
Event Service Specification

New Edition: June 2000
Version 1.0

Copyright 1993, DEC
Copyright 1993, Groupe Bull
Copyright 1993, Hewlett-Packard
Copyright 1993, HyperDesk
Copyright 1993, Itasca
Copyright 1993, Novell
Copyright 1993, O2
Copyright 1993, Object Design
Copyright 2000, Object Management Group, Inc.
Copyright 1993, Objectivity
Copyright 1993, Ontos
Copyright 1993, Oracle
Copyright 1993, Servio
Copyright 1993, SunSoft
Copyright 1993, Tivoli
Copyright 1993, Versant

The companies listed above have granted to the Object Management Group, Inc. (OMG) a nonexclusive, royalty-free, paid up, worldwide license to copy and distribute this document and to modify this document and distribute copies of the modified version. Each of the copyright holders listed above has agreed that no person shall be deemed to have infringed the copyright in the included material of any such copyright holder by reason of having used the specification set forth herein or having conformed any computer software to the specification.

PATENT

The attention of adopters is directed to the possibility that compliance with or adoption of OMG specifications may require use of an invention covered by patent rights. OMG shall not be responsible for identifying patents for which a license may be required by any OMG specification, or for conducting legal inquiries into the legal validity or scope of those patents that are brought to its attention. OMG specifications are prospective and advisory only. Prospective users are responsible for protecting themselves against liability for infringement of patents.

NOTICE

The information contained in this document is subject to change without notice. The material in this document details an Object Management Group specification in accordance with the license and notices set forth on this page. This document does not represent a commitment to implement any portion of this specification in any company's products.

WHILE THE INFORMATION IN THIS PUBLICATION IS BELIEVED TO BE ACCURATE, THE OBJECT MANAGEMENT GROUP AND THE COMPANIES LISTED ABOVE MAKE NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS MATERIAL INCLUDING, BUT NOT LIMITED TO ANY WARRANTY OF TITLE OR OWNERSHIP, IMPLIED WARRANTY OF MERCHANTABILITY OR WARRANTY OF FITNESS FOR PARTICULAR PURPOSE OR USE. In no event shall The Object Management Group or any of the companies listed above be liable for errors contained herein or for indirect, incidental, special, consequential, reliance or cover damages, including loss of profits, revenue, data or use, incurred by any user or any third party. The copyright holders listed above acknowledge that the Object Management Group (acting itself or through its designees) is and shall at all times be the sole entity that may authorize developers, suppliers and sellers of computer software to use certification marks, trademarks or other special designations to indicate compliance with these materials. This document contains information which is protected by copyright. All Rights Reserved. No part of this work covered by copyright herein may be reproduced or used in any form or by any means--graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems--without permission of the copyright owner.

RESTRICTED RIGHTS LEGEND. Use, duplication, or disclosure by government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Right in Technical Data and Computer Software Clause at DFARS 252.227.7013 OMG[®] and Object Management are registered trademarks of the Object Management Group, Inc. Object Request Broker, OMG IDL, ORB, CORBA, CORBAfacilities, CORBAservices, COSS, and IIOP are trademarks of the Object Management Group, Inc. X/Open is a trademark of X/Open Company Ltd.

ISSUE REPORTING

All OMG specifications are subject to continuous review and improvement. As part of this process we encourage readers to report any ambiguities, inconsistencies, or inaccuracies they may find by completing the issue reporting form at <http://www.omg.org/library/issuerpt.htm>.

Contents

Preface	v
About This Document	v
Object Management Group	v
0.0.1 What is CORBA?	vi
0.0.2 X/Open	vi
Intended Audience	vi
Need for Object Services	vi
What Is an Object Service Specification?	vii
Associated OMG Documents	vii
Service Design Principles	viii
Build on CORBA Concepts	viii
Basic, Flexible Services	ix
Generic Services	ix
Allow Local and Remote Implementations	ix
Quality of Service is an Implementation Characteristic ..	ix
Objects Often Conspire in a Service	ix
Use of Callback Interfaces	x
Assume No Global Identifier Spaces	x
Finding a Service is Orthogonal to Using It	xi
Interface Style Consistency	xi
Use of Exceptions and Return Codes	xi
Explicit Versus Implicit Operations	xi
Use of Interface Inheritance	xi
Acknowledgments	xi

1. Service Description	1-1
1.1 Overview	1-1
1.2 Event Communication	1-2
1.3 Example Scenario	1-3
1.4 Design Principles	1-4
1.5 Resolution of Technical Issues	1-5
1.6 Quality of Service	1-6
1.7 Generic Event Communication	1-7
1.7.1 Push Model	1-7
1.7.2 Pull Model	1-7
2. Modules and Interfaces	2-1
2.1 The CosEventComm Module	2-1
2.1.1 The PushConsumer Interface	2-2
2.1.2 The PushSupplier Interface	2-2
2.1.3 The PullSupplier Interface	2-3
2.1.4 The PullConsumer Interface	2-3
2.2 Event Channels	2-3
2.2.1 Push-Style Communication with an Event Channel	2-4
2.2.2 Pull-Style Communication with an Event Channel	2-4
2.2.3 Mixed Style Communication with an Event Channel	2-5
2.2.4 Multiple Consumers and Multiple Suppliers	2-5
2.2.5 Event Channel Administration	2-6
2.3 The CosEventChannelAdmin Module	2-8
2.3.1 The EventChannel Interface	2-9
2.3.2 The ConsumerAdmin Interface	2-10
2.3.3 The SupplierAdmin Interface	2-10
2.3.4 The ProxyPushConsumer Interface	2-10
2.3.5 The ProxyPullSupplier Interface	2-11
2.3.6 The ProxyPullConsumer Interface	2-11
2.3.7 The ProxyPushSupplier Interface	2-12
2.4 Typed Event Communication	2-12
2.4.1 Typed Push Model	2-12
2.4.2 Typed Pull Model	2-13
2.5 The CosTypedEventComm Module	2-14
2.5.1 The TypedPushConsumer Interface	2-14
2.5.2 The TypedPullSupplier Interface	2-15
2.6 Typed Event Channels	2-16

2.7	The CosTypedEventChannelAdmin Module	2-16
2.7.1	The TypedEventChannel Interface	2-17
2.7.2	The TypedConsumerAdmin Interface	2-17
2.7.3	The TypedSupplierAdmin Interface	2-18
2.7.4	The TypedProxyPushConsumer Interface	2-19
2.7.5	The TypedProxyPullSupplier Interface	2-19
2.8	Composing Event Channels and Filtering	2-20
2.9	Policies for Finding Event Channels	2-20
	Appendix A - Implementing Typed Event Channels	A-1
	Appendix B - An Event Channel Use Example	B-1

Contents

Preface

About This Document

Under the terms of the collaboration between OMG and X/Open Co Ltd, this document is a candidate for endorsement by X/Open, initially as a Preliminary Specification and later as a full CAE Specification. The collaboration between OMG and X/Open Co Ltd ensures joint review and cohesive support for emerging object-based specifications.

X/Open Preliminary Specifications undergo close scrutiny through a review process at X/Open before publication and are inherently stable specifications. Upgrade to full CAE Specification, after a reasonable interval, takes place following further review by X/Open. This further review considers the implementation experience of members and the full implications of conformance and branding.

Object Management Group

The Object Management Group, Inc. (OMG) is an international organization supported by over 800 members, including information system vendors, software developers and users. Founded in 1989, the OMG promotes the theory and practice of object-oriented technology in software development. The organization's charter includes the establishment of industry guidelines and object management specifications to provide a common framework for application development. Primary goals are the reusability, portability, and interoperability of object-based software in distributed, heterogeneous environments. Conformance to these specifications will make it possible to develop a heterogeneous applications environment across all major hardware platforms and operating systems.

OMG's objectives are to foster the growth of object technology and influence its direction by establishing the Object Management Architecture (OMA). The OMA provides the conceptual infrastructure upon which all OMG specifications are based.

What is CORBA?

The Common Object Request Broker Architecture (CORBA), is the Object Management Group's answer to the need for interoperability among the rapidly proliferating number of hardware and software products available today. Simply stated, CORBA allows applications to communicate with one another no matter where they are located or who has designed them. CORBA 1.1 was introduced in 1991 by Object Management Group (OMG) and defined the Interface Definition Language (IDL) and the Application Programming Interfaces (API) that enable client/server object interaction within a specific implementation of an Object Request Broker (ORB). CORBA 2.0, adopted in December of 1994, defines true interoperability by specifying how ORBs from different vendors can interoperate.

X/Open

X/Open is an independent, worldwide, open systems organization supported by most of the world's largest information system suppliers, user organizations and software companies. Its mission is to bring to users greater value from computing, through the practical implementation of open systems.

Intended Audience

The specifications described in this manual are aimed at software designers and developers who want to produce applications that comply with OMG standards for object services; the benefits of compliance are outlined in the following section, "Need for Object Services."

Need for Object Services

To understand how Object Services benefit all computer vendors and users, it is helpful to understand their context within OMG's vision of object management. The key to understanding the structure of the architecture is the Reference Model, which consists of the following components:

- **Object Request Broker**, which enables objects to transparently make and receive requests and responses in a distributed environment. It is the foundation for building applications from distributed objects and for interoperability between applications in hetero- and homogeneous environments. The architecture and specifications of the Object Request Broker are described in *CORBA: Common Object Request Broker Architecture and Specification*.
- **Object Services**, a collection of services (interfaces and objects) that support basic functions for using and implementing objects. Services are necessary to construct any distributed application and are always independent of application domains.
- **Common Facilities**, a collection of services that many applications may share, but which are not as fundamental as the Object Services. For instance, a system management or electronic mail facility could be classified as a common facility.

The Object Request Broker, then, is the core of the Reference Model. Nevertheless, an Object Request Broker alone cannot enable interoperability at the application semantic level. An ORB is like a telephone exchange: it provides the basic mechanism for making and receiving calls but does not ensure meaningful communication between subscribers. Meaningful, productive communication depends on additional interfaces, protocols, and policies that are agreed upon outside the telephone system, such as telephones, modems and directory services. This is equivalent to the role of Object Services.

What Is an Object Service Specification?

A specification of an Object Service usually consists of a set of interfaces and a description of the service's behavior. The syntax used to specify the interfaces is the OMG Interface Definition Language (OMG IDL). The semantics that specify a services's behavior are, in general, expressed in terms of the OMG Object Model. The OMG Object Model is based on objects, operations, types, and subtyping. It provides a standard, commonly understood set of terms with which to describe a service's behavior.

(For detailed information about the OMG Reference Model and the OMG Object Model, refer to the *Object Management Architecture Guide*).

Associated OMG Documents

The CORBA documentation is organized as follows:

- *Object Management Architecture Guide* defines the OMG's technical objectives and terminology and describes the conceptual models upon which OMG standards are based. It defines the umbrella architecture for the OMG standards. It also provides information about the policies and procedures of OMG, such as how standards are proposed, evaluated, and accepted.
- CORBA Platform Technologies
 - *CORBA: Common Object Request Broker Architecture and Specification* contains the architecture and specifications for the Object Request Broker.
 - *CORBA Languages*, a collection of language mapping specifications. See the individual language mapping specifications.
 - *CORBA Services*, a collection of specifications for OMG's Object Services. See the individual service specifications.
 - *CORBA Facilities*, a collection of specifications for OMG's Common Facilities. See the individual facility specifications.
- CORBA Domain Technologies
 - *CORBA Manufacturing*, a collection of specifications that relate to the manufacturing industry. This group of specifications defines standardized object-oriented interfaces between related services and functions.
 - *CORBA Med*, a collection of specifications that relate to the healthcare industry and represents vendors, healthcare providers, payers, and end users.

-
- *CORBA Finance*, a collection of specifications that target a vitally important vertical market: financial services and accounting. These important application areas are present in virtually all organizations: including all forms of monetary transactions, payroll, billing, and so forth.
 - *CORBA Telecoms*, a collection of specifications that relate to the OMG-compliant interfaces for telecommunication systems.

The OMG collects information for each book in the documentation set by issuing Requests for Information, Requests for Proposals, and Requests for Comment and, with its membership, evaluating the responses. Specifications are adopted as standards only when representatives of the OMG membership accept them as such by vote. (The policies and procedures of the OMG are described in detail in the *Object Management Architecture Guide*.)

To obtain print-on-demand books in the documentation set or other OMG publications, contact the Object Management Group, Inc. at:

OMG Headquarters
250 First Avenue, Suite 201
Needham, MA 02494
USA
Tel: +1-781-444-0404
Fax: +1-781-444-0320
pubs@omg.org
<http://www.omg.org>

Service Design Principles

Build on CORBA Concepts

The design of each Object Service uses and builds on CORBA concepts:

- Separation of interface and implementation
- Object references are typed by interfaces
- Clients depend on interfaces, not implementations
- Use of multiple inheritance of interfaces
- Use of subtyping to extend, evolve and specialize functionality

Other related principles that the designs adhere to include:

- Assume good ORB and Object Services implementations. Specifically, it is assumed that CORBA-compliant ORB implementations are being built that support efficient local and remote access to “fine-grain” objects and have performance characteristics that place no major barriers to the pervasive use of distributed objects for virtually all service and application elements.
- Do not build non-type properties into interfaces

A discussion and rationale for the design of object services was included in the HP-SunSoft response to the OMG Object Services RFI (OMG TC Document 92.2.10).

Basic, Flexible Services

The services are designed to do one thing well and are only as complicated as they need to be. Individual services are by themselves relatively simple yet they can, by virtue of their structuring as objects, be combined together in interesting and powerful ways.

For example, the event and life cycle services, plus a future relationship service, may play together to support graphs of objects. Object graphs commonly occur in the real world and must be supported in many applications. A functionally-rich Folder compound object, for example, may be constructed using the life cycle, naming, events, and future relationship services as “building blocks.”

Generic Services

Services are designed to be generic in that they do not depend on the type of the client object nor, in general, on the type of data passed in requests. For example, the event channel interfaces accept event data of any type. Clients of the service can dynamically determine the actual data type and handle it appropriately.

Allow Local and Remote Implementations

In general the services are structured as CORBA objects with OMG IDL interfaces that can be accessed locally or remotely and which can have local library or remote server styles of implementations. This allows considerable flexibility as regards the location of participating objects. So, for example, if the performance requirements of a particular application dictate it, objects can be implemented to work with a Library Object Adapter that enables their execution in the same process as the client.

Quality of Service is an Implementation Characteristic

Service interfaces are designed to allow a wide range of implementation approaches depending on the quality of service required in a particular environment. For example, in the Event Service, an event channel can be implemented to provide fast but unreliable delivery of events or slower but guaranteed delivery. However, the interfaces to the event channel are the same for all implementations and all clients. Because rules are not wired into a complex type hierarchy, developers can select particular implementations as building blocks and easily combine them with other components.

Objects Often Conspire in a Service

Services are typically decomposed into several distinct interfaces that provide different views for different kinds of clients of the service. For example, the Event Service is composed of *PushConsumer*, *PullSupplier* and *EventChannel* interfaces. This simplifies the way in which a particular client uses a service.

A particular service implementation can support the constituent interfaces as a single CORBA object or as a collection of distinct objects. This allows considerable implementation flexibility. A client of a service may use a different object reference to communicate with each distinct service function. Conceptually, these “internal” objects *conspire* to provide the complete service.

As an example, in the Event Service an event channel can provide both *PushConsumer* and *EventChannel* interfaces for use by different kinds of client. A particular client sends a request not to a single “event channel” object but to an object that implements either the *PushConsumer* and *EventChannel* interface. Hidden to all the clients, these objects interact to support the service.

The service designs also use distinct objects that implement specific service interfaces as the means to distinguish and coordinate different clients without relying on the existence of an object equality test or some special way of identifying clients. Using the event service again as an example, when an event consumer is connected with an event channel, a new object is created that supports the *PullSupplier* interface. An object reference to this object is returned to the event consumer which can then request events by invoking the appropriate operation on the new “supplier” object. Because each client uses a different object reference to interact with the event channel, the event channel can keep track of and manage multiple simultaneous clients. An event channel as a collection of objects conspiring to manage multiple simultaneous consumer clients.

Use of Callback Interfaces

Services often employ callback interfaces. Callback interfaces are interfaces that a client object is required to support to enable a service to *call back* to it to invoke some operation. The callback may be, for example, to pass back data asynchronously to a client.

Callback interfaces have two major benefits:

- They clearly define how a client object participates in a service.
- They allow the use of the standard interface definition (OMG IDL) and operation invocation (object reference) mechanisms.

Assume No Global Identifier Spaces

Several services employ identifiers to label and distinguish various elements. The service designs do not assume or rely on any global identifier service or global id spaces in order to function. The scope of identifiers is always limited to some context. For example, in the naming service, the scope of names is the particular naming context object.

In the case where a service generates ids, clients can assume that an id is unique within its scope but should not make any other assumption.

Finding a Service is Orthogonal to Using It

Finding a service is at a higher level and orthogonal to using a service. These services do not dictate a particular approach. They do not, for example, mandate that all services must be found via the naming service. Because services are structured as objects there does not need to be a special way of finding objects associated with services - general purpose finding services can be used. Solutions are anticipated to be application and policy specific.

Interface Style Consistency

Use of Exceptions and Return Codes

Throughout the services, exceptions are used exclusively for handling exceptional conditions such as error returns. Normal return codes are passed back via output parameters. An example of this is the use of a DONE return code to indicate iteration completion.

Explicit Versus Implicit Operations

Operations are always explicit rather than implied (e.g., by a flag passed as a parameter value to some “umbrella” operation). In other words, there is always a distinct operation corresponding to each distinct function of a service.

Use of Interface Inheritance

Interface inheritance (subtyping) is used whenever one can imagine that client code should depend on less functionality than the full interface. Services are often partitioned into several unrelated interfaces when it is possible to partition the clients into different roles. For example, an administrative interface is often unrelated and distinct in the type system from the interface used by “normal” clients.

Acknowledgments

The following companies submitted parts of the *Event Service* specification:

- DEC
- Groupe Bull
- Hewlett-Packard Company
- HyperDesk
- Itasca
- Novell
- O2
- Object Design
- Objectivity
- Ontos

-
- Oracle
 - Servio
 - SunSoft
 - Tivoli
 - Versant

Service Description

1

Contents

This chapter contains the following topics.

Topic	Page
“Overview	1-1
“Event Communication	1-2
“Example Scenario	1-3
“Design Principles	1-4
“Resolution of Technical Issues	1-5
“Quality of Service	1-6
“Generic Event Communication	1-7

1.1 Overview

A standard CORBA request results in the synchronous execution of an operation by an object. If the operation defines parameters or return values, data is communicated between the client and the server. A request is directed to a particular object. For the request to be successful, both the client and the server must be available. If a request fails because the server is unavailable, the client receives an exception and must take some appropriate action.

In some scenarios, a more decoupled communication model between objects is required. For example:

- A system administration tool is interested in knowing if a disk runs out of space. The software managing a disk is unaware of the existence of the system administration tool. The software simply reports that the disk is full. When a disk runs out of space, the system administration tool opens a window to inform the user which disk has run out of space.
- A property list object is associated with an application object. The property list object is physically separate from the application object. The application object is interested in the changes made to its properties by a user. The properties can be changed without involving the application object. That is, in order to have reasonable response time for the user, changing a property does not activate the application object. However, when the application object is activated, it needs to know about the changes to its properties.
- A CASE tool is interested in being notified when a source program has been modified. The source program simply reports when it is modified. It is unaware of the existence of the CASE tool. In response to the notification, the CASE tool invokes a compiler.
- Several documents are linked to a spreadsheet. The documents are interested in knowing when the value of certain cells have changed. When the cell value changes, the documents update their presentations based on the spreadsheet. Furthermore, if a document is unavailable because of a failure, it is still interested in any changes to the cells and wants to be notified of those changes when it recovers.

1.2 Event Communication

The Event Service decouples the communication between objects. The Event Service defines two roles for objects: the supplier role and the consumer role. *Suppliers* produce event data and *consumers* process event data. Event data are communicated between suppliers and consumers by issuing standard CORBA requests.

There are two approaches to initiating event communication between suppliers and consumers, and two orthogonal approaches to the form that the communication can take.

The two approaches to initiating event communication are called the *push model* and the *pull model*. The push model allows a supplier of events to initiate the transfer of the event data to consumers. The pull model allows a consumer of events to request the event data from a supplier. In the push model, the supplier is taking the initiative; in the pull model, the consumer is taking the initiative.

The communication itself can be either generic or typed. In the generic case, all communication is by means of generic push or pull operations that take a single parameter that packages all the event data. In the typed case, communication is via operations defined in OMG IDL. Event data is passed by means of the parameters, which can be defined in any manner desired.

An *event channel* is an intervening object that allows multiple suppliers to communicate with multiple consumers asynchronously. An event channel is both a consumer and a supplier of events. Event channels are standard CORBA objects and communication with an event channel is accomplished using standard CORBA requests.

1.3 Example Scenario

This section provides a general scenario that illustrates how the Event Service can be used.

The Event Service can be used to provide “change notification.” When an object is changed (its state is modified), an event can be generated that is propagated to all interested parties. For example, when a spreadsheet cell object is modified, all compound documents which contain a reference (link) to that cell can be notified (so the document can redisplay the referenced cell, or recalculate values that depend on the cell). Similarly, when an engineering specification object is modified, all engineers who have registered an interest in the specification can be notified that the specification has changed.

In this scenario, objects that can be “changed” act as suppliers, parties interested in receiving notifications of changes act as consumers, and one or more event channel objects are used as intermediaries between consumers and suppliers. *Either the push or the pull model can be used at either end.*

If the push model is used by suppliers, objects that can be changed support the **PushSupplier** interface so that event communication can be discontinued. Use the **EventChannel**, the **SupplierAdmin**, and the **ProxyPushConsumer** interfaces to register as suppliers of events, and use the **ProxyPushConsumer** interface to push events to event channels.

When a change occurs to an object, a changeable object invokes a push operation on the channel. It provides as an argument to the push operation information that describes the event. This information is of data type any - it can be as simple or as complex as is necessary. For example, the event information might identify the object reference of the object that has been changed, it might identify the kind of change that has occurred, it might provide a new displayable image of the changed object or it might identify one or more additional objects that describe the change that has been made.

If the pull model is used by consumers, all client objects that want to be notified of changes support the **PullConsumer** interface so communication can be discontinued, using the **EventChannel**, **ConsumerAdmin**, and **ProxyPullSupplier** interfaces to register as consumers of events, and using the **ProxyPullSupplier** interface to pull events from event channels.

The consumer may use either a blocking or non-blocking mechanism for receiving notification of changes. Using the **try_pull** operation, the consumer can periodically poll the channel for events. Alternatively, the consumer can use the pull operation which will block the consumer’s execution thread until an event is generated by some supplier.

Event channels act as the intermediaries between the objects being changed and objects interested in knowing about changes. The channels that provide change notification can be general purpose, well-known objects (e.g., “persistent server-based objects” that are run as part of a workgroup-wide framework of objects that provide “desktop services”) or specific-to-task objects (e.g., temporary objects that are created when needed). Objects that use event channels may locate the channels by looking for them in a persistently available server (e.g., by looking for them in a naming service) or they may be given references to these objects as part of a specific-to-task object protocol (e.g., when an “open” operation is invoked on an object, the object may return the reference to an event channel which the caller should use until the object is closed).

Event channels determine how changes are propagated between suppliers and consumers, i.e., the qualities of service. For example, an event channel determines the persistence of an event. The channel may keep an event for a specified period of time, passing it along to any consumer who registers with the channel during that period of time (e.g., it may keep event notifications about changes to engineering specifications for a week). Alternatively, the channel may only pass on events to consumers who are currently waiting for notification of changes (e.g., notifications of changes to a spreadsheet cell may only be sent to consumers who are currently displaying that cell).

This scenario exemplifies one way the event service described here forms a basic building block used in providing higher-level services specific to an application or common facilities framework of objects.

Instead of using the generic event channel, a typed event channel could also have been used.

1.4 Design Principles

The Event Service design satisfies the following principles:

- Events work in a distributed environment. The design does not depend on any global, critical, or centralized service.
- Event services allow multiple consumers of an event and multiple event suppliers.
- Consumers can either request events or be notified of events, whichever is more appropriate for application design and performance.
- Consumers and suppliers of events support standard OMG IDL interfaces; no extensions to CORBA are necessary to define these interfaces.
- A supplier can issue a single standard request to communicate event data to all consumers at once.
- Suppliers can generate events without knowing the identities of the consumers. Conversely, consumers can receive events without knowing the identities of the suppliers.
- The Event Service interfaces allow multiple qualities of service, for example, for different levels of reliability. It also allows for future interface extensions, such as for additional functionality.
- The Event Service interfaces are capable of being implemented and used in different operating environments, for example, in environments that support threading and those that do not.

1.5 Resolution of Technical Issues

This specification addresses the issues identified for event services in the OMG *Object Services Architecture*¹ document as follows:

- **Distributed environment:** The interfaces are designed to allow consumers and suppliers of events to be disconnected from time to time, and do not require centralized event identification, processing, routing, or other services that might be a bottleneck or a single point of failure.

Events themselves are *not* objects because the CORBA distributed object model does not support passing objects by value.

Event generation: The specification describes how events are generated and delivered in a very general fashion, with event channels as intermediate routing points. It does not require (or preclude) polling, nor does it require that an event supplier directly notify every interested party.

Events involving multiple objects: Complex events may be handled by constructing a notification tree of event consumer/suppliers checking for successively more specific event predicates. The specification does not require a general or global event predicate evaluation service as this may not be sufficiently reliable, efficient, or secure in a distributed, heterogeneous (potentially decoupled) environment.

Scoping, grouping, and filtering events: The specification takes advantage of CORBA's distributed scoping and grouping mechanisms for the identifier and type of events. Event filtering is easily achieved through event channels that selectively deliver events from suppliers to consumers. Event channels can be composed; that is, one event channel can consume events supplied by another. Typed event channels can provide filtering based on event type.

Registration and generation of events: Consumers and suppliers register with event channels themselves. Event channels are objects and they are found by any fashion that objects can be found. A global registration service is not required; any object that conforms to the IDL interface may consume an event.

Event parameters: The specification supports a parameter of type any that can be delivered with an event, used for application-specific data.

Forgery and secure events: Because event suppliers are objects, the specification leverages any ORB work on security for object references and communication.

1. *Object Services Architecture*, Document Number 92-8-4, Object Management Group, Framingham, MA, 1992.

Performance: The design is a minimalist one, and requires only one ORB call per event received. It supports both push-style and pull-style notification to avoid inefficient event polling. Since event suppliers, consumers, and channels are all ORB objects, the service directly benefits from a Library Object Adapter or any other ORB optimizations.

Formalized Event Information: For specific application environments and frameworks it may be beneficial to formalize the data associated with an event (defined in this specification as type any). This can be accomplished by defining a typed structure for this information. Depending on the needs of the environment, the kinds of information included might be a priority, timestamp, origin string, and confirmation indicator. This information might be solely for the benefit of the event consumer or might also be interpreted by particular event channel implementations.

Confirmation of Reception: Some applications may require that consumers of an event provide an explicit confirmation of reception back to the supplier. This can be supported effectively using a “reverse” event channel through which consumers send back confirmations as normal events. This obviates the need for any special confirmation mechanism. However, strict atomic delivery between all suppliers and all consumers requires additional interfaces.

1.6 *Quality of Service*

Application domains requiring event-style communication have diverse reliability requirements, from “at-most-once” semantics (best effort) to guaranteed “exactly-once” semantics, availability requirements, throughput requirements, performance requirements (i.e., how fast events are disseminated), and scalability requirements.

Clearly no single implementation of the Event Service can optimize such a diverse range of technical requirements. Hence, multiple implementations of event services are to be expected, with different services targeted toward different environments. As such, the event interfaces do not dictate *qualities of service*. Different implementations of the Event Service interfaces can support different qualities of service to meet different application needs.

For example, an implementation that trades at most once delivery to a single consumer in favor of performance is useful for some applications; an implementation that favors performance but cannot preclude duplicate delivery is useful for other applications. Both are acceptable implementations of the interfaces described in this chapter.

Clearly, an implementation of an event channel that discards all events is *not a useful* implementation. Useful implementations will at least support “best-effort” delivery of events.

Note that the interfaces defined in this chapter are incomplete for implementations that support strict notions of atomicity. That is, additional interfaces are needed by an implementation to guarantee that either all consumers receive an event or none of the consumers receive an event; and that all events are received in the same order by all consumers.

1.7 Generic Event Communication

There are two basic models for communicating event data between suppliers and consumers: the *push model* and the *pull model*.

1.7.1 Push Model

In the push model, suppliers “push” event data to consumers; that is, suppliers communicate event data by invoking push operations on the **PushConsumer** interface.

To set up a push-style communication, consumers and suppliers exchange **PushConsumer** and **PushSupplier** object references. Event communication can be broken by invoking a **disconnect_push_consumer** operation on the **PushConsumer** interface or by invoking a **disconnect_push_supplier** operation on the **PushSupplier** interface. If the **PushSupplier** object reference is nil, the connection cannot be broken via the supplier.

Figure 1-1 illustrates push-style communication between a supplier and a consumer.




Figure 1-1 Push-style Communication Between a Supplier and a Consumer

1.7.2 Pull Model

In the pull model, consumers “pull” event data from suppliers; that is, consumers request event data by invoking **pull** operations on the **PullSupplier** interface.

To set up a pull-style communication, consumers and suppliers must exchange **PullConsumer** and **PullSupplier** object references. Event communication can be broken by invoking a **disconnect_pull_consumer** operation on the **PullConsumer** interface or by invoking a **disconnect_pull_supplier** operation on the **PullSupplier** interface. If the **PullConsumer** object reference is nil, the connection cannot be broken via the consumer.

Figure 1-2 illustrates pull-style communication between a supplier and a consumer.




Figure 1-2 Pull-style Communication Between a Supplier and a Consumer

Contents

This chapter contains the following topics.

Topic	Page
“The CosEventComm Module”	2-1
“Event Channels”	2-3
“The CosEventChannelAdmin Module”	2-8
“Typed Event Communication”	2-12
“The CosTypedEventComm Module”	2-14
“Typed Event Channels”	2-16
“The CosTypedEventChannelAdmin Module”	2-16
“Composing Event Channels and Filtering”	2-20
“Policies for Finding Event Channels”	2-20

2.1 The CosEventComm Module

The communication styles shown in Chapter 1 are both supported by four simple interfaces: *PushConsumer*, *PushSupplier*, and *PullSupplier* and *PullConsumer*. These interfaces are defined in an OMG IDL module named **CosEventComm**, as shown below.

```
module CosEventComm {  
  
    exception Disconnected{};
```

```
interface PushConsumer {
    void push (in any data) raises(Disconnected);
    void disconnect_push_consumer();
};

interface PushSupplier {
    void disconnect_push_supplier();
};

interface PullSupplier {
    any pull () raises(Disconnected);
    any try_pull (out boolean has_event)
        raises(Disconnected);
    void disconnect_pull_supplier();
};

interface PullConsumer {
    void disconnect_pull_consumer();
};

};
```

2.1.1 The *PushConsumer* Interface

A push-style consumer supports the **PushConsumer** interface to receive event data.

```
interface PushConsumer {
    void push (in any data) raises(Disconnected);
    void disconnect_push_consumer();
};
```

A supplier communicates event data to the consumer by invoking the **push** operation and passing the event data as a parameter. If the event communication has already been disconnected, the **Disconnected** exception is raised.

The **disconnect_push_consumer** operation terminates the event communication; it releases resources used at the consumer to support the event communication. The *PushConsumer* object reference is disposed.

2.1.2 The *PushSupplier* Interface

A push-style supplier supports the **PushSupplier** interface.

```
interface PushSupplier {
    void disconnect_push_supplier();
};
```

The **disconnect_push_supplier** operation terminates the event communication; it releases resources used at the supplier to support the event communication. The *PushSupplier* object reference is disposed.

2.1.3 The PullSupplier Interface

A pull-style supplier supports the **PullSupplier** interface to transmit event data.

```
interface PullSupplier {
    any pull () raises(Disconnected);
    any try_pull (out boolean has_event)
        raises(Disconnected);
    void disconnect_pull_supplier();
};
```

A consumer requests event data from the supplier by invoking either the **pull** operation or the **try_pull** operation on the supplier.

- The **pull** operation blocks until the event data is available or an exception is raised.¹ It returns the event data to the consumer. If the event communication has already been disconnected, the **Disconnected** exception is raised.
- The **try_pull** operation does not block: if the event data is available, it returns the event data and sets the **has_event** parameter to *true*; if the event is not available, it sets the **has_event** parameter to *false* and the event data is returned as long with an undefined value. If the event communication has already been disconnected, the **Disconnected** exception is raised.

The **disconnect_pull_supplier** operation terminates the event communication; it releases resources used at the supplier to support the event communication. The *PullSupplier* object reference is disposed.

2.1.4 The PullConsumer Interface

A pull-style consumer supports the **PullConsumer** interface.

```
interface PullConsumer {
    void disconnect_pull_consumer();
};
```

The **disconnect_pull_consumer** operation terminates the event communication; it releases resources used at the consumer to support the event communication. The *PullConsumer* object reference is disposed.

2.2 Event Channels

The *event channel* is a service that decouples the communication between suppliers and consumers. The event channel is itself both a consumer and a supplier of the event data.

¹. This, of course, may be a standard CORBA exception.

An event channel can provide asynchronous communication of event data between suppliers and consumers. Although consumers and suppliers communicate with the event channel using standard CORBA requests, the event channel does not need to supply the event data to its consumer at the same time it consumes the data from its supplier.

2.2.1 Push-Style Communication with an Event Channel

The supplier pushes event data to the event channel; the event channel, in turn, pushes event data to the consumer. Figure 2-1 illustrates a push-style communication between a supplier and the event channel, and a consumer and the event channel.




Figure 2-1 Push-style Communication Between a Supplier and an Event Channel, and a Consumer and an Event Channel

2.2.2 Pull-Style Communication with an Event Channel

The consumer pulls event data from the event channel; the event channel, in turn, pulls event data from the supplier. Figure 2-2 illustrates a pull-style communication between a supplier and the event channel, and a consumer and the event channel.




Figure 2-2 Pull-style communication between a supplier and an event channel and a consumer and the event channel

2.2.3 Mixed Style Communication with an Event Channel

An event channel can communicate with a supplier using one style of communication, and communicate with a consumer using a different style of communication.

Figure 2-3 illustrates a push-style communication between a supplier and an event channel, and a pull-style communication between a consumer and the event channel. The consumer pulls the event data that the supplier has pushed to the event channel.




Figure 2-3 Push-style Communication Between a Supplier and an Event Channel, and Pull-style Communication Between a Consumer and an Event Channel

2.2.4 Multiple Consumers and Multiple Suppliers

Figure 2-1, Figure 2-2, and Figure 2-3 illustrate event channels with a single supplier and a single consumer. An event channel can also provide many-to-many communication. The channel consumes events from one or more suppliers, and supplies events to one or more consumers. Subject to the quality of service of a particular implementation, an event channel provides an event to all consumers.

Figure 2-4 illustrates an event channel with multiple push-style consumers and multiple push-style suppliers.

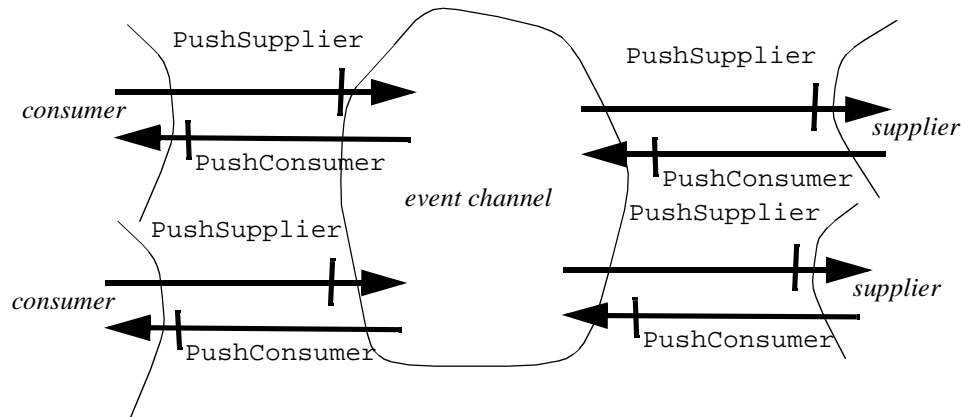


Figure 2-4 An Event Channel with Multiple Suppliers and Multiple Consumers

An event channel can support consumers and suppliers using different communication models.

If an event channel has at least one push-style consumer or at least one pending pull request, the event channel requires an event. If the event channel has pull suppliers, it will issue a request on a pull supplier to satisfy its requirement.

2.2.5 Event Channel Administration

The event channel is built up incrementally. When an event channel is created, no suppliers or consumers are connected to the event channel. Upon creation of the channel, the factory returns an object reference that supports the **EventChannel** interface, as illustrated in Figure 2-5.

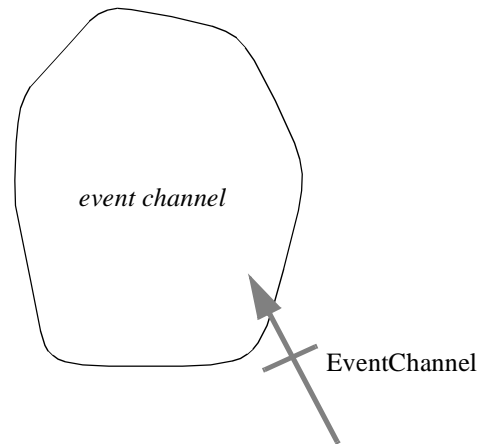


Figure 2-5 A newly created event channel. The channel has no suppliers or consumers.

The **EventChannel** interface defines three administrative operations: an operation returning a **ConsumerAdmin** object for adding consumers, an operation returning a **SupplierAdmin** object for adding suppliers, and an operation for destroying the channel.

The operations for adding consumers return *proxy suppliers*. A proxy supplier is similar to a normal supplier (in fact, it inherits the interface of a supplier), but includes an additional method for connecting a consumer to the proxy supplier.

The operations for adding suppliers return *proxy consumers*. A proxy consumer is similar to a normal consumer (in fact, it inherits the interface of a consumer), but includes an additional method for connecting a supplier to the proxy consumer.

Registration of a producer or consumer is a two step process. An event-generating application first obtains a proxy consumer from a channel, then “connects” to the proxy consumer by providing it with a supplier. Similarly, an event-receiving application first obtains a proxy supplier from a channel, then “connects” to the proxy supplier by providing it with a consumer.

The reason for the two-step registration process is to support composing event channels by an external agent. Such an agent would compose two channels by obtaining a proxy supplier from one and a proxy consumer from the other, and passing each of them a reference to the other as part of their connect operation.

Proxies are in one of three states: *disconnected*, *connected* or *destroyed*. Figure 2-6 gives a state diagram for a proxy. The nodes of the diagram are the states and the edges are labelled with the operations that change the state of the proxy. **Push/pull** operations are only valid in the *connected* state.




Figure 2-6 State diagram of a proxy.

2.3 The *CosEventChannelAdmin* Module

The **CosEventChannelAdmin** module defines the interfaces for making connections between suppliers and consumers. The **CosEventChannelAdmin** module is defined below.

```

#include "CosEventComm.idl"

module CosEventChannelAdmin {

    exception AlreadyConnected {};
    exception TypeError {};

    interface ProxyPushConsumer: CosEventComm::PushConsumer {
        void connect_push_supplier(
            in CosEventComm::PushSupplier push_supplier)
            raises(AlreadyConnected);
    };

    interface ProxyPullSupplier: CosEventComm::PullSupplier {
        void connect_pull_consumer(
            in CosEventComm::PullConsumer pull_consumer)
            raises(AlreadyConnected);
    };

    interface ProxyPullConsumer: CosEventComm::PullConsumer {
        void connect_pull_supplier(
            in CosEventComm::PullSupplier pull_supplier)
            raises(AlreadyConnected,TypeError);
    };

    interface ProxyPushSupplier: CosEventComm::PushSupplier {

```



```

        void connect_push_consumer(
            in CosEventComm::PushConsumer push_consumer)
            raises(AlreadyConnected, TypeError);
    };

    interface ConsumerAdmin {
        ProxyPushSupplier obtain_push_supplier();
        ProxyPullSupplier obtain_pull_supplier();
    };

    interface SupplierAdmin {
        ProxyPushConsumer obtain_push_consumer();
        ProxyPullConsumer obtain_pull_consumer();
    };

    interface EventChannel {
        ConsumerAdmin for_consumers();
        SupplierAdmin for_suppliers();
        void destroy();
    };
};

```

2.3.1 The EventChannel Interface

The **EventChannel** interface defines three administrative operations: adding consumers, adding suppliers, and destroying the channel.

```

interface EventChannel {
    ConsumerAdmin for_consumers();
    SupplierAdmin for_suppliers();
    void destroy();
};

```

Any object that possesses an object reference that supports the **EventChannel** interface can perform these operations:

- The **ConsumerAdmin** interface allows consumers to be connected to the event channel. The **for_consumers** operation returns an object reference that supports the **ConsumerAdmin** interface.
- The **SupplierAdmin** interface allows suppliers to be connected to the event channel. The **for_suppliers** operation returns an object reference that supports the **SupplierAdmin** interface.
- The **destroy** operation destroys the event channel.

Consumer administration and supplier administration are defined as separate objects so that the creator of the channel can control the addition of suppliers and consumers. For example, a creator might wish to be the sole supplier of event data but allow many consumers to be connected to the channel. In such a case, the creator would simply export the **ConsumerAdmin** object.

2.3.2 *The ConsumerAdmin Interface*

The **ConsumerAdmin** interface defines the first step for connecting consumers to the event channel; clients use it to obtain proxy suppliers.

```
interface ConsumerAdmin {  
    ProxyPushSupplier obtain_push_supplier();  
    ProxyPullSupplier obtain_pull_supplier();  
};
```

The **obtain_push_supplier** operation returns a **ProxyPushSupplier** object. The **ProxyPushSupplier** object is then used to connect a push-style consumer.

The **obtain_pull_supplier** operation returns a **ProxyPullSupplier** object. The **ProxyPullSupplier** object is then used to connect a pull-style consumer.

2.3.3 *The SupplierAdmin Interface*

The **SupplierAdmin** interface defines the first step for connecting suppliers to the event channel; clients use it to obtain proxy consumers.

```
interface SupplierAdmin {  
    ProxyPushConsumer obtain_push_consumer();  
    ProxyPullConsumer obtain_pull_consumer();  
};
```

The **obtain_push_consumer** operation returns a **ProxyPushConsumer** object. The **ProxyPushConsumer** object is then used to connect a push-style supplier.

The **obtain_pull_consumer** operation returns a **ProxyPullConsumer** object. The **ProxyPullConsumer** object is then used to connect a pull-style supplier.

2.3.4 *The ProxyPushConsumer Interface*

The **ProxyPushConsumer** interface defines the second step for connecting push suppliers to the event channel.

```
interface ProxyPushConsumer: CosEventComm::PushConsumer {  
    void connect_push_supplier(  
        in CosEventComm::PushSupplier push_supplier)  
        raises(AlreadyConnected);  
};
```

A nil object reference may be passed to the **connect_push_supplier** operation; if so a channel cannot invoke the **disconnect_push_supplier** operation on the supplier; the supplier may be disconnected from the channel without being informed.

If the **ProxyPushConsumer** is already connected to a **PushSupplier**, then the **AlreadyConnected** exception is raised.

2.3.5 *The ProxyPullSupplier Interface*

The **ProxyPullSupplier** interface defines the second step for connecting pull consumers to the event channel.

```
interface ProxyPullSupplier: CosEventComm::PullSupplier {
    void connect_pull_consumer(
        in CosEventComm::PullConsumer pull_consumer)
        raises(AlreadyConnected);
};
```

A nil object reference may be passed to the **connect_pull_consumer** operation; if so a channel cannot invoke a **disconnect_pull_consumer** operation on the consumer; the consumer may be disconnected from the channel without being informed.

If the **ProxyPullSupplier** is already connected to a **PullConsumer**, then the **AlreadyConnected** exception is raised.

2.3.6 *The ProxyPullConsumer Interface*

The **ProxyPullConsumer** interface defines the second step for connecting pull suppliers to the event channel.

```
interface ProxyPullConsumer: CosEventComm::PullConsumer {
    void connect_pull_supplier(
        in CosEventComm::PullSupplier pull_supplier)
        raises(AlreadyConnected, TypeError);
};
```

Implementations should raise the CORBA standard **BAD_PARAM** exception if a nil object reference is passed to the **connect_pull_supplier** operation.

If the **ProxyPullConsumer** is already connected to a **PullSupplier**, then the **AlreadyConnected** exception is raised.

An implementation of a **ProxyPullConsumer** may put additional requirements on the interface supported by the pull supplier. If the pull supplier does not meet those requirements, the **ProxyPullConsumer** raises the **TypeError** exception. (See Section 2.5.2, “The TypedPullSupplier Interface,” on page 2-15 for an example.)

2.3.7 The ProxyPushSupplier Interface

The **ProxyPushSupplier** interface defines the second step for connecting push consumers to the event channel.

```
interface ProxyPushSupplier: CosEventComm::PushSupplier {
    void connect_push_consumer(
        in CosEventComm::PushConsumer push_consumer)
        raises(AlreadyConnected, TypeError);
};
```

Implementations should raise the CORBA standard **BAD_PARAM** exception if a nil object reference is passed to the **connect_push_consumer** operation.

If the **ProxyPushSupplier** is already connected to a **PushConsumer**, then the **AlreadyConnected** exception is raised.

An implementation of a **ProxyPushSupplier** may put additional requirements on the interface supported by the push consumer. If the push consumer does not meet those requirements, the **ProxyPushSupplier** raises the **TypeError** exception. (See Section 2.5.1, “The TypedPushConsumer Interface,” on page 2-14 for an example.)

2.4 Typed Event Communication

Section 1.7, “Generic Event Communication,” on page 1-7 discusses generic event communication using **push** and **pull** operations. The next few sections describe how event communication can be described in OMG IDL and how typed event channels can support such typed event communication.

2.4.1 Typed Push Model

In the typed push model, suppliers call operations on consumers using some mutually agreed interface *I*. The interface *I* is defined in IDL, and may contain any operations subject to the following restrictions:

- All parameters must be in parameters only.
- No return values are permitted

These are the same restrictions as CORBA imposes on *oneway* operations, and for similar reasons: event communication is unidirectional, and does not directly support responses. The operations can be declared *oneway*, but need not be.

To set up typed push-style communication, consumers and suppliers exchange **TypedPushConsumer** and **PushSupplier** object references. (Note that the supplier interface is the same as the untyped case.) The supplier then invokes the **get_typed_consumer** operation of the **TypedPushConsumer** interface, which returns an object reference supporting the typed interface, *I*, referred to as an *I-reference*. The particular interface, *I*, that the reference supports is dependent on the particular **TypedPushConsumer**, and must be mutually agreed by supplier and consumer. Once the supplier has obtained the *I-reference*, it can call operations in interface *I* on the consumer.

As in the case of the generic push-style, event communication can be broken by invoking a **disconnect_push_consumer** operation on the **TypedPushConsumer** interface or by invoking a **disconnect_push_supplier** operation on the **PushSupplier** interface. If the **PushSupplier** object reference is nil, the connection cannot be broken via the supplier.

Figure 2-7 illustrates typed push-style communication between supplier and consumer.




Figure 2-7 Typed Push-style Communication Between a Supplier and a Consumer

2.4.2 Typed Pull Model

In the typed pull model, consumers call operations on suppliers, requesting event information, using some mutually agreed interface $Pull<I>$ ². For every interface I having the properties described in section 2.4.1, an interface $Pull<I>$ is defined as follows:

- For every operation o in I , $Pull<I>$ contains two operations:
 - **pull_o**, with all **in** parameters changed to **out** parameters. When called, this operation will return with the event data in the *out* parameters. If no *o*-event is currently available, it will block.
 - **boolean try_o**, with all in parameters changed to out parameters. When called, this operation will check whether an *o*-event is currently available. If so, it will return *true*, with the event data in the *out* parameters. If not, it will return *false*, with the out parameters undefined

The interface **Pull<I>** is designed to allow pulling of exactly the same events that can be pushed using interface **I**.

To set up typed pull-style communication, consumers and suppliers exchange **PullConsumer** and **TypedPullSupplier** object references. (Note that the consumer interface is the same as the untyped case.) The consumer then invokes the **get_typed_supplier** operation of the **TypedPullSupplier**, which returns an object reference supporting the typed interface, **Pull<I>**, referred to as a *Pull<I>-reference*.

². $Pull<I>$ is used as notation for a computed interface from interface I . Thus, if I is an interface *DocumentEvents*, $Pull<I>$ is an interface *PullDocumentEvents*.

The particular interface, **Pull<I>**, that the reference supports is dependent on the particular **TypedPullSupplier**, and must be mutually agreed by supplier and consumer. Once the consumer has obtained the *Pull<I>-reference*, it can call operations in interface **Pull<I>** on the supplier.

Figure 2-8 illustrates typed pull-style communication between supplier and consumer.

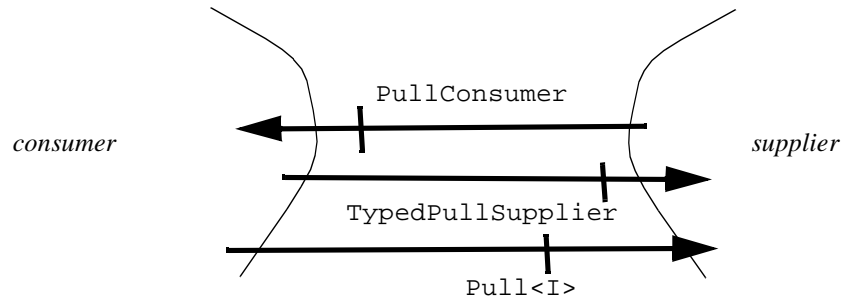


Figure 2-8 Typed Pull-style Communication Between a Supplier and a Consumer

2.5 The *CosTypedEventComm* Module

The typed communication styles shown in Figure 2-7 and Figure 2-8 are both supported by two new interfaces, **TypedPushConsumer** and **TypedPullSupplier** and two existing interfaces, **PushSupplier** and **PullConsumer**. The first two interfaces are defined in an OMG IDL module named **CosTypedEventComm**, as shown below. The last two are the same as for untyped event communication, and were defined in the **CosEventComm** module.

```
#include "CosEventComm.idl"
```

```
module CosTypedEventComm {
```

```
    interface TypedPushConsumer : CosEventComm::PushConsumer {
        Object get_typed_consumer();
    };
```

```
    interface TypedPullSupplier : CosEventComm::PullSupplier {
        Object get_typed_supplier();
    };
```

```
};
```

2.5.1 The *TypedPushConsumer* Interface

A typed push-style consumer supports the **TypedPushConsumer** interface both to receive event data in the generic manner, and to supply a specific typed interface through which to receive it in typed form.

```

interface TypedPushConsumer : CosEventComm::PushConsumer {
    Object get_typed_consumer();
};

```

The **TypedPushConsumer** can behave just like an untyped **PushConsumer**, described in Section 2.1.1, “The PushConsumer Interface,” on page 2-2. In addition, if the supplier wishes to communicate event data to the consumer in typed rather than generic form, it first invokes the **get_typed_consumer** operation. This returns an *I-reference* supporting an interface **I**. The particular interface **I** that the reference supports is dependent on the particular **TypedPushConsumer**. The return type of the operation is **Object**, because different **TypedPushConsumers** will return references of different types, so the actual type cannot be specified in a general definition. Once the supplier has obtained the **I-reference**, it can narrow it to **I**, and then call operations in interface **I** on the consumer. Mutual agreement about **I** is needed between the supplier and consumer. If they do not agree, the narrow operation will fail.

As noted above, a **TypedPushConsumer** must support the **push** operation, inherited from **CosEventComm::PushConsumer**. Implementing **push** fully is an unnecessary burden if the consumer is intended for typed use only. It is therefore permissible to implement a **TypedPushConsumer** with a null implementation of **push** that merely raises the standard CORBA exception **NO_IMPLEMENT**. Clearly, suppliers must know this and confine themselves to typed communication with such consumers.

2.5.2 The TypedPullSupplier Interface

A typed pull-style supplier supports the **TypedPullSupplier** interface both to allow consumers to pull event data in the generic manner, and to supply a specific typed interface through which they can pull it in typed form.

```

interface TypedPullSupplier : CosEventComm::PullSupplier {
    Object get_typed_supplier();
};

```

The **TypedPullSupplier** can behave just like an untyped **PullSupplier**, described in Section 2.1.3, “The PullSupplier Interface,” on page 2-3. In addition, if the consumer wishes to pull event data from the supplier in typed rather than generic form, it first invokes the **get_typed_supplier** operation. This returns a *Pull<I>-reference* supporting an interface **Pull<I>**. The particular interface, **Pull<I>**, that the reference supports is dependent on the particular **TypedPullSupplier**. The return type of the operation is **Object**, because different **TypedPullSuppliers** will return references of different types, so the actual type cannot be specified in a general definition. Once the consumer has obtained the *Pull<I>-reference*, it can narrow it to **Pull<I>**, and then call operations in interface **Pull<I>** on the supplier. Mutual agreement about **Pull<I>** is needed between the supplier and consumer. If they do not agree, the **narrow** operation will fail.

As noted above, a **TypedPullSupplier** must support the **pull** and **try_pull** operations, inherited from **CosEventComm::PullSupplier**. Implementing these operations fully is an unnecessary burden if the supplier is intended for typed use only. It is therefore

permissible to implement a **TypedPullSupplier** with null implementations of **pull** and **try_pull** that merely raise the standard CORBA exception **NO_IMPLEMENT**. Clearly, consumers must know this and confine themselves to typed communication with such suppliers.

2.6 *Typed Event Channels*

Typed event channels are analogous to generic event channels, but they support both typed and generic event communication. These forms can be mixed at will. A single channel can handle events supplied and consumed in any combination of the forms defined earlier (push/pull, generic/typed). An event supplied in typed form can be consumed in generic form, or vice versa.³

2.7 *The CosTypedEventChannelAdmin Module*

The **CosTypedEventChannelAdmin** module defines the interfaces for making connections between suppliers and consumers that use either generic or typed communication. Most of its interfaces are specializations of the corresponding interfaces in the **CosEventChannel** module.

```
#include "CosEventChannel.idl"
#include "CosTypedEventComm.idl"

module CosTypedEventChannelAdmin {

    exception InterfaceNotSupported {};
    exception NoSuchImplementation {};
    typedef string Key;

    interface TypedProxyPushConsumer :
        CosEventChannelAdmin::ProxyPushConsumer,
        CosTypedEventComm::TypedPushConsumer {};

    interface TypedProxyPullSupplier :
        CosEventChannelAdmin::ProxyPullSupplier,
        CosTypedEventComm::TypedPullSupplier {};

    interface TypedSupplierAdmin :
        CosEventChannelAdmin::SupplierAdmin {
            TypedProxyPushConsumer obtain_typed_push_consumer(
                in Key supported_interface)
                raises(InterfaceNotSupported);
            ProxyPullConsumer obtain_typed_pull_consumer (
                in Key uses_interface)
```

3. Doing this does require an understanding on the part of the generic suppliers and consumers of how the channel packages parameters of typed calls when converting them to generic form. Details of this packaging are dependent on the implementation of the channel.


```

        raises(NoSuchImplementation);
};

interface TypedConsumerAdmin :
    CosEventChannelAdmin::ConsumerAdmin {
        TypedProxyPullSupplier obtain_typed_pull_supplier(
            in Key supported_interface)
            raises (InterfaceNotSupported);
        ProxyPushSupplier obtain_typed_push_supplier(
            in Key uses_interface)
            raises(NoSuchImplementation);
};

interface TypedEventChannel {
    TypedConsumerAdmin for_consumers();
    TypedSupplierAdmin for_suppliers();
    void destroy ();
};
};

```

2.7.1 The TypedEventChannel Interface

```

interface TypedEventChannel {
    TypedConsumerAdmin for_consumers();
    TypedSupplierAdmin for_suppliers();
    void destroy ();
};

```

This interface is analogous to **CosEventChannelAdmin::EventChannel**. However, it returns typed versions of the consumer and supplier administration interfaces, which are capable of providing proxies for either generic or typed communication.

2.7.2 The TypedConsumerAdmin Interface

The **TypedConsumerAdmin** interface defines the first step for connecting consumers to typed event channel; clients use it to obtain proxy suppliers.

```

interface TypedConsumerAdmin :
    CosEventChannelAdmin::ConsumerAdmin {
        TypedProxyPullSupplier obtain_typed_pull_supplier(
            in Key supported_interface)
            raises (InterfaceNotSupported);
        ProxyPushSupplier obtain_typed_push_supplier(
            in Key uses_interface)
            raises(NoSuchImplementation);
};

```

The **obtain_typed_pull_supplier** operation takes a **Key** parameter that identifies an interface, **Pull<I>**. The scope of the key is the typed event channel. It returns a **TypedProxyPullSupplier** for interface **Pull<I>**. The **TypedProxyPullSupplier** will allow an attached pull consumer to pull events either in generic form or using operations in interface **Pull<I>**. It is up to the implementation of **obtain_typed_pull_supplier** to create or find an appropriate **TypedProxyPullSupplier**. If it cannot, it raises the exception **InterfaceNotSupported**.

The **obtain_typed_push_supplier** operation takes a **Key** parameter that identifies an interface, **I**. The scope of the key is the typed event channel. It returns a **ProxyPushSupplier** that calls operations in interface **I**, rather than **push** operations. It is up to the implementation of **obtain_typed_push_supplier** to create or find an appropriate **ProxyPushSupplier**⁴. If it cannot, it raises the exception **NoSuchImplementation**.

Such a **ProxyPushSupplier** is guaranteed only to invoke operations defined in interface **I**. Any event on the channel that does not correspond to an operation defined in interface **I** is not passed on to the consumer. Such a **ProxyPushSupplier** is therefore an event filter based on type.

2.7.3 The *TypedSupplierAdmin* Interface

The **TypedSupplierAdmin** interface defines the first step for connecting suppliers to the typed event channel; clients use it to obtain proxy consumers.

```
interface TypedSupplierAdmin :
    CosEventChannelAdmin::SupplierAdmin {
        TypedProxyPushConsumer obtain_typed_push_consumer(
            in Key supported_interface)
            raises(InterfaceNotSupported);
        ProxyPullConsumer obtain_typed_pull_consumer (
            in Key uses_interface)
            raises(NoSuchImplementation);
    };
```

The **obtain_typed_push_consumer** operation takes a **Key** parameter that identifies an interface, **I**. The scope of the key is the typed event channel. It returns a **TypedProxyPushConsumer** for **I**. An attached supplier can provide events by using operations in interface **I**. It is up to the implementation of **obtain_typed_push_consumer** to create or find an appropriate **TypedProxyPushConsumer**. If it cannot, it raises the exception **InterfaceNotSupported**.

The **obtain_typed_pull_consumer** operation takes a **Key** parameter that identifies an interface, **Pull<I>**. The scope of the key is the typed event channel. It returns a **ProxyPullConsumer** that calls operations in interface **Pull<I>**, rather than **pull**

4.see Appendix A for implementation considerations.

operations. It is up to the implementation of **obtain_typed_pull_consumer** to create or find an appropriate **ProxyPullConsumer**. If it cannot, it raises the exception **NoSuchImplementation**.

Such a **ProxyPullConsumer** is guaranteed only to invoke operations defined in interface **Pull<I>**. Any event request that does not correspond to an operation defined in interface **Pull<I>** is not pulled from the supplier. Such a **ProxyPullConsumer** is therefore an event filter based on type.

2.7.4 The *TypedProxyPushConsumer* Interface

The **TypedProxyPushConsumer** interface defines the second step for connecting push suppliers to the typed event channel.

```
interface TypedProxyPushConsumer :
    CosEventChannelAdmin::ProxyPushConsumer,
    CosTypedEventComm::TypedPushConsumer { };
```

By inheriting from both **CosEventChannelAdmin::ProxyPushConsumer** and **CosTypedEventComm::TypedPushConsumer**, this interface supports:

- connection and disconnection of push suppliers, exactly as in the generic event channel,
- generic **push** operation, and
- obtaining the typed view, so that the supplier can use typed push communication. The reference returned by **get_typed_consumer** has the interface identified by the **Key** used when this **TypedProxyPushConsumer** was obtained. (See Section 2.7.3, “The TypedSupplierAdmin Interface,” on page 2-18.)

2.7.5 The *TypedProxyPullSupplier* Interface

The **TypedProxyPullSupplier** interface defines the second step for connecting pull consumers to the typed event channel.

```
interface TypedProxyPullSupplier :
    CosEventChannelAdmin::ProxyPullSupplier,
    CosTypedEventComm::TypedPullSupplier { };
```

By inheriting from both **CosEventChannelAdmin::ProxyPullSupplier** and **CosTypedEventComm::TypedPullSupplier**, this interface supports:

- Connection and disconnection of pull consumers, exactly as in the generic event channel,
- generic **pull** and **try_pull** operations, and
- obtaining the typed view, so that the consumer can use typed pull communication. The reference returned by **get_typed_supplier** supports the interface identified by the **Key** used when this **TypedProxyPullSupplier** was obtained. (See Section 2.7.2, “The TypedConsumerAdmin Interface,” on page 2-17.)

2.8 *Composing Event Channels and Filtering*

The event channel administration operations defined in Section 2.3, “The CosEventChannelAdmin Module,” on page 2-8 support the composition of event channels. That is, one event channel can consume events supplied by another. This architecture allows the implementation of an event channel that filters the events supplied by another.

Since the **ProxyPushSupplier** for interface **I** of a typed event channel only pushes events that correspond to **I**, it acts as a filter based on type. Similarly, the **ProxyPullConsumer** for interface **Pull<I>** of a typed event channel only pulls events that correspond to **Pull<I>**, it also acts as a filter based on type.

2.9 *Policies for Finding Event Channels*

The Event Service does not establish a policy for finding event channels. Finding a service is orthogonal to using the service. Higher levels of software (such as the desktop) can make policies for using the event channel. That is, higher layers will dictate when an event channel is created and how references to the event channel are obtained. By representing the event channel as an object, it has all of the properties that apply to objects, including support by finding mechanisms.

For example, when a user performs a drag-and-drop or cut-and-paste operation, an event channel could be created and identified to suppliers and consumers. Alternatively, the event channel could be named in a naming context, or it could be exported through an operation on an object.

Implementing Typed Event Channels

A.1 Introduction

Note – Implementation details do not form part of an OMG specification, and should not be standardized. On the other hand, it is not obvious that typed channels can be implemented without extensions to CORBA. This section indicates *one* strategy for implementing typed event channels. It is included to show that typed event channels can be implemented; it is not intended in any way to constrain implementations. Optimized implementations are certainly possible.

Figure A-9 demonstrates a possible implementation of a typed event channel. This appendix concentrates on push style communication. The implementation of pull-style communication is analogous.

The implementation interposes an *encoder* between typed-style suppliers and the channel and a *decoder* between the channel and typed-style consumers.

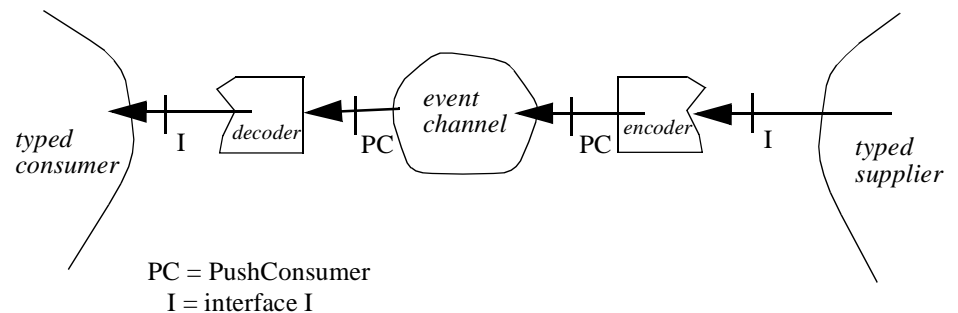


Figure A-9 A possible implementation of a typed event channel.

At the supplier end, an *encoder* converts operation calls to **push** calls.

At the consumer end, a *decoder* converts **push** calls back to operation calls.

The effect of such a communication is thus that the original operation is eventually called on the consumer, but the communication is routed via the channel. Of course, there can be multiple suppliers and multiple consumers on the same channel. Whenever one of the suppliers calls an operation, it is delivered by the channel to all consumers.

The encoder must package the operation identification and the parameters in a manner that the decoder can unpack them correctly.

Given the OMG IDL definition of an interface, **I**, an encoder generator could generate an implementation that supports the interface **I** and converts all calls on this interface to **push** calls on an event channel.

Similarly, it is possible to generate an I-decoder from the OMG IDL definition of **I**.

The typed event channel is responsible for finding, creating, or implementing the appropriate encoders. An appropriate encoder is found or created in response to the **obtain_typed_push_consumer** request on the typed event channel. The encoder is returned in response to the **get_typed_consumer** request.

Similarly, the typed event channel is responsible for finding, creating, or implementing the appropriate decoders. An appropriate decoder is found or created in response to the **connect_push_consumer** request on the typed event channel.

An Event Channel Use Example

B

This section illustrates an example use of the event channel, including the following:

- Creating an event channel
- Consumers and/or suppliers finding the channel
- Suppliers using the event channel
- In this example, the document object creates event channels and defines operations in its interface to allow consumers to be added.
- The **Document** interface defines two operations to return event channels:

```
interface Document {  
  
    ConsumerAdmin title_changed();  
  
    ConsumerAdmin new_section();  
  
    :  
  
};
```

The **title_changed** operation causes the document to generate an event when its title is changed; the **new_section** operation causes the document to generate an event when a new section is added. Both operations return **ConsumerAdmin** object references. This allows consumers to be added to the event channel.

-
- The **title_changed** implementation contains instance variables for using and administering the event channels.

```
/* Factory for creating event channels. */  
EventChannelFactoryRef ecf;  
  
/* For title changed event channel */  
EventChannelRef event_channel;  
  
ConsumerAdminRef consum_admin;  
SupplierAdminRef supplier_admin;  
  
ProxyPushConsumerRef proxy_push_consumer;  
PushSupplierRef doc_side_connection;
```

- At some point, the document implementation creates the event channel, gets supplier and consumer administrative references, and adds itself as a supplier¹.

```
event_channel = ecf->create_eventchannel(env);  
  
supplier_admin = event_channel->for_suppliers(env);  
consumer_admin = event_channel->for_consumers(env);  
proxy_push_consumer = supplier_admin->obtain_push_consumer(env);  
  
proxy_push_consumer->connect_push_supplier(env,  
doc_side_connection)
```

- The **title_changed** operation returns the **ConsumerAdmin** object reference.

```
return consumer_admin;
```

Clients of this operation can add consumers.

1. For readability, exception handling is omitted from these code fragments.

-
- When the title changes, the document implementation pushes the event to the channel.

```
proxy_push_consumer->push(env,data);
```

The document implementation similarly initializes, exports, and uses the event channel for reporting new sections.

- A**
 application object 1-2
- C**
 callback interface
 described x
 common facilities vi
 compound object ix
 concepts of viii
 connect 2-11
 Consolidated OMG IDL A-1, B-1
 consumer 1-2
 ConsumerAdmin interface 2-9, 2-10, 2-17
 for_consumers operation 2-9
 obtain_pull_supplier operation 2-10
 obtain_push_supplier operation 2-10
 CORBA viii
 contributors xi
 documentation set vii
 standard requests 1-1
 CORBA OMG IDL based Specification of the Trading
 Function A-1, B-1
 CosEventChannelAdmin module
 OMG IDL 2-8–2-9
 CosEventComm module
 OMG IDL 2-1–??
 CosTypedEventComm module
 OMG IDL 2-14
- E**
 event channel ix, x, 1-5, 2-6
 adding consumers 2-9
 adding consumers to 2-10
 adding consumers to typed 2-17
 adding pull consumer to typed 2-19
 adding pull consumers to 2-11
 adding pull suppliers to 2-11
 adding push consumers to 2-12
 adding push suppliers to 2-10
 adding push suppliers to typed 2-19
 adding suppliers 2-9
 adding suppliers to 2-10
 adding suppliers to typed 2-18
 and CORBA requests 2-4
 decoders A-2
 defined 1-3, 2-3
 encoders A-2
 filtering 2-20
 implementing typed A-1–A-2
 sample use B-1–B-3
 event communication
 mixed 2-5
 multiple 2-5
 pull model 1-2, 1-7, 2-4
 push model 1-2, 1-7, 2-4
 typed pull model 2-13
 typed push model 2-12
 event consumer 1-2, 1-7, 2-3
 proxy 2-7
 event service
 and CORBA scoping 1-5
 overview 1-1
 event supplier 1-2, 1-7, 2-3
 proxy 2-7
 EventChannel interface ix, 2-6, 2-7, 2-9
 exception 2-18, 2-19
 exceptions
 described xi
- G**
 global identifier x
- I**
 interface inheritance.see subtyping
- O**
 Object Management Group v
 address of viii
 object model vii
 object request broker vi, vii
 object service
 context vi
 specification defined vii
 OMG IDL vii, ix
- P**
 property list 1-2
 ProxyPullConsumer interface 2-11
 connect_pull_supplier operation 2-11
 ProxyPullSupplier 2-11
 ProxyPullSupplier interface 1-3, 2-11
 connect_pull_consumer operation 2-11
 ProxyPushConsumer interface 1-3, 2-10
 connect_push_supplier operation 2-11
 disconnect_push_supplier operation 2-11
 ProxyPushSupplier interface 2-12
 connect_push_consumer operation 2-12
 PullConsumer interface 1-3, 2-3, 2-13
 disconnect_pull_consumer operation 1-7
 PullSupplier interface ix, 1-7, 2-3
 disconnect_pull_supplier operation 1-7, 2-3
 pull operation 2-3
 try_pull operation 2-3
 PushConsumer interface ix, 1-7, 2-2
 disconnect_push_consumer operation 2-2
 push operation 2-2
 PushSupplier interface 1-3, 2-2
 disconnect_push_supplier operation 1-7, 2-2
- Q**
 quality of service ix, 1-4, 1-6, 2-5
- R**
 reference model vi
- S**
 subtyping viii, xi
 supplier 1-2
 SupplierAdmin interface 1-3, 2-9, 2-10
 for_suppliers operation 2-9
 obtain_pull_consumer operation 2-10
 obtain_push_consumer operation 2-10

Index

- T**
TypedConsumerAdmin interface
 obtain_typed_pull_supplier operation 2-18
 obtain_typed_push_supplier operation 2-18
TypedProxyPullSupplier interface 2-19
TypedProxyPushConsumer interface 2-19
TypedPullSupplier interface 2-13
- TypedPushConsumer interface 2-13
TypedSupplierAdmin interface 2-18
 obtain_typed_pull_consumer operation 2-18
 obtain_typed_push_consumer operation 2-18
- X**
X/Open vi