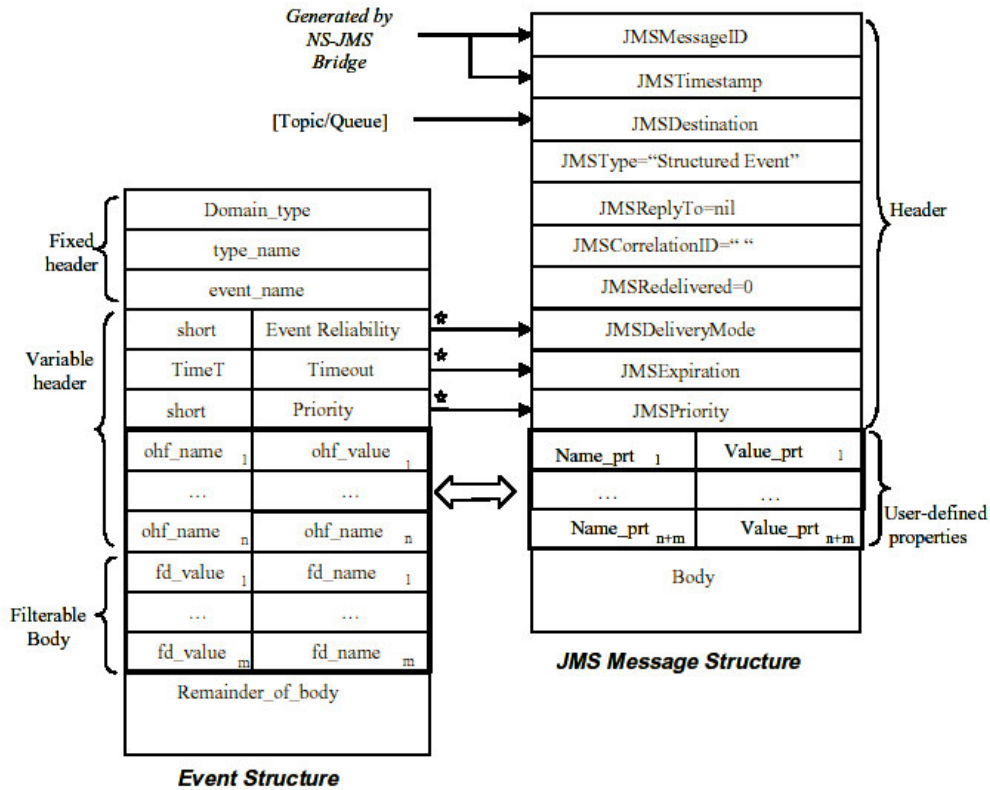


Figure 14. Structured Event to JMS Message mapping

This translation formula is described in more detail below.

Non-Structured Event fields

Values for the JMS message fields JMSMessageID, JMSTimestamp, JMSDestination, and JMSType are created by the JMS-NS bridge. Specifically -

- JMSMessageID is a String value which should be a unique key, prefixed by 'ID'. The exact scope of uniqueness is not defined.
- JMSTimestamp field contains the time the message was submitted to JMS for sending. It is in the format of a normal Java millisecond time value.
- JMSDestination contains the Topic or Queue name to which the message is being sent.
- JMSType is a String value that should be set to 'Structured Event'.

Structured Event Fixed Header mapping

The structured event's fixed header fields are mapped to user-defined property fields of the JMS message. (The name for these properties must obey the rules for a message selector identifier⁴ specified in Section 3.8.1.1 of the JMS Specification [7]). Specifically -

⁴ An identifier is an unlimited-length character sequence that must begin with a Java identifier start character; all following characters must be Java identifier part characters. An identifier start character is any character for which the method Character.isJavaIdentifierStart returns true. This includes '_' and '\$'.

- The domain_type event header field is mapped to a new JMS user-defined property whose name is “domain_type” and whose value is the event's domain_type field converted to java String.
- The type_name event header field is mapped to a new JMS user-defined property whose name is "type_name" and whose value is the event's type_name field converted to java String.
- The event_name event header field is mapped to a new JMS user-defined property whose name is "sevent_name" and whose value is the event's event_name field converted to java String.

Structured Event Variable Header mapping

If the structured event's variable header contains the fields EventReliability, Priority, or Timeout, then these are mapped to JMSDelivery, JMSpriority, and JMSTimetolive respectively. If these fields are not defined in the event, then the following values are used: JMSDelivery is set to PERSISTENT; JMSpriority is set to 4; and JMSTimetolive is set to Unlimited.

Structured Event Optional Header fields and Filterable Body mapping

Optional header fields (ohf_*) and filterable data fields (fd_*) of the Event are mapped to user-defined JMS properties with names \$ohf_* and \$fd_* respectively. These JMS property names must obey the rules for a message selector identifier specified in Section 3.8.1.1 of JMS specification. The content of the ohf_* or fd_* field is converted to a Java primitive data type.

If the optional header field or filterable data field is a complex IDL type then multiple JMS properties are used, one for each primitive element of the complex type. The complex structure is decomposed into the individual JMS properties as follows:

- The new \$ohf_* or \$fd_* field name is the concatenation of the event field name, the structure name(s) and the member of the structure name(s). The structure member operator '.' is used to delimitate each name component in the concatenation. When a sequence structure is encountered, then the ordinal value of the particular element is used as its name (e.g. 1, 2, 3, etc). This process is repeated if the structure contains nested data structures as elements, along with primitive elements.
- The content of the decomposed field is converted to a corresponding Java primitive data type.

For example, a structured event containing a field with name <Fd1> and value of type <CORBA::Any> which encapsulates an IDL struct named Alarm { string AlarmName; int Severity }, is transformed into two JMS user-defined properties : <\$Fd1.AlarmName, String> and <\$Fd1.Severity, Integer>.

Another example is a structured event containing an optional header field with name <Ohf1> and value of type <CosNaming::Name> (defined a sequence of CosNaming::NameComponent) whose content is NC1/NC2/NC3. This is transformed into the JMS user-defined properties: <\$Ohf1.Name.1.id, String>, <\$Ohf1.Name.1.kind, String>, <\$Ohr1.Name.2.id, String>, <\$Ohr1.Name.2.kind, String>, <\$Ohr1.Name.3.id, String>, <\$Ohr1.Name.3.kind, String>.

If the optional header or filterable date field has other complex data types, this is mapped to a byte stream. The JMS client will then need to use the appropriate CORBA Helper classes to unmarshal this user data.

Structured Event Remainder of Body mapping

The Structured Event's remainder_of_body is of type CORBA::Any. Its contents are mapped to one of the five JMS message body types (Text, Stream, Map, Bytes, or Object) depending on the complexity of the field's encapsulated data type. When the Event's remainder_of_body Any type involves:

- IDL basic type elements, each element is mapped to a Java primitive type using standard IDL to Java mapping. The resulting set are packaged in a JMS StreamMessage body.

- String type element only, is mapped to a java String type and packaged in a JMS TextMessage body.
- Sequence of Properties (PropertySeq), is mapped to a set of name-value pairs where names are Strings concatenated with the ordinal value of the element in the sequence, and values are Java primitives. This is packaged in a JMS MapMessage body.
- Octet Sequence element or other type such as user constructed types, are mapped to a JMS BytesMessage body. In order to reconstruct the IDL type, the JMS client that would receive a BytesMessage and would use the appropriate CORBA Helper class.

4.2.1.2 Considerations with this approach

An important characteristic of this approach is the decomposition or ‘flattening’ of the structured event’s non-primitive data type fields into a series of JMS properties in the counterpart JMS message. The need for this arises because the Java Messaging Service allows only primitive data types to appear in user-defined properties. However, the use of non-primitive data types in CORBA and in the telecom domain is common. Examples include the IDL structure TimeBase::UtcT (see ITU-T X.780) and the IDL sequence CosNaming::Name (see TMF 814A). Therefore, applying this approach to the telecom domain will result in message consumers frequently needing to deal with ‘flattened’ structured information.

Flattening this information into JMS message properties allows the JMS message to be filtered by the JMS service, in a similar way to the original structured event being filterable in the CosNotification domain. However, there are many disadvantages to flattening the structured event information, including as the following.

- a) Substantial programming effort is required to re-assemble the event information in the consumer. As this information it is no longer represented explicitly by a structured datatype in the JMS message but implicitly in the names of JMS properties, JMS message consumers will need to formulate property names (in dotted notation) in order to navigate and retrieve event information. Such an approach is less type-safe than using first-class data structures of a language (e.g. those produced by an IDL compiler). With no automated, compile-time checking to ensure the property names are formed correctly, there exists the potential for programming errors which will not become evident until runtime.
- b) Algorithms and designs from existing CORBA-based event consumers can not be reused in the JMS consumer unless the event is first reformatted into its original structure. Using existing algorithms and designs to interpret events can significantly decrease development costs and improve code quality since such artifacts have already been developed and used and therefore more likely to have defects identified and corrected.
- c) JMS message consumers will likely need to access CORBA Helper classes as well, even though the event’s filterable data has been decomposed by the translation process. This is because user-defined structures in the event’s remainder_of_body field are not decomposed during the translation process as filterable data is. Hence, the JMS message consumer will typically incorporate both CORBA Helper classes to access non-filterable event data as well as custom code to access filterable event data; i.e. it will contain code to work with two technology domains – JMS and CosNotification. Such requirements makes the resulting consumers complex and causes much of its implementation to be duplicated.

A final consideration with this approach is that it only relates to Structured events not Typed events (events from the Typed Event Notification Channel). As discussed later in section 4.2.2.2, Typed Channels are much more efficient and effective at event propagation which is why some telecom domain standards (e.g. ITU-T/T1M1) recommend their use. To employ this approach with Typed events would require they first

be converted into Structured Events before being translated to JMS. Such an arrangement adds yet another level of translation, effort and subsequent administration effort.

4.2.2 Approach 2: Interworking directly with CosNotification service

Another approach whereby EJBs can receive events from a CosNotification channel is to register EJBs with the channel as Event Consumers. In this approach, CosNotification events are delivered directly to EJBs.



Figure 15. Interworking with the CosNotification Channel

The following sections discuss how this is achieved for the Structured Event and Typed Event Channels.

4.2.2.1 Structured Event Channel Interworking

For EJBs to receive events from a Structured Event Channel⁵ they implement the CosNotification interface for a Structured Event Consumer and subscribe to the Notification Channel as an event consumer. In this way the Notification channel treats the EJB as an event consumer and delivers events to it directly.

Interface Requirements

There are two kinds of Structured Event Consumer defined by the Notification service; a PullConsumer and a PushConsumer. A PullConsumer is one which polls the Notification Channel for events; a PushConsumer is one which is sent events from the Notification Channel. The EJB as an event consumer should implement the PushConsumer interface⁶ in keeping with the intention of the EJB specification which desires that EJBs be passive targets of operations rather than self-initiated actors of activities. Being a Push Consumer requires the EJB to declare and implement the `push_structured_event()` operation through which events will be received from the channel. The IDL for this interface is shown below.

⁵ The Structured Event Channel is explained in more detail in Appendix 2.

⁶ Although an EJB could choose the PullConsumer interface instead, this would require it to be proactive in periodically polling the channel for events. Such autonomous behavior by EJBs is discouraged by the J2EE specification.

```
// File: CosNotifyComm.idl from the Notification Service spec.

module CosNotifyComm {

    interface StructuredPushConsumer : NotifyPublish {

        void push_structured_event ( in CosNotification::StructuredEvent
            notification ) raises ( CosEventComm::Disconnected );

        void disconnect_structured_push_consumer();

    };

    ...
};
```

Figure 16. IDL definition of StructuredPushConsumer

This IDL interface and associated datatype (StructuredEvent) can be mapped to Java via the IDL to Java language mapping [4]. Note, however, that the StructuredPushConsumerOperations interface produced by the mapping is not an RMI interface (as its declared methods do not declare the `java.rmi.RemoteException` exception in their throws clause), and therefore not suitable for being extended by the EJB's Remote interface (although could be extended by its bean interface). Hence the EJB must explicitly declare the `push_structured_event()` operation itself in its Remote interface. This is shown below.

```
// EJB StructConsumer Remote Interface

public interface StructConsumer extends javax.ejb.EJBObject {

    public void push_structured_event(
        org.omg.CosNotification.StructuredEvent event )
        throws java.rmi.RemoteException, org.omg.CosEventComm.Disconnected;

    public void disconnect_structured_push_consumer()
        throws java.rmi.RemoteException
}

}
```

Figure 17. Remote Interface of EJB Structured Event Consumer

Note also that the `push_structured_event()` operation declares an argument of type `org.omg.CosNotification.StructuredEvent`. The definition of this datatype is produced by the IDL to Java mapping. Since it is used within a method of an RMI interface, it will inherit RMI semantics; one of which is to allow the datatype to be null. This concept is represented in IDL by 'boxing' the datastructure in a valuetype, and expressed in IIOP by passing the datastructure using the Object-By-Value mechanism.

```
module org {
    module omg {
        module boxedIDL {
            module org {
                module omg {
                    module CosNotification {
                        valuetype StructuredEvent ::CosNotification::StructuredEvent;
                    };
                };
            };
        };
    };
};
```

Figure 18. BoxedIDL ValueType for org.omg.CosNotification.StructuredEvent

This fact is important because it means that the EJB container is expecting callers of this method to supply their IIOP argument as type `::org::omg::boxedIDL::org::omg::CosNotification::StructuredEvent`. If the caller derives its stubs from the EJB's Remote interface, then its invocations will conform to this requirement since they will be using RMI-IIOP. However, if the caller is using the interface found in `CosNotification.idl` for event suppliers, as would be typically be the case for pre-existing CORBA-based systems, then such requirements will not be met since that interface does not declare that the `StructuredEvent` is to be boxed in a `ValueType`. This is a special case of the same general problem encountered earlier in the paper when synchronous calls from CORBA-based systems to EJBs were discussed (see 4.1.1) where a solution was to introduce a proxy to provide the required RMI semantics. Fortunately, in this instance, the operation and datatype to be translated are known *a priori* to the Notification Channel. This means that the Notification Channel itself could box the argument if the consumer is determined to be using RMI-IIOP, and deliver the same event unmodified (un-boxed) to consumers which have directly implemented the `CosNotification IDL`. This is the approach taken by Borland's `VisiNotify` implementation of the CORBA Notification Channel. The alternative of providing a IIOP/ RMI-IIOP proxy between the EJB and Notification Channel, based on what was discussed in earlier sections, is also an option.

An illustration of the approach whereby the Notification Channel acts as the proxy to supply RMI semantics is shown in Example-4 of the accompanying code. It consists of

- EJB4 - an EJB which can be subscribed to the Notification Channel as an Event Consumer,
- StructPushConsumer - a CORBA-based Event Consumer which can also be subscribed to the channel, and
- StructPushSupplier - an event supplier which produces events based on simplified definitions of selected ITU-T TMN notifications (from the ITU-T X.700 recommendation series).

This example shows that ITU-T TMN notifications can be propagated via a `CosNotification` and directly received by both a CORBA-based consumer as well as an EJB-based consumer without the need for custom proxies.

Summary

Subscribing EJBs directly to the Structured Event Channel as event consumers is possible if the EJB implements the necessary operations. The RMI semantics of the EJB's Remote interface require that the event supplier push events which are boxed. Since the event consumer's interface is known *a priori* to the Notification channel, the channel can perform this value-boxing of those events that are to be delivered to RMI-based event consumers. This capability is demonstrated through example-4 using Borland's implementation of the `CosNotification` service - `VisiNotify`.

4.2.2.2 Typed Event Channel Interworking

The `CosNotification` service defines four different kinds of channel: the `Untyped-Event` channel, the `Structured-Event` channel, the `Sequenced-Event` channel, and the `Typed-Event` channel. The `Typed-Event` channel differs from the other kinds of channels in that it does not use a data-structure defined by the `CosNotification` service to represent events, nor does it use predefined IDL operations such as `push()` or `push_structured_event()` to transfer events between client and channel. Instead, suppliers propagate events to consumers via `Typed` channels using application-defined operations and datastructures. The Notification specification refers to these interfaces used by typed suppliers and consumers as `<I>` interfaces.

There are many advantages to using `Typed Event` channels over the other forms of event (notification) channels. Arguably the most significant advantage is that applications will be communicating via strongly typed application-level interfaces. In fact, the events being exchanged can be declared directly in IDL as data structures along with the operations used to convey them. In contrast, a `Structured`, `Sequenced` or `Untyped` channel uses a generic event type with predefined operations. This flexibility allows compile-time checking of typed suppliers and consumers to ensure the events they will be exchanging will be compatible,

something which isn't possible with the other types of channel other than asserting that clients correctly exchange a StructuredEvent datatype for instance. Support for compile-time checking results in being able to detect errors before operational use and doesn't rely on testing to expose such errors.

Another benefit of using Typed-Event channels, and one related to the above, is that clients require less code to parse event structures. Parsing data types that have statically defined structures is simpler than parsing structures which use meta-level descriptions such as property lists that contain name / value pairs. Consequently, parsing in Typed Event Channel clients requires less code and therefore fewer code defects are likely.

Since clients of typed event channels spend less time parsing event structures at runtime and typed events comprise smaller messages, there is the potential for greater throughput, scalability and overall efficiency. Examples of this compared to the other forms of event channels are shown below. These tests used Borland's implementation of the CosNotification service – VisiNotify version 5.2 executing on an 700MHz PIII running Windows 2000, with one local supplier and four local consumers.

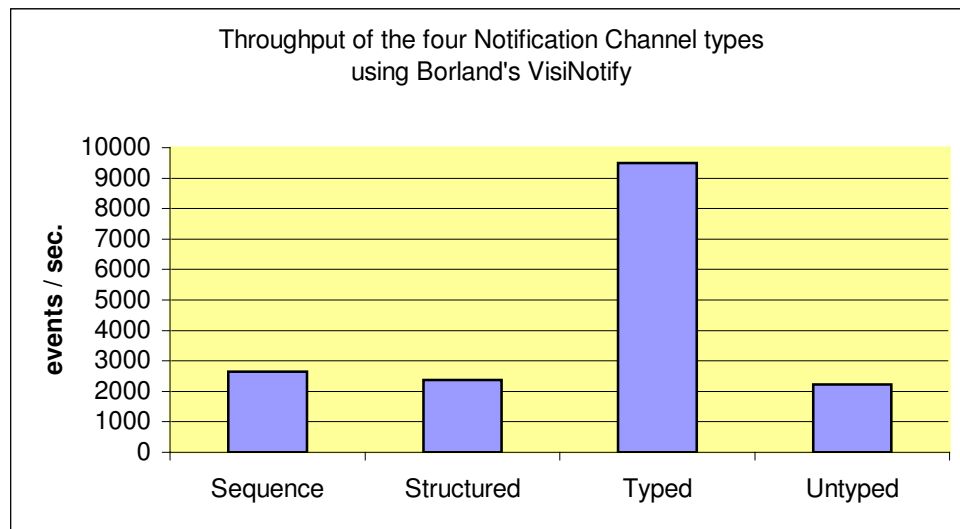


Figure 19. Comparing throughput of different Notification Channel types

For these reasons, ITU-T/T1M1 latterly recommends using typed events.

One restriction with the Typed-Event channel is that the <I> interface used by both suppliers and consumers must match. Therefore if EJBs are to be subscribed as typed consumers to the channel, then the suppliers will need to implement the RMI Remote interface of the EJB. This is the same issue faced by CORBA clients contacting EJBs discussed in section 4.1.1. The solutions discussed there can also be applied to this environment. EJBs which need to receive events from non-RMI suppliers could either use a proxy as described in section 4.1.1 or a simplified mapping of the EJB interface could be exposed to the supplier and subscribed to the channel.

Example 5 illustrates an EJB being subscribed to a Typed Event Channel.

Summary

The Typed Notification channel offers a more strictly typed environment for event suppliers and consumers to use. The Typed channel typically also provides the best performance of the four kinds of Notification channel. For these reasons the ITU-T/T1M1 recommends their use. Interfacing CORBA systems with EJBs via the Typed channel faces the same challenges as when these two domains are

interfaced for synchronous communication. Solutions to those problems (as discussed in section 4.1.1) are therefore applicable here too.

4.2.3 Comparison of Approaches

Two approaches were examined by which EJBs can receive events from the CosNotification service. The first approach bridges the CosNotification and JMS services, translating events from one domain to the other, thus delivering the events as JMS messages to a MDB and then to an EJB. The most significant challenge faced with this approach is the translation employed for structured data particularly structured data types defined by various telecommunications recommendations. Under this approach, parts of the event structure must be flattened into property lists of name value pairs. The disadvantages associated with this practice were enumerated in section 4.2.1.2.

The other approach is for an EJB to implement the CosNotification event consumer interface and be registered directly with the Notification channel so events can be delivered directly to the EJB. This requires an EJB to possess the operations of the Structured Consumer interface, and requires either the Notification Channel or some intermediary proxy to box events before being delivered to the EJB. With this approach, the event does not need to be flattened, and since no content translation is involved, the EJB code required to navigate and parse the event datatype will be very similar (if not identical) to that found in existing CORBA consumers. The benefit of this is that code reuse is more likely, or at least that the design can be reused thus leading to more confidence in the code quality.

For these reasons, the second approach to interworking EJBs with the Notification Service is considered more suitable for interfacing EJBs with the CosNotification service.

5. Appendix 1 - A Simplified IDL Mapping

As discussed in section 4.1.1, mapping the Home and Remote (RMI) interfaces of an EJB to non-Java languages (e.g. for use by non-Java clients of the EJB) can result in complex and complicated interfaces. Much of this complexity arises from trying to accommodate as many of the RMI semantics as possible in the resulting IDL and ultimately in the target implementation language. However, if the EJB's Home and Remote interfaces are relatively simple then such a comprehensive mapping (as specified by CORBA) is not necessary. In such cases, a trade-off can be made whereby some semantics of the RMI interface are not fully translated into the IDL and target language in order to simplify the resulting IDL; resulting in a simplified mapping of Java to IDL.

Such a simplified mapping is not standardized by the OMG, but is commercially available with the Borland Enterprise Server from Borland Software Corporation. Details of the assumptions made are explained in an example [3] which is included with the product.

As provided by that product, the simplified IDL mapping is able to map Java constructs to simpler IDL forms as compared to the regular Java to IDL mapping because it makes certain assumptions about what are acceptable semantics. For example, SIDL maps the datatype `java.util.Date` to a structure containing an integer value representing milliseconds since Jan 1, 1970 GMT; Java Collections, Enumerations, Lists and Vectors are all mapped directly to sequences of Any; and `java.lang.Class` is mapped to an IDL struct containing two strings - one being the name of the associated Java type and the other being the name of the associated IDL type. Such mappings convey the majority of the original datatype's meaning for most uses and are implemented by an accompanying library. The simplified IDL generated for EJB2 is shown below.

```
#include <sidl.idl>

typedef sequence<wstring> Sequence_of_Wstring;

module ossj {
  module example2 {
    module myEJBapp {
      interface EJB2 : ::sidl::javax::ejb::EJBObject {
        ::sidl::java::util::Date
          echoDate(in ::sidl::java::util::Date arg0);
        wstring echoString(in wstring arg0);
        ::Sequence_of_Wstring echoStrings(in ::Sequence_of_Wstring arg0);
      };
      ...
    };
  };
};
```

Figure 20. Some of the 'Simplified IDL' derived from EJB2's Remote Interface

The SIDL scheme thus presents a simplified IDL interface of an EJB. The scheme does not require the EJB to be modified or adapted, and is deployed to the Borland Enterprise Server's EJB container in the same way. Once deployed, the EJB container will offer the usual RMI-IIOP interface to the EJB based on the CORBA standard mapping of Java to IDL, and will also (if configured for an EJB) will offer another RMI-IIOP interface based on the simplified IDL mapping of the EJB. Then at runtime when a method of the EJB's simplified interface is invoked the Borland Enterprise Server's EJB container will translate the simplified types used into the appropriate Java types required by the EJB before dispatching the invocation to the EJB, and will perform the inverse translation for invocation replies. Since it is the EJB container which performs the translation to and from simplified constructs, the EJB itself does not need to be altered

to participate in this environment, with the caveat that the EJB's interface is expressible using simplified IDL. Similarly, the client operates as a standard CORBA client invoking a CORBA-compliant IDL, albeit a simplified one.

Thus, this approach to interworking non-Java callers with EJBs takes much of the complexity which previously would have existed in the client application and places it within the EJB container.

6. Appendix 2 - The OMG Notification Service

The Notification Service is a standardized service of the CORBA platform defined by the OMG ([8]). It provides applications with the ability to use decoupled communication between their elements instead of traditional direct/synchronous communication. The Notification Service supercedes a similar, but less sophisticated service—the Event Service[9]—by offering more features and maintaining backwards compatibility.

A key element of the Notification Service is the Notification Channel (referred to here as the Channel) whose role is to propagate events from suppliers to consumers. Once an event has been delivered to the Channel the Channel takes responsibility for delivering it to each subscribed consumer. This arrangement is shown in Figure 21.




Figure 21. The Notification / Event Channel

The default behavior of the Notification Channel is to deliver every event it receives to every subscribed consumer. This is also the behavior of the Event Channel. However, the Notification Channel has the facility to filter events and thereby provide selective delivery. To use this facility, consumers specify which events they are interested in receiving by registering a filter expression with the Notification Channel. The Channel then applies the filter expression to each event to determine whether it should be delivered to that consumer or not. This and other features, such as Quality of Service parameters, can be used to tailor the behavior of the Channel. However, these facilities are only mentioned here to provide an overview of the capabilities of the Channel. The examples discussed in this paper do not require such advanced facilities, even though their use in these cases is conceivable.

A consequence of application elements using the Notification Channel is that they no longer communicate directly with each other but indirectly via the Channel. There are many benefits arising from this decoupling, including the following:

- Supplier elements can deliver events at different rates to which consumer elements process them. And therefore, they can produce events at a different rate as well. In this respect, the Channel acts as a buffer, accommodating and leveling-out peaks in an application's processing activity.

- The absence or unavailability of consumer elements does not prevent supplier elements from delivering events. In this respect, the channel allows an application to continue functioning even when parts of that application are unavailable.
- A supplier can send an event to every consumer by creating a single event and delivering it to the Channel. In this capacity, the Channel acts as a broadcast medium for the application. If filtering is used in the Notification Channel, then the Channel acts as a multicast medium.
- The identity of consumers is not needed by suppliers in order to reach them; only the identity of the Channel is needed by consumers and suppliers. Because of this, suppliers and/or consumers can be introduced to a system without requiring reconfiguration of existing suppliers or consumers in order to accommodate them. This has enormous benefit for large distributed applications.

Distributed application architectures can use these characteristics of decoupled communication to improve their performance, reliability, scalability, and adaptability.

The four kinds of Notification Channel

The Notification Service specifies four different kinds of Notification Channel:

- the Untyped-Event Channel,
- the Structured-Event Channel,
- the Sequenced-Event Channel, and
- the Typed-Event Channel.

All four types of channel propagate events from suppliers to consumers, but each represents the event using a different form. This choice of representation allows applications to choose a channel most suited to their needs.

The Untyped-Event Notification Channel uses the CORBA::Any datatype to represent events. To deliver such datatypes, it also defines a set of predefined IDL interfaces with standard operations for communicating events between the channel and its clients. Typically, clients of the Untyped-Event Notification Channel will define their own data structure for storing event information and then package that into the Any datatype. This allows applications to continue using their own event datatype with the Notification Channel and is one of the primary advantages of the Untyped-Event Channel.

The Structured-Event and Sequenced-Event Notification Channels use predefined structured datatypes to represent events. For the Structured Channel, the datatype comprises a fixed header, plus a variable header, and a body. The headers can include standard fields, as well as name/value pairs. A complete definition of the structured event datatype is found in the Notification Service specification (Section 2.2, [8]) and an abbreviated version (applicable to the discussions in this document) is provided below.

```
// Definition of a Structured Event from the CosNotification.idl file

struct EventType {
    string domain_name;
    string type_name;
};

struct FixedEventHeader {
    EventType event_type;
    string event_name;
};

struct Property {
    PropertyName name;
    PropertyValue value;
};

typedef sequence<Property> PropertySeq;
typedef PropertySeq OptionalHeaderFields;
```

```

struct EventHeader {
    FixedEventHeader fixed_header;
    OptionalHeaderFields variable_header;
}; // continued over page

typedef PropertySeq FilterableEventBody;

// Define the Structured Event structure

struct StructuredEvent {
    EventHeader header;
    FilterableEventBody filterable_data;
    any remainder_of_body;
};

```

Figure 22. IDL definition of the CosNotification StructuredEvent datatype

The Sequenced-Event Channel's event datatype is a sequence of the Structured-Event Channel's event datatype. In this way, the Sequenced Channel can transport batches of events via a single invocation. A significant advantage of both Structured-and Sequenced-Event Channels is that their event filtering activities are much more efficient than the other kinds of channels because the Event Structure is known to the Channel *a priori* meaning that event demarshalling and inspection algorithms that must be implemented in the Channel can be significantly optimized at design time.

The fourth kind of Channel is the Typed-Event Notification Channel. This differs from the other kinds of channels in that it does not use an explicitly defined data-structure to represent events. Nor, therefore, does it provide predefined IDL operations such as `push()` or `push_structured_event()` to transfer events between client and channel. Instead, clients communicate with Typed Channels through application-defined interfaces. The Notification specification refers to these interfaces as <I> interfaces. Operations of an <I> interface which are to be used to transfer events must specify a void return type and not use any out or inout parameters. One or more event operations may be defined in an <I> interface. Calling an event operation of an <I> interface obtained from a Typed-Event Notification Channel causes the invocation together with its argument values to be sent to the Channel. The Channel then treats this invocation as a Typed-Event. Similarly, to receive Typed Events, consumer applications implement the <I> interface and subscribes that CORBA object to the Typed-Event Channel. The Channel then delivers the events by invoking the specified operation with the accompanying parameter values as received from the supplier. One significant advantage of using the Typed-Event Notification Channel is that an application's elements communicate via strongly typed application-level interfaces and therefore do not need to encode or decode to and from an event datatype when publishing or consuming events. In contrast, when using a Structured, Sequenced or Untyped Channel, an application must use a generic interface which requires events to be formatted into an explicit event structure, delivered using an infrastructure-level operation, and extracted from the event datatype. Non-Typed channels therefore force clients to use a contract which is semantically weaker, less object-oriented, and less type-safe than that offered by the Typed-Event Channel.

Although not covered here, it should be noted that the Notification Specification also defines rules for how events are to be mapped between their different forms—Untyped, Typed, Structured, and Sequenced. Hence, it is possible for suppliers and consumers to be bound to different kinds of Notification Channel yet interoperate with each other.

7. References

1. **Examples of Integrating EJBs with CORBA systems.** Java source code and documentation by Brenton Camac. <http://codecentral.borland.com/codecentral/ccWeb.exe/listing?id=20882>.
2. **Java Language Mapping to OMG IDL.** Version 1.3. (formal/03-09-04). The Object Management Group.
3. **A Simplified Java to IDL mapping for EJBs.** Java source code and documentation. Borland Software Corporation. <http://codecentral.borland.com/codecentral/ccWeb.exe/listing?id=20716>
4. **IDL to Java Language Mapping.** Version 1.2. (formal/2002-08-05). The Object Management Group.
5. **The Java Remote Method Invocation Specification.** Version 1.3. Sun Microsystems. <ftp://ftp.javasoft.com/docs/jdk1.2/rmi-spec-JDK1.2.pdf>
6. **Notification / JMS Interworking Specification.** Version 1.0 (dtc/03-06-01). The Object Management Group.
7. **The JMS Specification.** Version 1.0.2. Sun Microsystems. <http://docs.sun.com/db/doc/816-5904-10>
8. **Notification Service Specification.** Version 1.0. (formal/2000-06-20). The Object Management Group.
9. **Event Service Specification.** Version 1.1. (formal/2001-03-01). The Object Management Group.