



Discover the Future of CORBA

# Micro Focus OpenFusion CORBA Services Version 5.0.2

Log Service Guide

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

<https://www.microfocus.com>

Copyright © Micro Focus 2009-2021. All rights reserved.

MICRO FOCUS, the Micro Focus logo, and Micro Focus product names are trademarks or registered trademarks of Micro Focus Development Limited or its subsidiaries or affiliated companies in the United States, United Kingdom, and other countries. All other marks are the property of their respective owners.

2021-04-19

# Contents

<b>Preface .....</b>	<b>vii</b>
About the Log Service Guide.....	vii
Intended Audience .....	vii
Organisation .....	vii
Conventions.....	vii
Contacting Micro Focus .....	viii
Further Information and Product Support.....	viii
Information We Need .....	ix
Contact information .....	ix

## Part I Log Service

<b>Description .....</b>	<b>3</b>
Overview .....	3
Benefits.....	3
OMG Standard Features.....	4
OpenFusion Features.....	4
Instrumentation .....	4
Dependencies on Other Services.....	5
Concepts and Architecture.....	5
Basic Architecture .....	5
Log Types.....	5
Log Factories .....	7
Log Networks.....	9
Components and Features.....	9
Log Records.....	9
Log Generated Events .....	10
Lifecycle Operations.....	10
Log Management.....	11
Quality of Service .....	14
Log Features Comparison .....	14
<b>Using Specific Features .....</b>	<b>15</b>
Introduction .....	15
Import statements .....	16
The Basics and the BasicLog.....	16
Creating a BasicLog .....	17
Import Statements .....	17
Initialisation.....	17
Log Creation and Configuration.....	17
Locating the Log.....	19
BasicLog Exceptions.....	19
Creating a Client for the BasicLog .....	20
Import Statements .....	20
Client Initialisation.....	20
Preparing the Data .....	21
Sending the Data .....	21
Querying the Log Records.....	23
Client Exceptions for the BasicLog.....	24
The NotifyLog and Event-style Events .....	25
Creating a NotifyLog .....	25
Import Statements .....	25
NotifyLog Initialisation.....	26

NotifyLog Creation and Configuration .....	26
Locating the NotifyLog .....	28
NotifyLog Exceptions .....	28
Creating a Supplier Client for the NotifyLog.....	29
Import Statements.....	30
Supplier Client Initialisation .....	30
The Client as an Event Supplier.....	31
Sending Events.....	33
Terminating the Supplier Client .....	33
Supplier Client Exceptions .....	34
Creating a Consumer Client for the NotifyLog .....	34
Import Statements.....	35
Consumer Client Initialisation .....	35
The Client as an Event Consumer .....	36
Retrieving Records .....	37
Cleanup .....	39
Exceptions .....	39
NotifyLog, Structured Events and More.....	39
Creating a NotifyLog .....	39
Import Statements.....	40
NotifyLog Initialisation .....	40
NotifyLog Creation and Configuration .....	40
Filtering .....	42
Registering with the Naming Service.....	45
NotifyLog Exceptions .....	45
Creating a Supplier Client for the NotifyLog.....	46
Import Statements.....	46
Supplier Client Initialisation .....	46
Structured Events in Brief .....	48
Creating and Sending the Events .....	49
Supplier Client Exceptions .....	51
Creating a Consumer Client for the NotifyLog .....	51
Import Statements.....	51
Consumer Client Initialisation .....	52
Setting a Proxy Filter .....	53
Receiving the Structured Events.....	54
Retrieving Log Records .....	55
Cleanup .....	56
Consumer Client Exceptions .....	56
Monitoring Log Generated Events .....	56
Import Statements.....	57
Monitor Client Initialisation .....	57
Receiving the Generated Events .....	58
Monitor Client Exceptions .....	58
<b>Supplemental Information .....</b>	<b>59</b>
Exceptions.....	59

## Part II Configuration and Management

<b>Log Service Configuration .....</b>	<b>63</b>
Overview.....	63
Common Properties .....	63
LogFactorySingleton Configuration .....	63
CORBA Properties.....	63

Persistence Properties .....	65
General Properties .....	66
ProcessSingleton Configuration .....	68

<b>Log Service Manager .....</b>	<b>71</b>
Overview .....	71
Using the Log Service Manager .....	71
Log Object Settings.....	72
Log.....	72
Alarms .....	73
Scheduling .....	73
Miscellaneous.....	73
Channel .....	74
Events .....	75
Create a New Log .....	75
Copy Logs .....	76
Destroy Logs.....	76
Browser .....	76
Settings .....	77
Operations.....	77

## Part III Appendices

<b>Using the Naming Service to Locate a Log .....</b>	<b>81</b>
Creating a Name Binding .....	81
Obtaining the Object with Resolve .....	82
<b>Using with JacORB .....</b>	<b>83</b>
<b>Index.....</b>	<b>85</b>



# Preface

## About the Log Service Guide

The **Log Service Guide** is included with the OpenFusion Log Service' *Documentation Set*. The **Log Service Guide** explains how to use the OpenFusion Log Service.

The **Log Service Guide** is intended to be used with the **System Guide**, **Notification Service Guide** and other OpenFusion Log Service documents included with the product distribution; refer to the **Product Guide** for a complete list of OpenFusion documents.

## Intended Audience

The **Log Service Guide** is intended to be used by users and developers who wish to integrate the OpenFusion Log Service into products which comply with OMG or J2EE standards for object services. Readers who use this guide should have a good understanding of the relevant programming languages (such as Java, IDL) and of the relevant underlying technologies (J2EE, CORBA).

## Organisation

The **Log Service Guide** is organised into two main sections. The first section describes the OpenFusion Log Service. This section provides:

- a high level description and list of main features
- explanation of the architecture and concepts
- how to use specific features
- detailed explanations of the main interfaces and how to use them
- other information which is needed to use the service

The second section, "[Configuration and Management](#)", gives information on configuring and managing the Log Service using the OpenFusion Administration Manager, together with detailed descriptions of properties specific to the Log Service. This section should be read in conjunction with the **System Guide**.

## Conventions

The conventions listed below are used to guide and assist the reader in understanding the **Log Service Guide**.



Item of special significance or where caution needs to be taken.



Item contains helpful hint or special information.



Information applies to Windows systems only.



Information applies to Unix based systems (e.g. Solaris) only.



C language specific



C++ language specific



Java language specific

Hypertext links are shown as [blue](#).

Items shown as cross-references, such as "[Contact information](#)", act as hypertext links; click on the reference to go to the item.

```
% Commands or input which the user enters on the
command line of their computer terminal
```

Courier fonts indicate programming code and file names.

Extended code fragments are shown in shaded boxes

```
:
NameComponent newName[] = new NameComponent[1];
// set id field to "example" and kind field to an empty string
newName[0] = new NameComponent ("example", "");
```

*Italics* and ***Italic Bold*** indicate new terms or emphasise an item.

**Bold** indicates user related actions, e.g. **File | Save** from a menu.

Steps in a task are numbered:

1 One of several steps required to complete a task.

## Contacting Micro Focus

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

### Further Information and Product Support

Additional technical information or advice is available from several sources.

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

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

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

#### Note:

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

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

Also, visit:

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



## Information We Need

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

- The name and version number of all products that you think might be causing a problem.
- Your computer make and model.
- Your operating system version number and details of any networking software you are using.
- The amount of memory in your computer.
- The relevant page reference or section in the documentation.
- Your serial number. You can find this by either logging into your order via the Electronic Product Distribution email or via the invoice with the order.

## Contact information

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

Additional technical information or advice is available from several sources.

The product support pages contain considerable additional information, including the *Product Updates* section of the Micro Focus SupportLine Web site, where you can download fixes and documentation updates. To connect, enter <http://www.microfocus.com> in your browser to go to the Micro Focus home page, then click *Support*.

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

You may want to check in particular:

- <https://supportline.microfocus.com/productdoc.aspx>. (documentation updates and PDFs)

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

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



# Part I

## Log Service

### **In this part**

This part contains the following:

<a href="#">Description</a>	<a href="#">page 3</a>
<a href="#">Using Specific Features</a>	<a href="#">page 15</a>
<a href="#">Supplemental Information</a>	<a href="#">page 59</a>



# Description

## Overview

The *OpenFusion Log Service* is one of a range of services and interfaces included with the *OpenFusion CORBA Services* product.

The service can be used stand-alone or with other OpenFusion Log Service' interfaces and services.

*OpenFusion Log Service* is standards based, meaning it is fully compliant with recognised industry standards and specifications, and supports portability and interoperability, and conforms with the OMG's *Telecom Log Service Specification* (which is, itself, an extension of the telecommunication industry's *ITU-T X.735* specification for logging events).

The *Telecom Log Service* can log and query any type of Notification or Event Service-based event, including Event Style, Structured, and Typed events. The Telecom Log Service can also log data received from clients which have no knowledge of events.

The OpenFusion Log Service conforms to the OMG's Telecom Log Service Specification (which is, itself, an extension of the telecommunication industry's *ITU-T X.735* specification for logging events). The OpenFusion Log Service supports the Basic and Notify Log functionality of the OMG specification.

*i*

This version of the OpenFusion Log Service is a subset of the Telecom Log Service specification and *only* supports Basic and Notify Logs using *push* mode.

## Benefits

Some of the benefits of using the Log Service include:

- The ability to store and selectively query events and *Any* data
- The performance, scalability and control benefits of an event channel architecture, including decoupled asynchronous and synchronous event transmission
- Ease of maintenance when adding or removing event suppliers and consumers
- Ease of integration and use with the Notification Service
- Efficient use of network bandwidth between the suppliers and consumers

The OpenFusion Log Service is widely used in the telecommunications, finance, transport, travel and energy industries for applications which require logging of client generated events, such as faults, errors, or alarms. The Log Service can be used for storing and retrieving a variety of information in a central location, such as:

- Audit trails of alarms for Operational Support Systems (OSS)
- Audit trails of diagnostics
- Configuration and management information

## OMG Standard Features

This release of the OpenFusion Log Service includes the standard OMG features, such as:

- The ability to selectively store and query any Notification Service-based events
- The ability to store data, sent as *Anys*, from clients which have no knowledge of events
- Decoupling event transmission from suppliers to consumers by using *event channels* and *proxies*. The events may be *unstructured* (simple data), *structured* (containing details about the event)
- Push event communication model for basic events
- Leverage of the Notification Service's Quality of Service (QoS) support
- Support for multiple suppliers and consumers
- Provision of filters and the *Extended Trader Constraint Language* for controlling event transmission to and from clients using the service and for controlling which events are saved to the persistent store
- The ability to forward events to consumers (such as applications or other *log objects*), which can be used to create log networks with *log-and-forward* ability
- The ability to report (using self generated events) when the state or condition of a log changes

## OpenFusion Features

The OpenFusion Log Service provides many enhancements over the standard OMG specification. These enhancements include:

- Provision of external graphical user interfaces, as part of the OpenFusion Graphical Tools, for run time administration of the service
- Rich administrative interface
- Support for persisting events to commercial databases
- Provisions for improved performance and scalability, such as:
  - Multi-threading
  - Provision of persistence for events, channels and connections to selected commercial databases through the use of optimized stored procedures
  - Automatic service activation on demand
- Provision of service monitoring functions, *instrumentation* (described below)

## Instrumentation

OpenFusion provides both general and service-specific instrumentation features which can be used for system monitoring, which in turn aids in problem identification, performance tuning, etc. OpenFusion Instrumentation consists of a set of properties that can be monitored either using the Administration Manager (one of the OpenFusion Graphical Tools) or remotely using SNMP.

In addition to properties that are read-only at run time, OpenFusion provides some properties that can be set and reset at run time as required,

such as when a particular threshold value is reached or a time period has elapsed. Note that there is virtually no performance overhead involved in using any of the OpenFusion Instrumentation features.

## Dependencies on Other Services

The Log Service needs the OpenFusion Notification Service to be installed.

The Log Service examples included in this guide use the OpenFusion Naming Service. Although the Log Service itself does not require this service, the Naming Service is highly recommended.

# Concepts and Architecture

There are numerous situations where it is necessary or desirable to be able to dynamically store information about a running system and its components: it may not be sufficient for a system to simply respond to events which occur. It may also be necessary to have a record of those events. The Log Service enables records of events to be stored which can later be retrieved for analysis.

## Basic Architecture

The main component of the Telecom Log Service is the *log* or *log object*. Log objects store events and data as log records. Logs can transmit events from suppliers to consumers, like event or notification channels do. Logs can also report their own state or condition through the generation and transmission of their own events.

There are different types of log. All but one of these types can forward events to consumer clients and generate their own events. Further, logs allow clients, which have no knowledge of events or notifications, to store *Any*-based data directly to the log's persistent store.

Logs can also:

- Be linked together as networks, similar to the federation of event channels
- Support the *copy* and *destroy* Lifecycle operations

Logs are created with *log factories*. Log factories also perform other functions, such as transmitting log generated events to clients (when required).

## Log Types

The Telecom Log Service Specification defines five log types. The OpenFusion implements two of these, the *BasicLog* and *NotifyLog*. All log types are inherited from an abstract `Log` interface. Four of the types also inherit from one of four event channel types (Event, Typed Event, Notify, or Typed Notify).

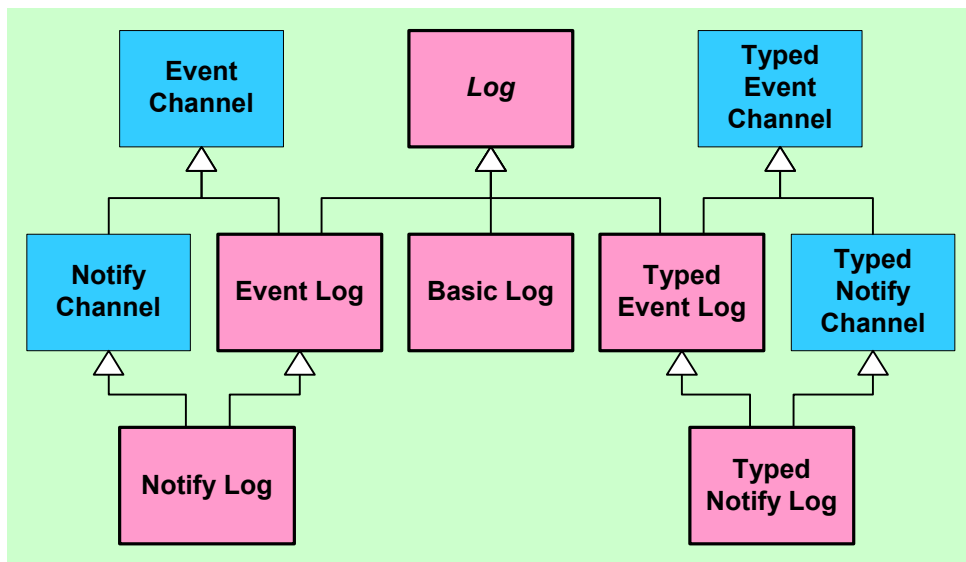
*i* This release of the OpenFusion Log Service only supports *Basic* and *Notify Logs* (described below): descriptions of the other log types as supported by the full Telecom Log Service Specification are included for information and general understanding.

The five log types are:

- 1 *BasicLog* - allows *event unaware* clients to store *Any*s directly to the log's persistent store. BasicLogs do not support event forwarding nor can

they generate events which report the log's state. BasicLogs only inherit from the `Log` interface.

- 2 *EventLog* - receives, saves to its persistent store, and forwards untyped Event Style events. EventLog can generate its own events which report the log's state. EventLog inherits from both `Log` and `CosEventChannelAdmin::EventChannel`.
- 3 *TypedEventLog* - receives, saves to its persistent store, and forwards typed Event Style events (Anys). TypedEventLog can generate its own events which report the log's state. TypedEventLog inherits from both `Log` and `CosTypedEventChannelAdmin::TypedEventChannel`.
- 4 *NotifyLog* - receives, saves to its persistent store, and forwards untyped events (including structured events). NotifyLog can filter events and generate its own events which report the log's state. NotifyLog inherits from both `EventLog` and `CosNotifyChannelAdmin::EventChannel`.
- 5 *TypedNotifyLog* - receives, saves to its persistent store, and forwards typed events. TypedNotifyLog can generate its own events which report the log's state and filter events. TypedNotifyLog inherits from both `TypedEventLog` and `CosTypedNotifyChannelAdmin::TypedEventChannel`.



**Figure 1 Log Types and Inheritance**

All logs, except for the `BasicLog`, inherit the architectural components of event channels. Figure 2 and Figure 3 show the basic architectures for event channel-based and notification channel-based logs, respectively.

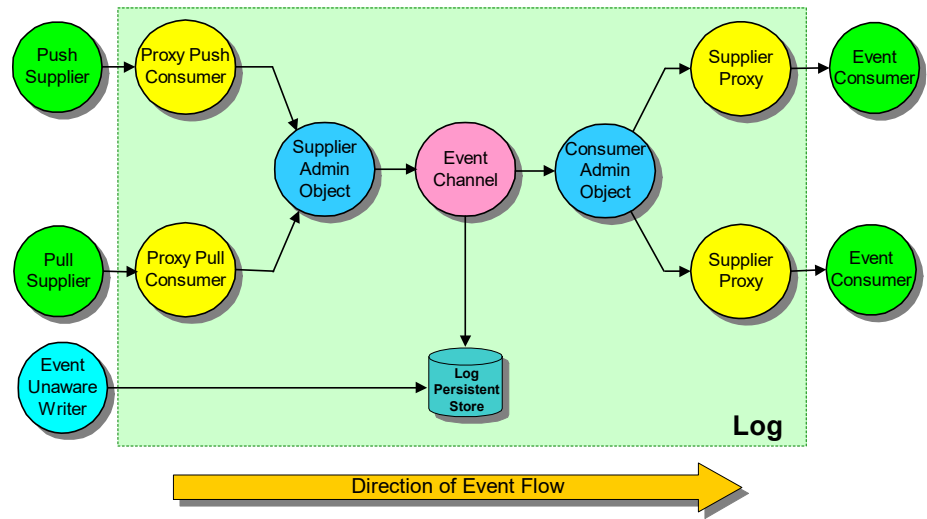
*i*

The `Log` inheritance hierarchy, shown in Figure 1, allows clients to access a log directly or through one of their underlying services, namely Event Service, Typed Event Service, Notification Service, and Typed Notification Service.

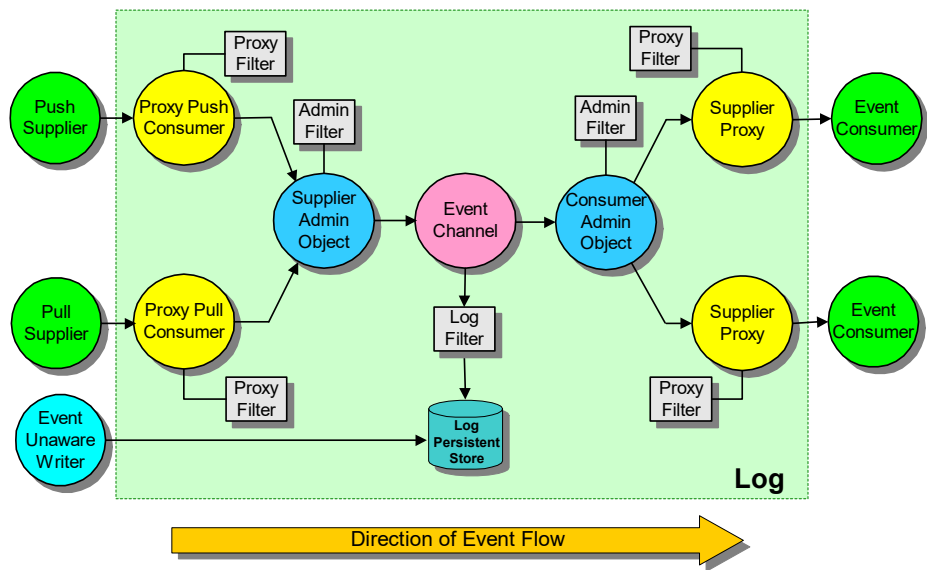
Further, a client can 'widen' a log to any of its bases classes. For example, a `NotifyLog` could be widened to a `Log` for use with event unaware applications.

A summary of the different features for each log type is listed in Table 2, [Features of Different Log Types](#), on page 14.





**Figure 2 Basic Architecture of Event Channel-based Log Types**



**Figure 3 Basic Architecture of Notify Channel-based Log Types**

## Log Factories

Logs are created by log factories. Log factories perform several functions, including:

- Log creation
- Performing as a log collection manager
- Generating log events
- Transmitting log events to subscribed consumers

Log factory interfaces are inherited from the *LogMgr* interface plus either a *ConsumerAdmin* or *NotifyConsumerAdmin* interface (with the exception of the *BasicLogFactory*) as shown in [Figure 4](#).

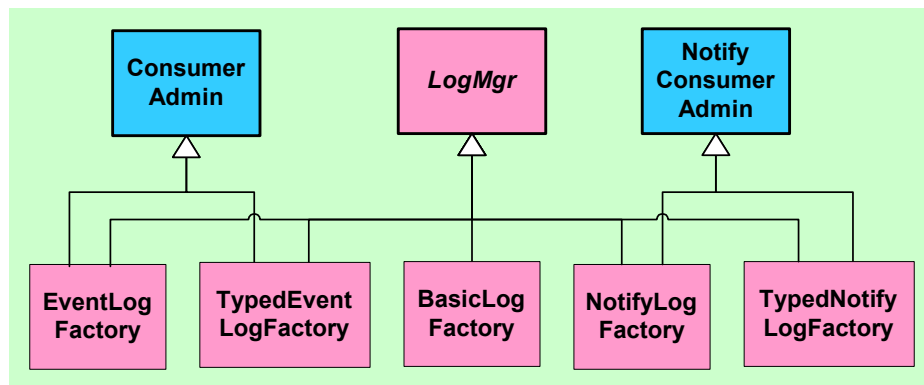
There are five different log factory types; each type corresponds to the log type that it creates:

- 1 *BasicLogFactory* - simply creates BasicLogs: they cannot emit nor forward log events since they only inherit from the *LogMgr* interface.
- 2 *EventLogFactory* - creates EventLogs plus emits log events as untyped events. *EventLogFactory* inherits from both *LogMgr* and *CosEventChannelAdmin::ConsumerAdmin*.
- 3 *TypedEventLogFactory* - creates TypedEventLogs plus emits log events as untyped events. *TypedEventLogFactory* inherits from both *LogMgr* and *CosEventChannelAdmin::ConsumerAdmin*.
- 4 *NotifyLogFactory* - creates NotifyLogs plus emits log events as untyped events, which can be filtered. *NotifyLogFactory* inherits from both *LogMgr* and *CosNotifyChannelAdmin::ConsumerAdmin*.
- 5 *TypedNotifyLogFactory* - creates TypedNotifyLogs plus emits log events as untyped events, which can be filtered. *TypedNotifyLogFactory* inherits from both *LogMgr* and *CosNotifyChannelAdmin::ConsumerAdmin*.

The *LogMgr* interface is an abstract interface which provides the following operations common to all log factory interfaces:

- Listing logs that have been created or copied
- Listing logs by their id
- Looking up logs by their id and returning a reference to the log

The *LogMgr* also provides log factories with log collection manager functionality.



**Figure 4 Log Factory Types and Inheritance**

Log factories which inherit from the *ConsumerAdmin* or *NotifyConsumerAdmin* interfaces enable clients to receive *log generated events*. Log generated events contain information about the state or condition of a log. These events are transmitted by the factory that created the log. A log factory transmits log generated events via its *Consumer Admin* object for all of the logs it has created, as shown in [Figure 5](#) below.

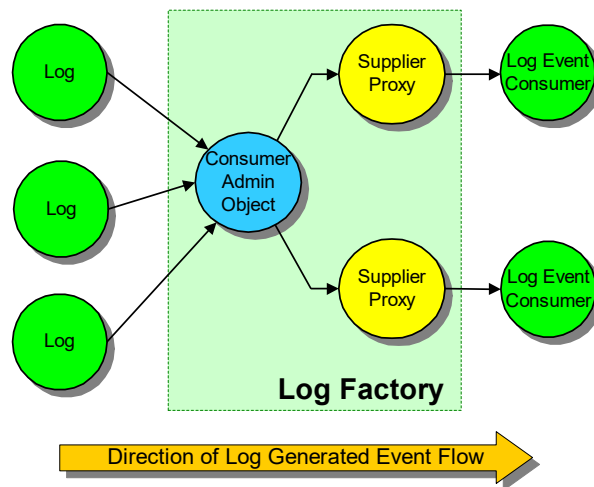


Figure 5 Log Generated Event Transmission

## Log Networks

Logs can be formed into networks by connecting one log's consumer side (*ConsumerAdmin* interface) to the other logs' supplier side (*SupplierAdmin* interface), then forwarding events from each log to the others.

## Components and Features

### Log Records

Events or data sent to a log is stored as a log record in the log's persistent store. Events travel to the log's persistent store via the log's event channel; data sent from event-unaware clients is written directly to the persistent store.

There are two types of log record, *LogRecord* and *TypedLogRecord*. *LogRecord* is supported by all log interfaces. *TypedLogRecord* is only supported by the typed log interfaces. A typed log can be stored as a (non-typed) *LogRecord*, subject to the operations that a client uses to retrieve it.

A *LogRecord* consists of:

- *id* - a unique id number for the individual record
- *time* - a time stamp of when the event was logged (i.e. stored)
- *attr\_list* (*attribute list*) - a user-defined list of attributes, which are not part of the event received by the log, containing log record related information; this information can be queried or modified
- *info* - the event or data stored as an *Any*, noting that structured or typed events can be wrapped in a CORBA *Any* before the event is stored

A *TypedLogRecord* consists of:

- *id* - a unique id number for the individual record
- *time* - a time stamp of when the event was logged (i.e. stored)
- *attr\_list* (*attribute list*) - a user-defined list of attributes containing log record related information

- *interface\_id* - repository id of the interface that sent the typed event
- *operation\_name* - name of the operation that emitted the typed event
- *arg\_list* - argument list that contains the event data

## Features

Log records can be queried, retrieved, or deleted according to each record's id, log time, attributes, or event contents. Details of working with log records is described in ["Using Specific Features"](#).

## Log Generated Events

Logs which inherit from an event channel are able to generate their own events (referred to as *log events*) which can be transmitted to consumer clients. Log generated events are not transmitted directly to consumers, but through the log factory that created the log. Log factories are also able to generate events.

Logs or log factories generate the following events:

- *ObjectCreation* - when a log is created
- *ObjectDeletion* - when a log is deleted
- *ThresholdAlarm* - the log's persistent store approaches its defined capacity or *wrapping* condition (*wrapping* is when the oldest records in the log are deleted in order to free space for new records)
- *AttributeValueChange* - when certain log attributes change (listed in [Table 1 on page 13](#))
- *StateChange* - a log's administrative or operational state changes (see ["Log Management"](#) below)
- *ProcessingErrorAlarm* - generated by the log factory when one of its logs generates an error

Consumers can receive log generated events by subscribing to the log factory.

## Lifecycle Operations

Log supports the *copy* and *destroy* Lifecycle operations.

There are two versions of the copy operation. The first, *copy*, creates an empty log, generates an id, and initializes the new logs attributes to the same values as the log on which it was invoked. The second version, *copy\_with\_id*, likewise creates a empty log with its attributes initialized to the same values as the log on which it was invoked, but gives the log an id supplied as a parameter - provided the id does not already exist.

The *destroy()* method disconnects any consumers or suppliers connected to the log's event or notify channel, then destroys the logs and its log records. *destroy()* is inherited from *CosEventChannelAdmin::EventChannel*.



However, if an event or notification-based client terminates, then *it* must disconnect itself from the log (by destroying the relevant proxy object), since the log has no means of knowing that the client no longer exists. Accordingly, the client should call its associated proxy's disconnect method. For example, if the client is a push supplier connected to a *ProxyPushConsumer* (suppliers connect to consumer proxies, consumers connect to supplier proxies), then the *disconnect\_push\_consumer()*

method for its `ProxyPushConsumer` object should be called prior to termination.

*i*

This need to remove proxy objects when a client terminates also applies to filter and other, similar objects - if a client or application creates an object, then it is responsible for removing the object, otherwise these "hanging" or unattached objects will remain in the system and will be a source of memory leaks.

## Log Management

Logs have management operations for monitoring and setting their behaviour. Most of these operations are common to all logs; the event- and notify-based logs have additional operations for management of their particular features (listed under "[Additional Management Operations](#)").

Several of the management operations cause *AttributeValueChange* events to be generated. The operations or changes which generate *AttributeValueChange* events are listed in [Table 1](#) on [page 13](#).

### Common Management Operations

The management capabilities which are common to all logs fall into two groups: those which can only obtain them and those which can change and obtain settings, values, and attributes.

#### Read-Only Settings

The following log settings, values, and attributes can only be obtained or read:

- *id* - the log's id
- *Factory id* - the id of the log factory that created the log
- *Operational state* - whether a log is operational or not (due to a run time problem, for example)

#### Read and Write Settings

The following log settings, values and attributes can be set as well as read:

- *Administrative state* - turns logging on or off and termed *unlocked* and *locked*:
  - When *unlocked*, log records are allowed to be created, retrieved, and deleted subject to the values of other states or settings
  - When *locked*, log records are allowed to be retrieved, and deleted, but new records are not allowed to be created
  - The administrative state does not affect the forwarding of events
- *Maximum log size* - the maximum size, in bytes, that a log is allowed to store:
  - A value of zero (0) sets the size as unlimited
- *Log full actions* - what actions should be performed when the maximum log size (the size of its persistent store) has been reached. Possible actions are:
  - *Wrap*, whereby the oldest records in the log are deleted in order to free space for new records
  - *Halt*, whereby no more records will be logged and all incoming events are discarded

- *Log capacity alarm threshold* - the level of the log's storage capacity at which an alarm is triggered to warn clients:
  - The alarm threshold is set as a percentage of the maximum log size
  - The alarm threshold can be used to warn clients that the oldest records may be overwritten soon when the log full action is set to wrap; the alarm is triggered each time the log wraps
  - When the log full action is set to halt, the alarm can warn clients that the log will soon stop accepting events, provided:
    - The client is registered to receive a ThresholdAlarm event (with the log's factory) and
    - The alarm threshold has been set to a value lower than 100% of the max log size
- *Log duration* - the period of time period during which a log will store events, and where:
  - The administrative and operational states are, respectively, unlocked and enabled
  - If the duration's *stop* value is set to zero (0), then the duration is indefinite
  - The duration is sufficiently long to avoid a race condition occurring between the activation of the log and the receiving of events
  - The default log duration is for the lifetime of the log (i.e. from the time the log is created until it is destroyed)
  - Log duration does not affect the status of event forwarding
- *Log scheduling* - the fine-grained time intervals that a log will store records of events, and where:
  - The administrative and operational states are, respectively, unlocked and enabled
  - Time intervals can be specified according to minutes, hours, days, and weeks
  - The time intervals must be within the log duration time interval
  - The default is disabled, which means that the log can store records at any time
  - Log scheduling does not affect the status of event forwarding
- *Availability status* - whether the combination of various states and conditions allows or enables events to be logged and/or retrieved; a log is available for logging if all of the following states are true:
  - The log is not full (i.e. the capacity of log store has not been reached)
  - The administrative state is unlocked
  - The operational state is enabled
  - The current time is within the log duration time
  - The current time is within the log's scheduled time(s)
- *Log record lifetime* - the maximum length of time, in seconds, that a log record is retained:
  - This is used to remove old, stale records from the log's store
  - The default lifetime value is zero (0), which allows records to remain indefinitely and until the log is destroyed

- *Quality of service* - Quality of Service properties specific to logs (see "Quality of Service" on page 14); logs can also use the Notification Service's QoS properties provided by a log's inherited event or notify channel

### Additional Management Operations

Logs which inherit from the event or notify channels provide additional management capabilities.

EventLog and TypedEventLog provide management of log forwarding.

NotifyLog and TypedNotifyLog provide management of log forwarding plus filtering of events:

- Received from suppliers and forwarded to consumers
- Sent to the log record stores

These operations (for reading and setting) are briefly described below:

- *Log record filtering* - sets the filter to be used for filtering log records (as distinct from the filters used by the log's channel for filtering incoming and forwarded events):
  - Logs do not have a log record filter when they are created and therefore log all events (until a log record filter is set)
- *Forwarding state* - controls whether a log will forward incoming events from all suppliers to all consumers which are currently connected:
  - The default is set to *on* (forward all events)

### Attribute Change Events

The following table lists properties or settings which cause *AttributeValueChange* events to be generated.

**Table 1 Settings Generating AttributeValueChange Events**

Setting or Value	Condition Generating Event
log full action	when set
log size	when set
log duration	when set
log scheduling	when <i>week</i> mask is set
max record lifetime	when set
log-specific quality of service	when log QoS set (excludes QoS inherited from event or notify channels)
log capacity alarm threshold	when set
forwarding state	when set
log filter	when set

## Quality of Service

The Log Service provides Quality of Service (QoS) properties specifically for logs. In addition, logs which inherit from the Notify or TypedNotify channels can use the Notification Service's QoS framework (refer to the *OMG Notification Service Specification* or the **OpenFusion Notification Service Guide**).

The OMG specifies three `Log` specific QoS levels:

- *QoSNone* - no quality of service is promised
- *QoSFlush* - log records are made persistent:
  - The `QoSFlush` property does not affect the persistence setting for events which are *forwarded* by the log
- *QoSReliability* - all log records are guaranteed to be available



The OpenFusion implementation of the Log Service provides `QoSReliability` as the default (and only) `Log`-specific QoS level: this level guarantees that all log records sent to a log will be available and enables log records to be recovered if, for example, the system crashes.

Also, the OpenFusion `NotifyLog` inherits from the OpenFusion Notification Service, including all of its standard OMG and OpenFusion extended QoS properties.

## Log Features Comparison

The following table provides a summary of the different features for each log type.

**Table 2 Features of Different Log Types**

Log Type	write (in-bound)	store (log record)	forward (out-bound)	emit log events	QoS
BasicLog	write operation	no filtering	no	no	log QoS
EventLog	untyped events push-pull model no filtering	no filtering	untyped events push-pull model no filtering	yes no filtering	log QoS
TypedEvent Log	typed events push-pull model no filtering	no filtering	typed events push-pull model no filtering	yes no filtering	log QoS
NotifyLog	untyped events structured events. push-pull model event filtering	log filter	untyped events structured events. push-pull model filtering	yes event filtering	log QoS Notification QoS
TypedNotify Log	typed events push-pull model event filtering	log filter	typed events push-pull model filtering	yes event filtering	log QoS Notification QoS



# Using Specific Features

## Introduction

The main tasks which are normally performed when using logs include:

- Creation of the log for storing records of data or events sent by clients
- Creation of clients which supply the data or events
- Creation of clients which use the log's records

It may also be necessary to perform other tasks, such as:

- Creating a client which can receive events sent via the log
- Creating a client which can receive events generated by the log, itself
- Enabling the log to record only selected events or only for specific time periods
- Manage the log, in terms of its reliability, size, lifetime, etc.
- Perform necessary clean-up operations

These tasks depend on the particular capabilities and implementation requirements of the type of log used (either *BasicLog* or *NotifyLog*).

This section, "[Using Specific Features](#)", describes how the specific features and requirements for the *BasicLog* and *NotifyLog* types can be used to achieve the relevant tasks listed above. The section is organised into a sequence of topics which:

- Describes basic aspects common to both log types
- Highlights aspects which are specific to the simple *BasicLog*, then progressing to Notification-based logs
- Progressively introduces methods of performing specific tasks

Each topic uses examples to illustrate how tasks can be achieved. Additional examples, complete with source code and descriptions of how to compile and run them, are supplied separately as part of the OpenFusion product distribution.



### Note

- All of the example code used in this section requires that the OpenFusion Log and Naming Services are installed and running.
- There is little or no error-checking provided in the examples for the sake of clarity. Code to deal with exceptions has generally been omitted, again, for clarity and brevity, appreciating that exceptions must be properly caught and handled in a working system.
- Full details of the OpenFusion Log Service API is provided in the OpenFusion IDL Documentation.

## Import statements

The following packages are required to be imported into classes which use the log types listed below.



This list is not exhaustive: additional packages may also be required, depending on requirements and the specific features of each service which are used.

### BasicLog

Classes which create the log or clients which access the log:

*org.omg.DsLogAdmin.\**

### NotifyLog

Classes which create the log:

*org.omg.DsLogAdmin.\**

*org.omg.DsNotifyLogAdmin.\**

*org.omg.CosNotification.\**

Clients which access the log:

*org.omg.DsLogAdmin.\**

*org.omg.DsNotifyLogAdmin.\**

*org.omg.CosNotification.\**

*org.omg.CosNotifyComm.\**

*org.omg.CosNotifyChannelAdmin.\**

### OpenFusion Extensions

The following package is needed when using the OpenFusion Log Service extensions:

*com.prismt.cos.CosLogging.LogExtensions.\**

## The Basics and the BasicLog

This topic explains how to:

- 1 Create a BasicLog and perform rudimentary configuration of its record storage availability and behaviour.
- 2 Use the Naming Service to easily locate specific log instances by clients.
- 3 Create an event-unaware client which supplies the log with data.
- 4 Create a client which can query log records containing specific, desired information.

The `BasicLog` possesses features which are common to all logs. This log type is also relatively simple and straight-forward to use, since it inherits only from the Log interface. Consequently, it is a convenient log type to use for learning about the basic aspects of all log types.

An example application is employed here to demonstrate how to use a BasicLog. The application consists of two component programs: one creates a log, the other sends data to the log, then queries the log's records.

## Creating a BasicLog

The first example program creates, configures and instantiates a BasicLog as a CORBA object. The program also uses the Naming Service to associate or *bind* the log object to a *name object*: this name object is subsequently used by the client program to obtain a reference (i.e. IOR) to the log object, thereby enabling the client to communicate with the log.

### Import Statements

The example BasicLog creation program uses the services and associated APIs listed below. Applications or classes that use these services or APIs must include the import statements shown below.

#### OpenFusion Orb Adaptor

```
com.prismt.orb.ORBAdapter
```

#### BasicLog

```
org.omg.DsLogAdmin.*
```

#### Naming Service

```
org.omg.CosNaming.*
```

```
org.omg.CosNaming.NamingContextPackage.*
```

```
org.omg.CosNaming.NamingContextExt.*
```

```
org.omg.CosNaming.NamingContextExtPackage.*
```

### Initialisation

The program declares *orb* and *log* objects, plus an *org.omg.CORBA.IntHolder*, *logId*. The *logId* is used to hold the log's id (generated by the log factory which creates the log).

```
org.omg.CORBA.ORB orb = null;
BasicLog log = null;
org.omg.CORBA.IntHolder logId = null;
```

OpenFusion's ORB-vendor independent abstraction layer is used to initialise and manage the orb object: the *ORBAdapter.init* method initialises the ORB:

```
orb = ORBAdapter.init (args);
System.out.println ("got orb");
```

After the *orb* object has been initialised, the steps required for creating the BasicLog object can be performed.

### Log Creation and Configuration

A log is created by a log factory's *create()* method. The *create()* method requires some configuration information, including the record store's maximum size and what happens when the store is full (i.e. overwrite the oldest records or reject new ones).

Any configuration properties or settings not explicitly set when the log is created are set to default values. All of a log's configuration properties or settings can be changed after the log is created using various Log methods, such as *set\_interval()*. It is usually necessary to change one or more configuration properties, since many of the default settings disable potentially useful or required control mechanisms.

The following steps describe how to create and configure a BasicLog object.

### 1 Obtain a BasicLog factory.

The following example code obtains a CORBA object for a *BasicLogFactory* using the ORB's *resolve\_initial\_references()* method, then narrows the object to a *BasicLogFactory*. Code should be included to test if the factory was not created, for example when the Log Service is not running.

```
// Create a BasicLog using the BasicLogFactory
org.omg.CORBA.Object object =
    orb.resolve_initial_references ("BasicLogFactory");
System.out.println ("initial references BasicLog factory");
BasicLogFactory factory = BasicLogFactoryHelper.narrow (object);
System.out.println ("resolved BasicLog factory.");
if(factory == null)
{
    System.out.println ("BasicLog factory not found.");
    System.exit(1);
}
```

### 2 Create the log.

A BasicLog is created with the *BasicLogFactory.create()* method. In this example, the log, called *log*, will be created with a record store having a maximum size of 10,000 bytes and which will overwrite the oldest records when full (set using *wrap.value*). The record store size used here is for illustrative purposes only: the size used in a real environment should reflect the approximate size (in bytes) of the records to be stored multiplied by the number of records the log is required to hold.

The *logId* variable (an *IntHolder*) is used to hold the log's id value (assigned by the *create()* method).

```
// create a log which will remove oldest records when its
// size of 10,000 bytes is exceeded
logId = new org.omg.CORBA.IntHolder();
log = factory.create(wrap.value, 10000, logId);
```

### 3 Configure the log.

Administrative or other property settings can now be set, as required. It is recommended that the log's ability to record data or events be disabled whenever an administrative or similar property is set using *set\_administrative\_state()* with *AdministrativeState.locked*.

The example uses the log's *set\_interval()* method to set the length of time that the log will be allowed to store records. This method takes a *TimeInterval* parameter; *TimeInterval* has two values, *start* and *stop*, which define the beginning and end times of the interval. *TimeInterval* is based on *Gregorian* time: any time value used should be relative to the start of the Gregorian calendar, i.e. the 15th of October 1582. This time can easily be obtained using Java's *GregorianCalendar* and *Date* classes as shown in the example.

When all necessary properties have been set, then re-enable record storing using *set\_administrative\_state()* with *AdministrativeState.unlocked*.

```
// lock the log prior to setting its operational interval
log.set_administrative_state(AdministrativeState.locked);

// Set length of time for log to accept data
GregorianCalendar calendar = new GregorianCalendar();
Date beginning = calendar.getGregorianCalendarChange();
```

```

long basetime = - beginning.getTime();
long now = System.currentTimeMillis();

TimeInterval interval = new TimeInterval();
interval.start = now + basetime;
interval.start *= 10000; // convert to 100ns

// Stop logging from now plus duration
interval.stop = now + basetime + duration;
interval.stop *= 10000; // convert to 100ns
log.set_interval(interval);

// unlock the log after setting the operational interval
log.set_administrative_state(AdministrativeState.unlocked);

```

## Locating the Log

Any client or application which uses a CORBA object must be able to obtain a reference to it using its IOR. One method is to save its IOR to a file, then have the client read the file in order to obtain the IOR. However, from the viewpoint of a distributed system, this method is very crude and has obvious failings, such as “How does the client locate the file containing the IOR?”.

An alternate, much preferred method is to use the Naming Service to locate the object on behalf of the client. The Naming Service associates objects with a name: this association is called a *name binding* and can be used by clients to efficiently obtain a reference to the object. The only prerequisites for using name bindings are that the Naming Service must be running and the object must be bound to a name before clients try to use it.

This example binds the log object to a name using the Naming Service: clients can then use the name to locate the log. Describing how to use the Naming Service and create name bindings is outside the scope of this topic. For convenience however, a brief description is provided in the Appendix “[Using the Naming Service to Locate a Log](#)”. Refer to the ***OpenFusion Naming Service Guide*** for complete details on using the Naming Service.

```

// bind the log object to a name with Naming Service
org.omg.CORBA.Object namingService = null;
NamingContextExt rootContext = null;
NameComponent basicLog[];

namingService = orb.resolve_initial_references ("NameService");
rootContext = NamingContextExtHelper.narrow(namingService);
basicLog = new NameComponent[1];

basicLog[0] = new NameComponent ("BasicLog", "log");
rootContext.rebind(basicLog, log);

System.out.println("BasicLog created.");

```

The log object can now be located, from clients located anywhere in the system, using the context’s *resolve()* method (described below under “[Creating a Client for the BasicLog](#)”).

## BasicLog Exceptions

The example code shown above, which creates the BasicLog object and Naming Service name binding, must either throw or catch the exceptions listed below. Log exception details, for this and all other examples, are described in “[Exceptions](#)”.

### ORB

```
org.omg.CORBA.ORBPackage.InvalidName
```

### BasicLog (including Interval)

```
org.omg.DsLogAdmin.InvalidLogFullAction
org.omg.DsLogAdmin.InvalidTime
org.omg.DsLogAdmin.InvalidTimeInterval
```

### Naming Service

```
org.omg.CosNaming.NamingContextPackage.InvalidName
org.omg.CosNaming.NamingContextPackage.NotFound
org.omg.CosNaming.NamingContextPackage.CannotProceed
```

## Creating a Client for the BasicLog

The following client program repeatedly sends data to a BasicLog instance and then queries the log's records to selectively obtain the records of data previously sent. The client utilises separate threads, via J2SE's *Timer* and *TimerTask* classes, to send data and query the log.

### Import Statements

The program uses the services and associated APIs listed below. Applications or classes that use these services or APIs must include the import statements shown below.

#### OpenFusion Orb Adaptor

```
com.prismt.orb.ORBAdapter
```

#### BasicLog

```
org.omg.DsLogAdmin.*
```

#### Naming Service

```
org.omg.CosNaming.*
org.omg.CosNaming.NamingContextPackage.*
org.omg.CosNaming.NamingContextExt.*
org.omg.CosNaming.NamingContextExtPackage.*
```

#### J2SE'S Timer and TimerTask

```
java.util.*
java.util.TimerTask
```

### Client Initialisation

The program declares *orb* and *log* objects, plus an *org.omg.CORBA.Any* array, *data*, which will contain the data that will be sent to the log. Since the Naming Service will be used to retrieve the *log* object, variables required to retrieve the Naming Service instance's *context* and *name component* for the log are also declared.

As mentioned above, this example uses threads provided by the *Timer* and *TimerTask* classes. Two *Timer* objects, *sendTimer* and *receiveTimer*, are declared. These thread objects respectively send data and queries at set intervals for a fixed duration of time, as defined by the *interval* and *duration* variables. The *delay* variable sets a time delay, in milliseconds, for the commencement of the query thread. Please note that *Timer* and *TimerTask* are not the only approaches which can be used to send data or events to a log using threads.

```
org.omg.CORBA.ORB orb = null;
org.omg.CORBA.Object obj = null;

BasicLog log = null;
org.omg.CORBA.Any[] data;

NamingContextExt rootContext = null;
```

```

NameComponent basicLog[];

private Timer sendTimer;
private Timer receiveTimer;

private final long duration = 3 * 60 * 1000; // activity time in
millisecs
private final long interval = 2000; // 2 seconds
private final long delay = 2000;

```

The example client obtains references to and performs initialisation of the ORB, Naming Service instance, naming context, and name component in a similar manner as was performed by the previous program which created the log.

A reference to the log is obtained by passing a `NameComponent` object (containing the log's `id` and `kind` values) to the context's `resolve()` method, then narrowing the object to the `BasicLog` type (`resolve()` returns an *Object* type).

```

// obtain reference to the orb
orb = ORBAdapter.init(args);

// obtain reference to log using Naming Service
obj = orb.resolve_initial_references("NameService");
rootContext = NamingContextExtHelper.narrow(obj);
basicLog = new NameComponent[1];
basicLog[0] = new NameComponent("BasicLog", "log");
obj = rootContext.resolve(basicLog);
log = BasicLogHelper.narrow(obj);

```

## Preparing the Data

The `BasicLog` type can not accept events - it can only accept data which is sent in an array of type `Any`. The following code creates an `Any` array and initialises it with four different primitive types: `string`, `double`, `long` and `char`. If a single data element is to be sent, then an `Any` array with one element should be used.

```

// initialise sample data to be logged
data = new org.omg.CORBA.Any[4];

for(int i = 0; i < data.length; i++)
{
    data[i] = orb.create_any();
}
data[0].insert_string("Hello world!");
data[1].insert_double(42.424242);
data[2].insert_long(123456789);
data[3].insert_char('a');

```

## Sending the Data

The example code shown below:

- Checks that a log is available to accept records
- Writes data to the log

The client sends data to the log every two seconds for five minutes: the frequency and duration is controlled by the `Timer.schedule()` method in combination with the client's implementation of `TimerTask.run()`. The `run()` method is executed by `schedule()` according to the `delay` and `interval` values. `run()` has been implemented so that it exits when a timeout value, `stopTime`, is reached.

*i*

It should be remembered that the use of `Timer` and `TimerTask.run()` is a client specific approach to controlling *when* the log is accessed. (Details for

using *Timer* and *TimerTask* are provided in Oracle Corporation's J2SE API documentation and the *Java Tutorials*, available on Oracle's web site).

The following code is taken from the client's `TimerTask.run()` implementation.

```
// thread for sending data
class Sender extends TimerTask
{
    long stopTime = System.currentTimeMillis() + duration;

    public void run()
    {
        if((System.currentTimeMillis() < stopTime) &&
            (log.get_availability_status().off_duty == false))
        {
            try
            {
                log.write_records(data);
                System.out.println ("\nWrote " + data.length + " records");
            }
            catch (org.omg.DsLogAdmin.LogFull ex)
            {
                System.err.println("Log is full.");
                System.exit(1);
            }
            catch (org.omg.DsLogAdmin.LogLocked ex)
            {
                System.err.println("Log is locked.");
                System.exit(1);
            }
            catch (org.omg.DsLogAdmin.LogDisabled ex)
            {
                System.err.println("Log is disabled.");
                System.exit(1);
            }
            catch (org.omg.DsLogAdmin.LogOffDuty ex)
            {
                System.err.println("Log is off duty.");
                System.exit(1);
            }
        }
        else
        {
            if(log.get_availability_status().off_duty)
            {
                System.out.println("The log is off duty " +
                    "- cannot log data.");
            }
            if(System.currentTimeMillis() >= stopTime)
            {
                System.out.println("Finished sending data.\n");
                sendTimer.cancel();
            }
        }
    } // end run
} // end inner class Sender
```

The `get_availability_status().off_duty` value is used to determine if the log is available to store records: the client uses this method to test to whether it can write data to the log. A log can become *off duty*, that is unavailable to receive data or events, for several reasons, including:

- The log's operational state is disabled
- The log's administrative state is locked (refer to the code showing `log.set_administrative_state(AdministrativeState.locked)` under "[Log Creation and Configuration](#)" above)



- The current time is outside the log's *duration time* (refer to the code showing `log.set_interval(interval)` under “[Log Creation and Configuration](#)” above)
- The current time is outside the log's scheduled times (scheduled times are set using the Log's `set_week_mask()` method)

Provided the log is available, the data (as an array of `Any`s) is sent to the log by simply using:

```
log.write_records(data);
```

Each `Any` in the array will be stored as a separate record; if the array contains four elements, then four records will be stored.

## Querying the Log Records

The code from the example client shown below:

- Creates a simple query which retrieves selected log records
- Retrieves information from the record

The client begins querying the log's records two minutes after it starts sending data, every two seconds for eight minutes: the frequency and duration is controlled using the same approach that is used for sending the data, that is by using `Timer.schedule()` and a `TimerTask.run()` implementation as explained under “[Sending the Data](#)” above.

A query consists of two elements:

- A constraint expression, stated according to a supported grammar and constraint language, e.g. `EXTENDED_TCL`, and
- Making the query using Log's `query()` method with the constraint expression, grammar, and iterator parameters (`query()` sets the iterator via an iterator holder)

The query will return a `LogRecord` array of those records which satisfy the conditions defined in the constraint.

A `LogRecord` is a structure containing the following elements:

- *id* - a unique number which is assigned by the log
- *time* - the time the log was created (based on the Gregorian calendar, expressed in 100 nanosecond units (IDL type `TimeT`))
- *attr\_list* - an (optional) attribute list of user defined name-value pairs which can provide information about the date or event, but which is not part of the data, as such
- *info* - the actual data or event being stored

### 1 Creating a Query Expression

A `LogRecord` array is created to hold records which are returned by the query. The default grammar (the OpenFusion Log Service uses `EXTENDED_TCL`) is assigned to a string which will be passed to `query()`. The constraint expression states, in this case, that only records that contain the string “Hello World!” should be retrieved. Note the single quotes surrounding the `'Hello world!'` data string. The `query()` method returns the number of query results (via `IteratorHolder`).

```
// thread for querying log records
class Receiver extends TimerTask
{
    // initialise sample query
    LogRecord[] records;
```

```
String grammar = default_grammar.value;
String constraint = "$.info == 'Hello world!'";
IteratorHolder iter = new IteratorHolder();

long stopTime = System.currentTimeMillis() + duration;
```

## 2 Sending the Query

The following code executes the query, saves the result to the `LogRecord` array, called *records*, then retrieves and prints the id number and the data which was stored in the last record stored in *records*.

The query is called within a *try-catch* block in order to catch the *InvalidGrammar* and *InvalidConstraint* exceptions.

```
public void run()
{
    if(System.currentTimeMillis() < stopTime)
    {
        try
        {
            records = log.query (grammar, constraint, iter);
            if (records.length >= 1)
            {
                int last = records.length - 1;
                System.out.println ("The last log record\'s id is: " +
                    records[last].id + " and contains " +
                    records[last].info);
            } else
            {
                System.out.println ("No records in log store.\n");
            }
        }
        catch (org.omg.DsLogAdmin.InvalidGrammar ex)
        {
            System.err.println ("Invalid grammar.");
            System.exit (1);
        }
        catch (org.omg.DsLogAdmin.InvalidConstraint ex)
        {
            System.err.println ("Invalid constraint.");
            System.exit (1);
        }
    }
    else
    {
        System.out.println("Finished querying log records.\n");
        receiveTimer.cancel();
    }
}
} //end inner class Receiver
```

## Client Exceptions for the BasicLog

The code used in the example client program must throw or catch the following exceptions.

### ORB

*org.omg.CORBA.ORBPackage.InvalidName*

### BasicLog (including Write and Query Related)

*org.omg.DsLogAdmin.LogDisabled*

*org.omg.DsLogAdmin.LogOffDuty*

*org.omg.DsLogAdmin.LogFull*

*org.omg.DsLogAdmin.LogLocked*

*org.omg.DsLogAdmin.InvalidGrammar*

*org.omg.DsLogAdmin.InvalidConstraint*

### **Naming Service**

```
org.omg.CosNaming.NamingContextPackage.NotFound  
org.omg.CosNaming.NamingContextPackage.CannotProceed  
org.omg.CosNaming.NamingContextPackage.InvalidName
```

## **The NotifyLog and Event-style Events**

This topic describes the essential aspects for working with the NotifyLog. Understanding these aspects is a precursor to understanding the NotifyLog's more powerful features. The topic explains how to:

- 1 Create an NotifyLog and perform initial configuration for this log type.
- 2 Set log scheduling using the `set_week_mask()` method.
- 3 Create a client which supplies the log with Event-style events.
- 4 Create a client which receives Event-style events forwarded by the log and also retrieves the latest log records.
- 5 Perform clean-up operations when terminating NotifyLog clients.

The NotifyLog inherits from both the `Log` interface and the Notification Service's event channel interfaces. (see [Figure 1, Log Types and Inheritance](#)) Consequently, it is more complex than the BasicLog and requires an understanding of the Notification Service's proxies and channels. Developing an understanding of proxies and channels is important generally since they are used by all log types except the BasicLog.

An example application consisting of three program components is used here to demonstrate how to create and interact with a NotifyLog. The three component model used by the example application reflects an architecture which may likely be used in a working environment.

### **Creating a NotifyLog**

This example program creates, configures and instantiates a NotifyLog as a CORBA object. The program also uses the Naming Service to bind the log object to a name, as was used in the BasicLog example program, to facilitate location of the log by the application's client programs.

### **Import Statements**

This example NotifyLog creation program uses the services and associated APIs listed below. Applications or classes that use these services or APIs must include the import statements shown below.

#### **OpenFusion Orb Adaptor**

```
com.prismt.orb.ORBAdapter
```

#### **NotifyLog**

```
org.omg.DsLogAdmin.*  
org.omg.DsNotifyLogAdmin.*  
org.omg.CosNotification.*
```

#### **Naming Service**

```
org.omg.CosNaming.*  
org.omg.CosNaming.NamingContextPackage.*  
org.omg.CosNaming.NamingContextExt.*  
org.omg.CosNaming.NamingContextExtPackage.*
```

## NotifyLog Initialisation

The orb object for the program is declared and initialised in the same manner as was shown under “[The Basics and the BasicLog](#)”.

## NotifyLog Creation and Configuration

A NotifyLog is created by obtaining a reference to a *NotifyLogFactory*, then using the factory to create a *NotifyLog* object with the requisite configuration information. The *NotifyLogFactory*'s *create\_with\_id()* method requires:

- The record store's maximum size, in bytes
- Whether to wrap (overwrite the oldest records) or reject new records when the record store is full (*halt.value*)
- The levels when alarms are generated as the record store becomes full, expressed as a percentage of the record store's maximum size
- The initial Quality of Service (QoS) property settings
- The initial Admin property settings

It is possible to avoid setting the QoS or Admin properties by passing zero length *Property* arrays to the create methods.

Any remaining configuration properties which are required to be set to non-default values can be set after the log is created.

*i*

The OMG specifies three Log specific QoS levels: *QoSNone*, *QoSFlush*, and *QoSReliability*. The OpenFusion implementation of the Log Service provides *QoSReliability* as the default (and only) Log-specific QoS level: this level guarantees that all log records sent to a log will be available and enables log records to be recovered if, for example, the system crashes.

Also, the OpenFusion NotifyLog inherits from the OpenFusion Notification Service, including all of its standard OMG and OpenFusion extended QoS properties.

### 1 Obtain a NotifyLog factory.

The following example code obtains a CORBA object for a *NotifyLogFactory* using the ORB's *resolve\_initial\_references()* method, then narrows the object to a *NotifyLogFactory*. Code should be included to test whether the factory was successfully created.

```
orb = ORBAdapter.init(args);

org.omg.CORBA.Object object =
    orb.resolve_initial_references("NotifyLogFactory");
NotifyLogFactory factory = NotifyLogFactoryHelper.narrow(object);

if(factory == null)
{
    System.out.println("Can not resolve initial reference " +
        "of NotifyLogFactory. ");
    System.exit(1);
}
```

## 2 Create the log.

The `NotifyLog` is created with `EventLogFactory.create_with_id()`. In this example, the `NotifyLog`, called `log`, will:

- a Have a record store having a maximum size of 10,000 bytes. The size used in a real environment should reflect the approximate size (in bytes) of each record to be stored multiplied by the number of records the log is required to hold.
- b Overwrite the oldest records when full (using `wrap.value`)
- c Generate alarms when the log is 25, 50, and 75 percent full (defined with an array of shorts called `thresholds`)

The `logId` variable (an `IntHolder`) is used to hold the log's id value (which is assigned by `create_with_id()`).

```
NotifyLog log = null;
int logId = 0;
short[] thresholds = {25, 50, 75};

Property[] qos = new Property[0];
Property[] admin = new Property[0];

// create new log, replacing existing log having same id if it exists
org.omg.CORBA.Object corbaObj = null;
corbaObj = factory.find_log(logId);
if ( corbaObj != null )
{
    System.out.println("Existing log found and being replaced.");
    log = NotifyLogHelper.narrow(corbaObj);
    log.destroy();
}

log = factory.create_with_id(logId, wrap.value, 10000,
    thresholds, qos, admin);
```

Note that `factory.find_log(logId)` is used to determine if a log with the same id already exists - if does, then it is destroyed (with `destroy()`).

*i*

`destroy()` removes the log, with its log store, and disconnects any clients which are connected to the log: it should be used to remove the log, but with caution, since the careless use of `destroy()` could result with the unintentional loss of records.

## 3 Configure the log.

The process for configuring a `NotifyLog` is similar to the process shown above for the `BasicLog`. This example shows how to configure the intervals that a log is allowed to store records using `WeekMask`. `WeekMask` is more flexible than `TimeInterval` (used in the `BasicLog` example), such that multiple start and stop times can be set to occur on specified days of the week. (`WeekMask` uses 24 hour time in the range of 0-23 hours and 0-59 minutes, as opposed to the Gregorian time used by `TimeInterval`.)

The following code creates a schedule that allows the log to store records on Mondays, Tuesdays, Wednesdays, Thursdays, and Fridays, for two intervals, from 0900 to 1200 and from 1300 to 1700, each day. After the day and time values are set, `log.set_week_mask(masks)` is called to set the schedule.

Note that the log's administrative state is locked while the `WeekMask` is being set. This is done in order to stop records being stored while the

properties are being changed.

```
// lock the log prior to setting its operational interval
log.set_administrative_state(AdministrativeState.locked);

// Set length of days and times for log to accept data
WeekMaskItem[] masks = new WeekMaskItem[1];
masks[0] = new WeekMaskItem();
masks[0].days = (short) (Monday.value | Tuesday.value |
Wednesday.value
                        | Thursday.value | Friday.value);
masks[0].intervals = new Time24Interval[2];

masks[0].intervals[0] = new Time24Interval();
masks[0].intervals[0].start = new Time24((short) 9, (short) 0);
masks[0].intervals[0].stop = new Time24((short) 12, (short) 0);

masks[0].intervals[1] = new Time24Interval();
masks[0].intervals[1].start = new Time24((short) 13, (short) 0);
masks[0].intervals[1].stop = new Time24((short) 17, (short) 30);

log.set_week_mask(masks);

// unlock the log after setting the operational interval
log.set_administrative_state(AdministrativeState.unlocked);
```

## Locating the NotifyLog

This example, like the BasicLog example, uses the Naming Service to facilitate the locating of the log object by clients. The code used here is nearly identical to the equivalent code used for the BasicLog example, however the `id` value of the `NameComponent` has been changed to `"NotifyLogA"` in order to distinguish this log object from the BasicLog object used previously and the NotifyLog object to be used in the ["NotifyLog, Structured Events and More"](#) example on [page 39](#). (Although `"NotifyLogA"` is used here, any string value could be used, provide the `id` and `kind` combination is unique within the naming context.)

```
// bind log object to a name using the Naming Service
org.omg.CORBA.Object namingService = null;
NamingContextExt rootContext = null;
NameComponent notifyLog[];

namingService = orb.resolve_initial_references ("NameService");
rootContext = NamingContextExtHelper.narrow(namingService);
notifyLog = new NameComponent[1];
notifyLog[0] = new NameComponent("NotifyLogA", "log");
rootContext.rebind(notifyLog, log);
System.out.println("NotifyLogA created");
```

## NotifyLog Exceptions

The example code shown above must either throw or catch the exceptions listed below. Log exception details, for this and all other examples, are described in ["Exceptions"](#).

### ORB and General (connecting to services)

*org.omg.CORBA.ORBPackage.InvalidName*

### NotifyLog and Duration

*org.omg.DsLogAdmin.LogIdAlreadyExists*  
*org.omg.DsLogAdmin.InvalidLogFullAction*  
*org.omg.DsLogAdmin.InvalidTime*  
*org.omg.DsLogAdmin.InvalidTimeInterval*  
*org.omg.DsLogAdmin.InvalidThreshold*  
*org.omg.DsLogAdmin.InvalidMask*  
*org.omg.CosNotification.UnsupportedQoS*

`org.omg.CosNotification.UnsupportedAdmin`

### Naming Service

`org.omg.CosNaming.NamingContextPackage.InvalidName`

`org.omg.CosNaming.NamingContextPackage.NotFound`

`org.omg.CosNaming.NamingContextPackage.CannotProceed`

## Creating a Supplier Client for the NotifyLog

The following supplier client program repeatedly sends Event-style events to a NotifyLog instance. This example client utilises J2SE's `Timer` and `TimerTask` classes to control how long and at what interval to send the events.

The most important difference between this client and the client used for the BasicLog example is that this client connects to the log as an *event supplier* - that is, it supplies events to the NotifyLog's *event channel* via a *proxy* (remembering that the NotifyLog inherits, via the `NotifyChannel`, from the `EventChannel` - as well as from `Log`). The client for the BasicLog wrote data directly to the log store, whereas this client sends events to the NotifyLog's event channel; the log, itself, writes the events to the record store, as well as forwarding them to consumer clients.

Accordingly, the supplier client uses two `CosNotify` interfaces (see "Import Statements" below), plus it must implement `PushSupplierOperations`.

*i*

Classes which implement any of the `Operations` interfaces must use `com.prismt.orb.ObjectAdapter` methods. Accordingly, unlike the previous examples, this client uses `ObjectAdapter` instead of `ORBAdapter` to connect with and initialise the ORB. However, there may be situations in working applications where methods from *both* `ObjectAdapter` and `ORBAdapter` are needed.

The `PushSupplierOperations` methods which are required to be implemented by the supplier are `disconnect_push_supplier()` and `subscription_change()`. These are callback methods which enables the NotifyLog's event channel to respectively inform the client that it has been disconnected or that there has been a subscription change. (Note that the Notification Service does not currently implement the `NotifyPublish` or `NotifySubscribe` interfaces, and so will never call the `subscription_change()` or `offer_change()` methods.)

```
// PushSupplierOperations callback method
public void disconnect_push_supplier()
{
    System.out.println ("Disconnected by proxy");
} // end disconnect_push_supplier

public void subscription_change(org.omg.CosNotification.EventType[]
added,
    org.omg.CosNotification.EventType[] removed)
{
    System.out.println ("Supplier subscription changed");
} // end subscription_change
```

## Import Statements

The example supplier program uses the services and associated APIs listed below. Applications or classes that use these services or APIs must include the import statements shown below.

### OpenFusion Orb Adaptor

```
com.prismt.orb.ObjectAdapter
```

### NotifyLog

```
org.omg.DsLogAdmin.*
org.omg.DsEventLogAdmin.*
org.omg.CosEventComm.*
org.omg.CosEventChannelAdmin.*
```

### Naming Service

```
org.omg.CosNaming.*
org.omg.CosNaming.NamingContextPackage.*
org.omg.CosNaming.NamingContextExt.*
org.omg.CosNaming.NamingContextExtPackage.*
```

### J2SE'S Timer and TimerTask

```
java.util.*
java.util.TimerTask
```

## Supplier Client Initialisation

The program declares, as in the previous examples, *orb* and *log* objects, plus an *org.omg.CORBA.Any*, called *event*, which is the Event-style event that will be sent to the log. A *ProxyPushConsumer* object is also declared, since the program will connect to the consumer client via the log's push consumer proxy (see ["The Client as an Event Supplier"](#) and ["Creating a Consumer Client for the NotifyLog"](#)).

Since the Naming Service will be used to retrieve the *log* object, associated variables required to locate the log via the Naming Service are also declared (refer to ["The Basics and the BasicLog"](#) and the Appendix ["Using the Naming Service to Locate a Log"](#) for details).

Like the previous examples, the *Timer* and *TimerTask* classes control the interval and duration when events are sent, as specified by the *interval* and *duration* variables.

```
// event suppliers must implement supplier operations
public class NotifySupplierA implements PushSupplierOperations
{
    org.omg.CORBA.ORB orb = null;
    org.omg.CORBA.Object obj = null;

    NotifyLog log = null;
    org.omg.CORBA.Any event;
    ProxyPushConsumer consumerProxy = null;

    NamingContextExt rootContext = null;
    NameComponent notifyLog[];

    private Timer sendTimer;

    private final long duration = 3 * 60 * 1000; // 3 seconds in
millisecs
    private final long interval = 2000; // 2 seconds
    private final long delay = 0;
```



The client program then:

- Initialises the ORB
- Establishes the client as a transient CORBA object
- Obtains a reference to the NotifyLog instance using the Naming Service
- Establishes itself as a *PushSupplier* and
- Obtains a reference to the log's *push consumer proxy*, thereby connecting itself to the log (described under “[The Client as an Event Supplier](#)” below)

When these tasks are completed, the client informs the ORB that it is ready to accept incoming requests (the only requests that might be received by this example will be from `disconnect_push_supplier()`).

```
orb = ObjectAdapter.init(args);

// instantiate this client as a transient CORBA object
ObjectAdapter.createTransient(this);

// obtain reference to log using Naming Service
obj = orb.resolve_initial_references("NameService");
rootContext = NamingContextExtHelper.narrow(obj);
notifyLog = new NameComponent[1];
notifyLog[0] = new NameComponent("NotifyLogA", "log");
log = NotifyLogHelper.narrow(rootContext.resolve(notifyLog));

// establish this object as a push supplier and
// connect to the log's proxy consumer
PushSupplier supplier =
    PushSupplierHelper.narrow(ObjectAdapter.getObject(this));
SupplierAdmin supplierAdmin = log.default_supplier_admin();
org.omg.CORBA.IntHolder id = new org.omg.CORBA.IntHolder();
ProxyConsumer proxy =
    supplierAdmin.obtain_notification_push_consumer(
        ClientType.ANY_EVENT, id);

consumerProxy = ProxyPushConsumerHelper.narrow(proxy);
consumerProxy.connect_any_push_supplier(supplier);

// enable incoming requests, without blocking, for log callbacks
ObjectAdapter.ready(false);
```

## The Client as an Event Supplier

Any client which sends or receives events to or from an event channel must be narrowed to an (Event Service) supplier or consumer type. Further, suppliers and consumers will use one of two models to send or receive events: *push* or *pull*.

*i*

- The push model is the only model supported by the OpenFusion Log Service implementation since it model which is typically or generally used. Briefly:
- *push suppliers* actively push events to the event channel
- *push consumers* passively receive events sent by the channel, i.e. a push consumer will receive any event that is sent to it by the event channel
- *pull suppliers* have their events pulled from them by the channel, i.e. the supplier must supply an event (if it has one), whenever the channel makes a request
- *pull consumers* actively pull events from the event channel, i.e. they retrieve events whenever they want them

Our supplier client is a *PushSupplier*. Its consumer client, described below, is a *PushConsumer*.

Clients connect to an event channel through a *proxy*: the proxy represents the consumer or supplier object(s) that a client intends to transmit events to or receive events from: a *supplier* connects to a *consumer proxy*, and visa versa.

The proxy, in combination with an event channel, enables a client to connect transparently with any number of other clients, without the need to obtain a reference to each client that is added to the system.

*i*

A detailed discussion of these issues is beyond the scope of this topic, however the OMG's *Event Service Specification* and the **OpenFusion Notification Service Guide** provide a complete description of push and pull models, suppliers, consumers, proxies and related topics.

### Connecting a Client to an Event Channel

A connection between a client and an event channel requires the following:

- 1 Narrow the client to the appropriate type, for example:

```
PushSupplier supplier =  
    PushSupplierHelper.narrow(ObjectAdapter.getObject(this));
```

where *this*, in this example, is the client itself

- 2 Obtain a reference to the channel's *admin* object for the respective supplier or consumer. Since the log is a subtype of a notification channel (which is a subtype of an event channel), call `NotifyLog.default_supplier_admin()` to obtain the channel's default supplier admin object:

```
SupplierAdmin supplierAdmin =  
    log.default_supplier_admin();
```

- 3 Obtain a reference from the admin object to the proxy (remembering that suppliers connect to consumer proxies and visa versa):

```
ProxyConsumer proxy =  
    supplierAdmin.obtain_notification_push_consumer(  
        ClientType.ANY_EVENT, id);
```

```
ProxyPushConsumer consumerProxy =  
    ProxyPushConsumerHelper.narrow(proxy);
```

- 4 Connect the client to the proxy:

```
consumerProxy.connect_any_push_supplier(supplier);
```

- 5 Inform the ORB that the client is ready to supply events:

```
ObjectAdapter.ready(false);
```

## Sending Events

The client sends a series of Event-style events, *Anys*, to the log. Like the BasicLog example, the time interval and duration of when the events are sent is controlled by *Timer* and an implementation of the abstract *TimerTask* class' *run()* method.

The *TimerTask* implementation, *Sender.run()* shown below, uses *stopSendingTime* to define when to stop sending events; the *num* variable is used to provide a data value for each event and is incremented each time an event is sent.

```
// thread for sending data
class Sender extends TimerTask
{
    // data value to be placed in each event
    int num = 0;

    long stopSendingTime = System.currentTimeMillis() + duration;
```

Provided the current time is earlier than *stopSendingTime* and the log's *off\_duty* property is false, then a *num* value will be inserted into each event. The event is then sent to the log via the log's consumer proxy, *consumerProxy.push(event)*.

An *org.omg.CosEventComm.Disconnected* exception is thrown if the client becomes disconnected from the proxy.

```
public void run()
{
    if((System.currentTimeMillis() < stopSendingTime) &&
        (log.get_availability_status().off_duty == false))
    {
        try
        {
            num++;
            event = orb.create_any();
            event.insert_long(num);
            consumerProxy.push(event);
            System.out.println("Event sent with value " + num + ".");
        }
        catch (org.omg.CosEventComm.Disconnected ex)
        {
            System.exit(1);
        }
    }
    else
    {
        if(log.get_availability_status().off_duty)
        {
            System.out.println("The log is off duty " +
                "- cannot log event.");
        }
        if(System.currentTimeMillis() >= stopSendingTime)
        {
            System.out.println("Finished sending events.\n");
        }
    }
}
```

## Terminating the Supplier Client

Before an event-based client terminates it must destroy any proxy objects that it has created. If these proxies are not destroyed, then they will become a source of leakage. Also, logs have no means of knowing if a client terminates unless the proxy is destroyed: it is therefore the client's responsibility to disconnect from a log by destroying the proxy. A client disconnects from event-based logs and destroys the associated proxy using an appropriate proxy disconnect method.

*i*

The OMG naming convention for disconnecting a proxy is the reverse of its equivalent connection method, whereby the *supplier* part of the method's name is replaced by *consumer* - which can be confusing. The easiest means of clarifying the convention is by example:

- A push supplier called *supplier connects* to a `ProxyPushConsumer` instance called *consumerProxy* using `connect_push_supplier(supplier)`
- A push supplier called *supplier disconnects* from a `ProxyPushConsumer` instance called *consumerProxy* using `disconnect_push_consumer()`

The example client disconnects from the proxy when `Sender.run()` has finished sending events.

```
// cleanup - destroy proxy, stop incoming requests
consumerProxy.disconnect_push_consumer();
ObjectAdapter.shutdown();

// stop thread
sendTimer.cancel();
}
} // end run
} // end inner class Sender
```

## Supplier Client Exceptions

The example code shown above, which creates a `NotifyLog` supplier client and uses a Naming Service name binding, must either throw or catch the exceptions listed below.

### ORB

`org.omg.CORBA.ORBPackage.InvalidName`

### NotifyLog

`org.omg.CosNotifyChannelAdmin.AdminLimitExceeded`  
`org.omg.CosEventChannelAdmin.AlreadyConnected`

### Naming Service

`org.omg.CosNaming.NamingContextPackage.InvalidName`  
`org.omg.CosNaming.NamingContextPackage.NotFound`  
`org.omg.CosNaming.NamingContextPackage.CannotProceed`

## Creating a Consumer Client for the NotifyLog

The example consumer client program:

- Passively receives events, which have been forwarded by the log from the supplier
- Retrieves the most recent ten records that the log is storing at a point in time five minutes from the time the consumer client was invoked

The client is implemented as a *push consumer*, able to passively accept events pushed to it by the log's event channel (see "[The Client as an Event Supplier](#)"). In many ways, the consumer client is similar to the supplier:

- They both implement `org.omg.CosEventComm`'s callback operations. Notwithstanding that the consumer client is a *push consumer* and not a *push supplier*, it implements the `PushConsumerOperations` instead of the `PushSupplierOperations` - and with one additional, important difference: there is an additional callback method, `push()`, which must be implemented and which receives the events transmitted by the channel.

- They both use a proxy to communicate with the log; the proxy, a *ProxyPushSupplier* in the consumer's case, is obtained using methods similar to those that were used by the supplier client
- They both use the Naming Service to locate the log object

## Import Statements

The program uses the services and associated APIs listed below. Applications or classes that use these services or APIs must include the import statements shown below.

### ORB

```
com.prismt.orb.ObjectAdapter
```

### NotifyLog and for Time values

```
org.omg.DsLogAdmin.*
org.omg.DsNotifyLogAdmin.*
org.omg.CosNotification.EventType
org.omg.CosNotifyComm.*
org.omg.CosNotifyChannelAdmin.*
java.util.GregorianCalendar
java.util.Date
```

### Naming Service

```
org.omg.CosNaming.*
org.omg.CosNaming.NamingContextPackage.*
org.omg.CosNaming.NamingContextExt.*
org.omg.CosNaming.NamingContextExtPackage.*
```

## Consumer Client Initialisation

This client declares approximately the same type of objects and variables as declared for the supplier client, with the exception that *Timer* is not used - since the *PushConsumerOperations.push()* method (implemented by the client) receives events whenever they arrive. (Compare this with the supplier client initialisation code shown under [“Supplier Client Initialisation”](#)).

```
org.omg.CORBA.ORB orb = null;
org.omg.CORBA.Object obj = null;

NotifyLog log = null;
org.omg.CORBA.Any event;
ProxyPushSupplier supplierProxy;

NamingContextExt rootContext = null;
NameComponent notifyLog[];

private final long duration = 3 * 60 * 1000; // 3 minutes in millisecs
```

The following code example shows, as in the supplier client example, that the ORB is initialised, the client is declared as a transient CORBA object, a reference to the NotifyLog object is obtained using the Naming Service, and the client makes itself available to receive requests.

```
orb = ObjectAdapter.init(args);

// instantiate this client as a transient CORBA object
ObjectAdapter.createTransient(this);

// obtain reference to log using Naming Service
obj = orb.resolve_initial_references("NameService");
rootContext = NamingContextExtHelper.narrow(obj);
notifyLog = new NameComponent[1];
notifyLog[0] = new NameComponent("NotifyLogA", "log");
```

```

log = NotifyLogHelper.narrow(rootContext.resolve(notifyLog));

// establish this object as a push consumer and
// connect to the log's proxy supplier
PushConsumer consumer =
    PushConsumerHelper.narrow(ObjectAdapter.getObject(this));

ConsumerAdmin consumerAdmin = log.default_consumer_admin();
org.omg.CORBA.IntHolder id = new org.omg.CORBA.IntHolder();

ProxySupplier proxy =
    consumerAdmin.obtain_notification_push_supplier (
        ClientType.ANY_EVENT, id);

supplierProxy = ProxyPushSupplierHelper.narrow(proxy);
supplierProxy.connect_any_push_consumer(consumer);

// enable incoming requests, without blocking, for log callbacks
ObjectAdapter.ready(false);

```

After the client is established as a *PushConsumer*, it is connected to the log's supplier proxy using:

```

PushConsumer consumer =
PushConsumerHelper.narrow(ObjectAdapter.getObject(this));

ConsumerAdmin consumerAdmin = log.default_consumer_admin();
org.omg.CORBA.IntHolder id = new org.omg.CORBA.IntHolder();

ProxySupplier proxy = consumerAdmin.obtain_notification_push_supplier
(ClientType.ANY_EVENT, id);

supplierProxy = ProxyPushSupplierHelper.narrow(proxy);
supplierProxy.connect_any_push_consumer(consumer);

```

## The Client as an Event Consumer

Similarly, as for the clients which are supplying the events, the consumer clients connect to an event channel through a proxy using the *push* model (remembering that the OpenFusion implementation only supports the push model). The steps required to establish an event consumer are the same as those for establishing an event supplier (see [“The Client as an Event Supplier”](#)), except that:

- A *ConsumerAdmin* object is used to obtain a *supplier* proxy
- The consumer connects to the *supplier* proxy using `obtain_push_supplier()` for the push model and `obtain_pull_supplier()` for the pull model
- The `push()` callback method is the method that will retrieve events when using the push model - the developer must implement this method to suit their particular requirements

## Receiving Events

The consumer client's `push()` method implementation, shown below, receives the events which have been forwarded by the `NotifyLog`.

```
// PushConsumerOperations callback methods
public void disconnect_push_consumer()
{
    System.out.println ("Disconnected by proxy.");
} // end disconnect_push_consumer

public void push(org.omg.CORBA.Any event)
{
    System.out.println("Event received contains: " +
event.extract_long());
} // end push

public void offer_change(EventType[] added, EventType[] removed)
{
    System.out.println ("Consumer offer changed.");
} // end offer_change
```

This particular implementation extracts a `long` from the event using the event's `extract_long()` method, then displays the value.

Note that the log's forwarding state must be set to `on` (which is the default value) in order for the log to forward events to the consumer. If this state is `off`, it can be set to `on` using:

```
log.set_forwarding_state(ForwardingState.on)
```

## Retrieving Records

The `Log.retrieve()` method (inherited by all log types) retrieves a sequential list of records, based on the time that the first or last record in the list was placed in the record store. Note that `retrieve()` uses Gregorian time.

`retrieve()` takes three parameters (listed in order):

- `from_time` - the Gregorian time, in 100 nanosecond units, from when the first record in the required list was stored
- `how_many` - the number of records to retrieve: if this value is positive, then the records will be retrieve in a positive direction from the oldest record to the newest; if the value is negative, then the newest record will be retrieved first followed by the next newest and so forth in a negative direction
- `i` - an iterator holder which is returned by `retrieve()` reporting the number of records which have been returned; the iterator can be declared as null if this value is not needed

The code example below retrieves the latest ten records which have been stored in the log: the `from_time` parameter is set to the current time and the `how_many` parameter is set to `-10`. This will retrieve records starting with the record which was stored at the current time and then work backwards, retrieving successively older records. The retrieved records will be held in a `LogRecord` array.

```
// array for retrieved records
LogRecord[] records;

IteratorHolder iter = new IteratorHolder();

// establish current time and convert ms to 100 ns
GregorianCalendar calendar = new GregorianCalendar();
long basetime = - calendar.getGregorianChange().getTime();
long currentTime = (System.currentTimeMillis () + basetime) * 10000;

System.out.println(
    "\nRetrieving the latest 10 records backwards from the current time.");
records = log.retrieve(currentTime, -10, iter);
if (records.length < 1 )
{
    System.out.println("No records retrieved.\n");
}
else
{
    // display information for each retrieved record
    System.out.println ("\n-----\n" +
        "ID\tTime\t\t\t\t\tInfo");
    for(int i = 0; i < records.length; i++)
    {
        System.out.println(records[i].id +
            "\t" + new Date(records[i].time/10000 - basetime) +
            "\t" + records[i].info);
    }
    System.out.println ("-----");
}
}
```

The client then displays the time the retrieved records were stored, along with their event data (held in the record's `info` field).



## Cleanup

Before the client terminates it must destroy the proxy it is connected to and stop receiving requests from the ORB: left-over proxies which are not destroyed contribute to resource leakage.

```
supplierProxy.disconnect_push_supplier();
ObjectAdapter.shutdown();
```

## Exceptions

The example code shown above, which creates a NotifyLog consumer client and uses a Naming Service name binding, must either throw or catch the exceptions listed below.

### ORB

*org.omg.CORBA.ORBPackage.InvalidName*

### NotifyLog

*org.omg.CosNotifyChannelAdmin.AdminLimitExceeded*

*org.omg.CosEventChannelAdmin.TypeError*

*org.omg.CosEventChannelAdmin.AlreadyConnected*

### Naming Service

*org.omg.CosNaming.NamingContextPackage.InvalidName*

*org.omg.CosNaming.NamingContextPackage.CannotProceed*

*org.omg.CosNaming.NamingContextPackage.NotFound*

## NotifyLog, Structured Events and More

This topic introduces some of the NotifyLog's more advanced features and explains how to:

- 1 Set Quality of Service (QoS) and Administration properties not covered in [“The NotifyLog and Event-style Events”](#).
- 2 Use a log filter for storing only selected records.
- 3 Use a filter on the log's notification channel to filter events sent to consumer clients.
- 4 Implement a supplier client which
  - a creates and sends structured events to the log
  - b sets a QoS property for its proxy.
- 5 Implement a consumer client which can
  - a receive structured events forwarded by the log
  - b retrieve information held in the NotifyLog's record store.
- 6 Monitor events generated by the log and log factory.

## Creating a NotifyLog

This example shows how to create a NotifyLog which will use structured events. This example extends information provided in the previous BasicLog and NotifyLog topics, especially regarding items which are specific to using structured events.

## Import Statements

The example NotifyLog creation program uses the services and associated APIs listed below. Applications or classes that use these services or APIs must include the import statements shown below.

### ORB

```
com.prismt.orb.ORBAdapter
```

### NotifyLog

```
org.omg.DsLogAdmin.*
org.omg.DsNotifyLogAdmin.*
org.omg.CosNotification.*
org.omg.CosNotifyFilter.Filter;
import org.omg.CosNotifyFilter.FilterFactory;
import org.omg.CosNotifyFilter.ConstraintExp;
import org.omg.CosNotifyFilter.InvalidConstraint;
import org.omg.CosNotifyFilter.InvalidGrammar
```

### Naming Service

```
org.omg.CosNaming.*
org.omg.CosNaming.NamingContextPackage.*
org.omg.CosNaming.NamingContextExt.*
org.omg.CosNaming.NamingContextExtPackage.*
```

## NotifyLog Initialisation

The ORB object for the program is declared and initialised in the same manner as described in [“The Basics and the BasicLog”](#). Like the previous examples, the Naming Service is used to facilitate location of the log object and therefore Naming Service objects are also declared.

```
org.omg.CORBA.ORB orb = null;

org.omg.CORBA.Object namingService = null;
NamingContextExt rootContext = null;
NameComponent notifyLog[];
```

## NotifyLog Creation and Configuration

A NotifyLog object, called *log*, is created using the `create_with_id()` method, as similarly shown under [“The NotifyLog and Event-style Events”](#). Additional QoS and Admin properties to those used in the [“The NotifyLog and Event-style Events”](#) are configured and set.

As shown in the following code extract, the log object in this example is assigned an id of 10. A test is made to determine if a log with an id of 10 already exists: this examples replaces the existing log: a working application may decide to simply change the value of *logId* and create a new, different log.

```
orb = ORBAdapter.init(args);

org.omg.CORBA.Object object =
    orb.resolve_initial_references("NotifyLogFactory");
NotifyLogFactory factory = NotifyLogFactoryHelper.narrow(object);

if(factory == null)
{
    System.out.println("Can not resolve initial reference " +
        "of NotifyLogFactory.");
    System.exit(1);
}

NotifyLog log = null;
```

```

int logId = 10;

// If Notify log with specified id exists, remove it, then
// create a new one with an id of 0, which discards oldest events
// when full, has a size limit of 50,000 bytes, and generates
// threshold alarms when 25%, 50%, and 75% capacity reached.
org.omg.CORBA.Object corbaObj = null;
corbaObj = factory.find_log(logId);
if ( corbaObj != null )
{
    System.out.println("Existing log found and being replaced.");
    log = NotifyLogHelper.narrow(corbaObj);
    log.destroy();
}

```

The QoS and Admin properties are now defined using two `Property` arrays, `qos` and `admin`:

- The `ConnectionReliability` and `EventReliability` properties are both set to `Persistent` to ensure the reliability of the notify channel connection and that events can be recovered in case of a system failure
- `MaxQueueLength` is set to limit the channel's event queue buffer to hold a maximum of 100 events
- `MaxConsumers` is set to limit (to 10) the maximum number of consumers that can be connected at one time to the log

```

// Log QoS set automatically to guarantee log availability
// Notification QoS set to guarantee connection and
// event reliability
Property[] qos = new Property[2];
qos[0] = new Property();
qos[0].name = ConnectionReliability.value;
qos[0].value = orb.create_any();
qos[0].value.insert_short(Persistent.value);

qos[1] = new Property();
qos[1].name = EventReliability.value;
qos[1].value = orb.create_any();
qos[1].value.insert_short(Persistent.value);

// Notification Admin set to have a maximum queue length of
// 100 events and a maximum of 10 consumers
Property[] admin = new Property[2];
admin[0] = new Property();
admin[0].name = MaxQueueLength.value;
admin[0].value = orb.create_any();
admin[0].value.insert_long(100);

admin[1] = new Property();
admin[1].name = MaxConsumers.value;
admin[1].value = orb.create_any();
admin[1].value.insert_long(10);

// set values for threshold alarms
short[] thresholds = {25, 50, 75};

// create the log
log = factory.create_with_id(logId, halt.value, 50000,
                             thresholds, qos, admin);

```

The log record size, threshold levels, and log full behaviour are set.

Note that the `halt` behaviour specified when the record store is full requires clients, which supply events to the log, to test for a log full condition (using `Log.get_availability_status().log_full`): an exception is thrown if a supplier attempts to send an event to a log when the record store is full and `halt` has been specified.

The remaining log properties can now be set. The following code sets the maximum record life to five hours (in units of one second) and ensures that the log will forward events to consumers.

```
// enable log forwarding and set the maximum log record lifetime
// to 5 hours in seconds - prevent any logging while doing so
log.set_administrative_state(AdministrativeState.locked);

log.set_forwarding_state(ForwardingState.on);
log.set_max_record_life (60 * 60 * 5);
```

Note that the log's administrative state has not yet been unlocked, since a filter will be added to the log: the administrative state will be unlocked after the filter has been added.

## Filtering

Structured events which are sent to NotifyLogs can be filtered in order to remove unwanted events. Filtering can be used to control event transmission and to improve overall performance. For example, unwanted events:

- Can be stopped from entering the event channel, thereby improving the log channel's effective bandwidth and performance
- Can be prevented from being placed in the record store, thereby conserving resources
- Can be filtered out on a consumer-by-consumer basis, providing fine-grained control and performance tuning on the client-side of the log's event channel

Logs are able to filter events in one of two ways: on the event channel or immediately before events are received by the record store, which are referred to here, respectively, as *channel-based* filtering or *log-based* filtering. Both of these methods are briefly described below.

*i*

Example code for performing log-based filtering is provided as a part of the example NotifyLog creation program. Example code for performing channel-based filtering is provided under "[Creating a Consumer Client for the NotifyLog](#)", as a part of the example consumer program, and demonstrates how to create a filter on a consumer-by-consumer basis.

Filtering support is provided by the *CosNotifyFilter* module and related interfaces, inherited from the Notification Service. Developers who use filters should have a good understanding of *CosNotifyFilter* and should refer to the sections in the OMG's *Notification Service Specification* and the OpenFusion **Notification Service Guide** which describe the *Filter*, *FilterFactory*, and *FilterAdmin* interfaces. (In addition to the standard *Filter* interface described here, there is also a *MappingFilter* interface: information on Mapping Filters is provided in the referenced texts.)

Although a detailed discussion of filtering is beyond the scope of this manual, basic information and guidance on creating and managing filters are provided on this topic under "[Creating a Consumer Client for the NotifyLog](#)".

## The Filter Object

Filtering is performed with a filter object. A filter can be added to a proxy object, channel admin object, or to a log itself. A single filter object can be added to more than one of these objects at a time: for example a single filter can be used by several proxies, or by a proxy and an admin - however this can lead to unmanageable deployment situations (see warning note shown immediately below).



Filter objects should be destroyed when the objects that use them are destroyed, otherwise they will become a source of memory leakage. However, care must be taken when destroying filter objects that are used by multiple admin/proxy/log objects in order to avoid inadvertently destroying a filter which is still in use.

### Filter Location

The events which are filtered depend on where the filter is added, as follows:

- *consumer proxy* (connects to a supplier) - filters the events received directly from the supplier and therefore events sent to the channel, log, and all supplier proxies (and their associated consumers) are filtered
- *supplier proxy* (connects to a consumer) - filters the events received from the channel before being sent to the proxy's consumer
- *supplier admin* - filters all events received from all consumer proxies connected to the supplier admin and therefore filters events sent to the log and all supplier proxies
- *consumer admin* - filters all events received from the channel and therefore filters events sent to any supplier proxies connected to the consumer admin
- *log* - only filters the events received from the channel which are sent to the log's record store

### Main Filter Components

The main components of a filter are an array of constraint expressions and a set of methods which are used to manage the filter.

A constraint expression defines the criteria by which events are filtered. The expression is written using the Notification Service's constraint language: the service's default grammar is the *Extended Trader Constraint Language*, ETCL. Filtering is performed on the *event header* and *filterable body* components of structured events.

The essential filter methods to be familiar with include:

- those which add, modify, and remove constraints, that is:
  - *add\_constraints()*
  - *modify\_constraints()*
  - *get\_constraints()*
  - *get\_all\_constraints()*
  - *remove\_all\_constraints()*
- the *destroy()* method which destroys the filter object

### Filter Creation

Filters are created as follows:

- 1 Obtain a reference to a `FilterFactory` from an event channel or the log using its `default_filter_factory()` method:

```
FilterFactory factory = Admin.default_filter_factory();
```

or

```
FilterFactory factory = log.default_filter_factory();
```

- 2 Create the filter using the factory's `create_filter()` method, passing it a string defining the grammar to be used - this is normally `default_grammar.value`, which is the Notification Service's default grammar, "EXTENDED\_TCL":

```
Filter filter =
factory.create_filter(default_grammar.value).
```

- 3 Create one or more constraint expressions, of type *ConstraintExp*, then create a *ConstraintExp* array containing the constraints. A constraint expression consists of two elements:
  - a a two element array, of type *EventType*, containing the *domain\_name* and *type\_name* components of the structured event's fixed header - this identifies the type of event to be filtered
  - b an expression, using the same grammar that was passed to *create\_filter()*, which defines the filtering criteria, e.g. "\$*switch* == 'open'", which means that a property called *switch* must contain a value of 'open' in order to be passed through the filter

For example, a simple constraint expression could be constructed as follows:

```
EventType[] type = {"teleco", "mobile"};
ConstraintExp exp =
    new ConstraintExp(type,
        "$.filterable_data(switch) == 'open'");

ConstraintExp[] expressions = {exp};
```

This would allow structured events to pass through the filter which have a fixed header with a *domain\_name* of 'telco', *type\_name* of 'mobile', and a filterable body property called *switch* which contains a value of 'open'.

- 4 Add the constraints array to the filter using its *add\_constraints()* method:

```
filter.add_constraints(expressions);
```

- 5 Associate the filter to the required object(s), noting that:

- a if the filter is to be used as a channel-based filter, i.e. associated with an *admin* or *proxy* object, then the *add\_filter()* method should be used:

```
FilterID id = Admin.add_filter(filter);
```

and where *add\_filter()* returns a unique id (a long); this id is needed by the *remove\_filter()* and *get\_filter()* methods.

- b if the filter is to be used as a log-based filter, i.e. filter the events sent *only* to the record store, then the *set\_filter()* method should be used:

```
log.set_filter(filter);
```

The following code example demonstrates how to create a log-based filter.

```
// filter out all events sent to the record store except where
// the "number" property is an even number (value of 0)
FilterFactory filterFactory = log.default_filter_factory();
Filter filter = filterFactory.create_filter(default_grammar.value);

EventType type = new EventType("software", "example");
EventType[] types = {type};
ConstraintExp exp =
    new ConstraintExp(types, "$.filterable_data(number) == 0");

ConstraintExp[] expressions = {exp};
filter.add_constraints(expressions);
log.set_filter(filter);
```

```
log.set_administrative_state(AdministrativeState.unlocked);
```

(Note that the code shows the log's administrative state being unlocked after setting the filter since configuration of the log is complete at this stage.)

### Filter Removal and Destruction

Removing a filter disassociates the filter from an object. Removing a filter does not destroy it - the filter can still be referenced by other objects which the filter has been added to.

Destroying a filter actually destroys it and releases any resources used by the filter. No object will be able to use the filter after the filter has been destroyed.

- A filter is removed by calling the `remove_filter()` method from the object it is to be removed from, e.g. `Admin.remove_filter(id)`, where `id` is the `FilterID` of the filter to be removed.
- A filter is destroyed by calling its `destroy()` method, e.g. `filter.destroy()`.

A filter does not need to be removed from objects before it is destroyed, i.e. it is not a prerequisite to call `remove_filter()` before calling `destroy()`.

### Registering with the Naming Service

Continuing with the log creation example, the final operation is to bind the log instance to a name using the Naming Service. The `id` component of the name is defined as "NotifyLog" in order to distinguish it from the other logs used in the examples.

```
// register the log object with the Naming Service
namingService = orb.resolve_initial_references ("NameService");
rootContext = NamingContextExtHelper.narrow(namingService);
notifyLog = new NameComponent[1];
notifyLog[0] = new NameComponent("NotifyLogB", "log");
rootContext.rebind(notifyLog, log);
System.out.println("NotifyLogB created");
```

### NotifyLog Exceptions

The example code shown above must either throw or catch the exceptions listed below.

#### ORB (and connecting to services)

```
org.omg.CORBA.ORBPackage.InvalidName
```

#### NotifyLog

```
org.omg.DsLogAdmin.LogIdAlreadyExists
org.omg.DsLogAdmin.InvalidLogFullAction
org.omg.DsLogAdmin.InvalidThreshold
org.omg.CosNotification.UnsupportedAdmin
org.omg.CosNotification.UnsupportedQoS
org.omg.CosNotifyFilter.InvalidGrammar
org.omg.CosNotifyFilter.InvalidConstraint
```

#### Naming Service

```
org.omg.CosNaming.NamingContextPackage.InvalidName
org.omg.CosNaming.NamingContextPackage.NotFound
org.omg.CosNaming.NamingContextPackage.CannotProceed
```

## Creating a Supplier Client for the NotifyLog

The supplier client example for the NotifyLog shows how to:

- establish a *structured push supplier* client which can send structured events to a NotifyLog
- set a Quality of Service (QoS) property for the client's proxy
- create a structured event which contains header information which can be filtered by the log and consumer clients
- handle the situation where the log's record store is full and refuses to store events

The structure of this supplier client example is similar to the supplier client for the example that uses event-style events: events are created and sent at regular intervals for a fixed duration (using the `Timer` and `TimerTask` thread mechanism). Note that the creation of the structured event varies in many ways from the creation of the simple event-style event.

### Import Statements

The program uses the services and associated APIs listed below. Applications or classes that use these services or APIs must include the import statements shown below.

#### ORB

```
com.prismt.orb.ObjectAdapter
```

#### NotifyLog

```
org.omg.DsLogAdmin.*
org.omg.DsNotifyLogAdmin.*
org.omg.CosNotifyComm.*
org.omg.CosNotifyChannelAdmin.*
org.omg.CosNotification.*
com.prismt.cos.CosNotification.NotificationExtensions.*
```

#### Naming Service

```
org.omg.CosNaming.*
org.omg.CosNaming.NamingContextPackage.*
org.omg.CosNaming.NamingContextExt.*
org.omg.CosNaming.NamingContextExtPackage.*
```

#### J2SE'S Timer and TimerTask

```
java.util.Timer.*
java.util.TimerTask
```

### Supplier Client Initialisation

As with the previous examples, objects and variables are declared which are necessary to connect to the ORB, log, proxy, Naming Service, etc. Note that since this client will be sending structured events using the push model, a reference to a structured push consumer proxy, a *StructuredProxyPushConsumer*, will therefore need to be obtained

```
.
// structured push suppliers must implement
StructuredPushSupplierOperations
public class NotifySupplierB implements
StructuredPushSupplierOperations
{
    org.omg.CORBA.ORB orb = null;
    org.omg.CORBA.Object obj = null;
```



```

NotifyLog log = null;
StructuredProxyPushConsumer pushConsumerProxy = null;

NamingContextExt rootContext = null;
NameComponent notifyLog[];

private Timer sendTimer;

private final long duration = 3 * 60 * 1000; // activity time in
millisecs
private final long interval = 1000; // 1 second

```

Also, as in the previous examples, a reference to the ORB is obtained and the client is instantiated as a transient CORBA object:

```

orb = ObjectAdapter.init(args);

// instantiate this client as a transient CORBA object
ObjectAdapter.createTransient(this);

```

A reference to the desired NotifyLog instance is obtained using the Naming Service, as before:

```

// obtain reference to log using Naming Service
obj = orb.resolve_initial_references ("NameService");
rootContext = NamingContextExtHelper.narrow(obj);
notifyLog = new NameComponent[1];
notifyLog[0] = new NameComponent("NotifyLogB", "log");
log = NotifyLogHelper.narrow(rootContext.resolve(notifyLog));

```

The client will now establish itself as a structured push supplier and connect to the consumer proxy, as shown in the code given below. Note that Notification Service's API is used to do this. Briefly, the procedure accommodates the features of the Notification Service, such as the provision of multiple channels, ability to handle different event types and transmission models (resulting in numerous proxy types), plus provision of Quality of Service. In particular, note that

- The NotifyLog's *default* supplier admin object is obtained (using `default_supplier_admin()`), since NotifyLog can have multiple channels and admin objects, and
- The StructuredProxyPushConsumer which will be used is obtained by narrowing the more general ProxyConsumer proxy and using `ClientType.STRUCTURED_EVENT` to identify which type of proxy will be created.

The example code also sets the `MaxInactivityInterval` for the proxy in order to demonstrate how to set a QoS property for a proxy. Although it is not necessary to set QoS properties when connecting to a proxy, setting this particular QoS property is useful as a mechanism for destroying a proxy when its associated client does not or can not explicitly destroy the proxy itself, such as when the client is terminated by a **Ctrl-C** executed on the command line, i.e. this mechanism can prevent the service filling up with 'dead' proxies.

`MaxInactivityInterval` destroys the proxy when there has been no activity detected by the client after a fixed time interval. The time-out for the interval is reset each time client activity is detected.

```

// establish this object as a structured push supplier
StructuredPushSupplier supplier =

StructuredPushSupplierHelper.narrow(ObjectAdapter.getObject(this));

// obtain the log's structured push consumer proxy and connect to it

```

```

SupplierAdmin supplierAdmin = log.default_supplier_admin();
org.omg.CORBA.IntHolder id = new org.omg.CORBA.IntHolder();

ProxyConsumer proxyConsumer =
    supplierAdmin.obtain_notification_push_consumer(
        ClientType.STRUCTURED_EVENT, id);

pushConsumerProxy =
    StructuredProxyPushConsumerHelper.narrow(proxyConsumer);

// set proxy's MaxInactivityInterval QoS to timeout after
// 20 seconds if this supplier is prematurely terminated
Property[] qos = new Property[1];
qos[0] = new Property();
qos[0].name = MaxInactivityInterval.value;
qos[0].value = orb.create_any();
qos[0].value.insert_ulongLong(10000 * 1000 * 5);

pushConsumerProxy.set_qos(qos);

// connect to the proxy
pushConsumerProxy.connect_structured_push_supplier(supplier);

// enable incoming requests, without blocking, for log callbacks
ObjectAdapter.ready(false);

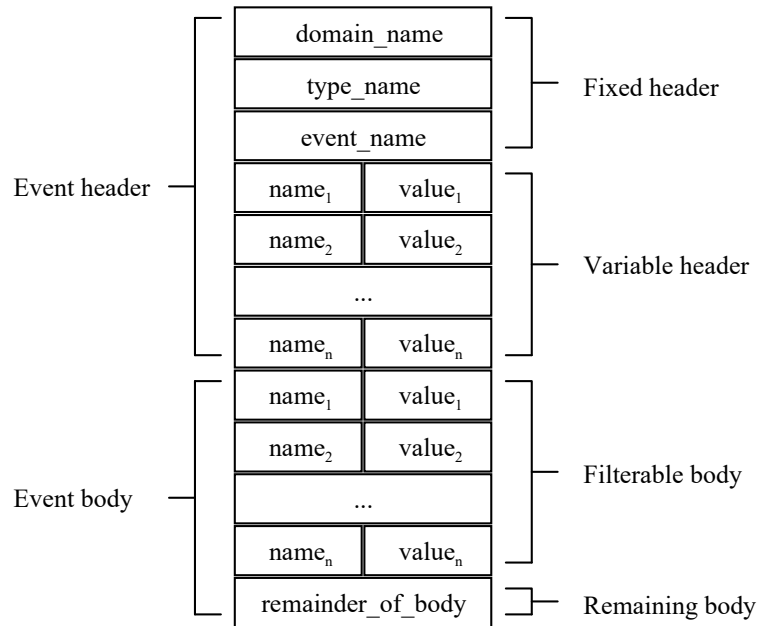
```

After the client is connected to the proxy and the QoS properties set, the ORB is notified that it is able to accept callbacks using `ObjectAdapter.ready(false)`. The supplier is now ready to create and send structured events.

## Structured Events in Brief

A structured event consists of two main parts:

- An *event header* which contains a *fixed header* and *variable header*; these contain:
  - *event domain* (*domain\_name*) - the domain of a particular vertical industry where the event type is defined, such as *telecommunications*, *finance*, *transportation*, etc.
  - *event type* (*type\_name*) - the type of particular event within the domain, for example *StockQuote* within the *finance* domain
  - *event name* (*event\_name*) - a unique name for the particular event instance being transmitted
  - QoS property settings as a list of *name-value* pairs
- An *event body* containing:
  - a *filterable\_body* which is a list of *name-value* pairs which are used to filter the event
  - a *remainder\_of\_body* which is the event itself, an Any



**Figure 6 Structured Event**

## Creating and Sending the Events

The following code example is the client's implementation of the `TimerTask`'s `run()` method. This method creates and sends a series of structured events to the `NotifyLog` instance for a fixed duration (determined by comparing the value returned by `System.currentTimeMillis()` with the value held in `stopSendingTime`).

The code example shows how to:

- Define the domain name ("software"), event type ("example") and event name ("Notify Log") fields of the structured event's fixed header
- Create and define two filterable properties ("number" and "time") in the filterable body header
- Define and add an event to the event's `remainder_of_body`
- Send the event to the proxy

The fixed header will be used by the log filter and consumer client's proxy to identify the event; the properties in the filterable body will be used to filter the event.

### Points to Note

- A check is made, before an event is created and sent, to ensure that the log's record store is *not full* - this is in addition to checking that the log is not *off duty*. This additional check, that the log is not full, is necessary since the particular log instance which will receive the events is configured to reject events when the record store is full: this configuration will cause a `LogFull` exception to be thrown when trying to store an event when the store is full: using `log.get_availability_status().log_full` in the `if()` expression avoids this exception being thrown.
- In order to provide an example mechanism for filtering, a counter is incremented each time an event is sent: if the value of the counter is

even, then a 0 (zero) is assigned to the event's *number* property (held in the event's filterable body); if the counter is odd, a 1 is assigned. The log and consumer will use these values to filter the events:

```
counter++;
num = counter % 2;
...
event.filterable_data[0] = new Property();
event.filterable_data[0].name = "number";
event.filterable_data[0].value = orb.create_any();
event.filterable_data[0].value.insert_long(num);
```

- `pushConsumerProxy.disconnect_structured_push_consumer()` is called when `run()` has stopped sending events; this is used to destroy the proxy (since it is no longer needed).

```
public void run()
{
    // check timeout and record store availability
    if((System.currentTimeMillis() < stopSendingTime) &&
        (log.get_availability_status().off_duty == false) &&
        (log.get_availability_status().log_full == false))
    {
        try
        {
            counter++;
            num = counter % 2; // determine if counter is odd or event
            // create the structured event to be sent to the log
            StructuredEvent event = new StructuredEvent();

            // create header
            EventType type = new EventType("software", "example");
            FixedEventHeader fixedHeader =
                new FixedEventHeader(type, "Notify Log");
            Property variableHeader[] = new Property[0];
            event.header = new EventHeader(fixedHeader, variableHeader);

            // create a filterable body containing properties to hold
            // the event's number and the time it was sent
            event.filterable_data = new Property[2];

            event.filterable_data[0] = new Property();
            event.filterable_data[0].name = "number";
            event.filterable_data[0].value = orb.create_any();
            event.filterable_data[0].value.insert_long(num);

            event.filterable_data[1] = new Property();
            event.filterable_data[1].name = "time";
            event.filterable_data[1].value = orb.create_any();

            // add a message to the body
            event.remainder_of_body = orb.create_any();
            event.remainder_of_body.insert_string("Your Message Here.");

            // add the time that event is sent
            event.filterable_data[1].value.insert_ulonglong(
                System.currentTimeMillis());

            // send the event
            pushConsumerProxy.push_structured_event(event);

            System.out.println ("The number value of the event sent is
" +
                event.filterable_data[0].value.extract_long() + ".");
        }
        catch (org.omg.CosEventComm.Disconnected ex)
        {
            System.out.println("Disconnected");
            System.exit(1);
        }
    }
}
```

```

    }
  }
  else
  {
    if(log.get_availability_status().off_duty)
    {
      System.out.println("The log is off duty " +
        "- cannot log event.");
    }
    if(log.get_availability_status().log_full)
    {
      System.out.println("The log is full " +
        "- cannot log event.");
    }
    if(System.currentTimeMillis() >= stopSendingTime)
    {
      System.out.println("Supplier finished sending events.\n");
    }

    // cleanup - destroy proxy, stop incoming requests
    pushConsumerProxy.disconnect_structured_push_consumer();
    ObjectAdapter.shutdown();

    // stop thread
    sendTimer.cancel();
  }
} // end run

```

## Supplier Client Exceptions

The example code shown above must either throw or catch the exceptions listed below.

### ORB

*org.omg.CORBA.ORBPackage.InvalidName*

### NotifyLog

*org.omg.CosEventChannelAdmin.AlreadyConnected*

*org.omg.CosEventComm.Disconnected*

*org.omg.CosNotifyChannelAdmin.AdminLimitExceeded*

*org.omg.CosNotification.UnsupportedQoS*

### Naming Service

*org.omg.CosNaming.NamingContextPackage.InvalidName*

*org.omg.CosNaming.NamingContextPackage.NotFound*

*org.omg.CosNaming.NamingContextPackage.CannotProceed*

## Creating a Consumer Client for the NotifyLog

The consumer client sets a filter on the proxy which connects it to the log's notify channel, receives events which are allowed to pass through the filter, and retrieves the last ten records held in the log's record store.

This example consumer is similar to the supplier in terms of obtaining a connection to the log, and similar to the log creation example in terms of creating a filter. Since this example builds on previous ones, only the relevant differences will be shown or highlighted.

## Import Statements

The program uses the services and associated APIs listed below. Applications or classes that use these services or APIs must include the import statements shown below.

### OpenFusion Orb Adaptor

*com.prismt.orb.ObjectAdapter*

## NotifyLog and for Time values

```
org.omg.DsLogAdmin.*
org.omg.DsNotifyLogAdmin.*
org.omg.CosNotifyComm.*
org.omg.CosNotifyChannelAdmin.*
org.omg.CosNotification.*

org.omg.CosNotifyFilter.Filter
org.omg.CosNotifyFilter.FilterFactory
org.omg.CosNotifyFilter.ConstraintExp
org.omg.CosNotifyFilter.InvalidConstraint
org.omg.CosNotifyFilter.InvalidGrammar

java.util.GregorianCalendar
java.util.Date
```

## Naming Service

```
org.omg.CosNaming.*
org.omg.CosNaming.NamingContextPackage.*
org.omg.CosNaming.NamingContextExt.*
org.omg.CosNaming.NamingContextExtPackage.*
```

## Consumer Client Initialisation

The consumer performs similar initialisation steps to the previous examples, and as shown in the code examples below. Note that the consumer uses a proxy of type `StructuredProxyPushSupplier` (which is obtained by narrowing a `ProxySupplier` with a client type of `STRUCTURED_EVENT`).

```
org.omg.CORBA.ORB orb = null;
org.omg.CORBA.Object obj = null;

NotifyLog log = null;

StructuredProxyPushSupplier pushSupplierProxy = null;
Filter filter = null;

org.omg.CORBA.Any event;

NamingContextExt rootContext = null;
NameComponent notifyLog[];

private final long duration = 3 * 60 * 1000; // activity time in
millisecs

orb = ObjectAdapter.init(args);

// instantiate this client as a transient CORBA object
ObjectAdapter.createTransient (this);

// obtain reference to log using Naming Service
obj = orb.resolve_initial_references ("NameService");
rootContext = NamingContextExtHelper.narrow(obj);
notifyLog = new NameComponent[1];
notifyLog[0] = new NameComponent("NotifyLogB", "log");
log = NotifyLogHelper.narrow(rootContext.resolve(notifyLog));

// establish this object as a structured push consumer
StructuredPushConsumer consumer =
StructuredPushConsumerHelper.narrow(ObjectAdapter.getObject(this));

// obtain the log's structured push supplier proxy and connect to it
ConsumerAdmin consumerAdmin = log.default_consumer_admin();
```

```

org.omg.CORBA.IntHolder id = new org.omg.CORBA.IntHolder();

ProxySupplier proxySupplier =
    consumerAdmin.obtain_notification_push_supplier(
        ClientType.STRUCTURED_EVENT, id);

pushSupplierProxy =
    StructuredProxyPushSupplierHelper.narrow(proxySupplier);

```

## Setting a Proxy Filter

After obtaining the proxy, the client creates and adds a filter to it. This filter is configured to only accept events which have a property called *number* (in the filterable body) which has a value of *1* (i.e. odd numbered events - see "Points to Note" on page 49).

### Points to Note

- A filter factory is used to create the filter; this factory is obtained from the log.
- The *domain name* and *type name* of the events to be filtered are used to create an `EventType` object, called *type* in the example. The `EventType` is added to an `EventType` array, which is used by a `ConstraintExp` object to identify which events to filter. A constraint expression, also used by the `ConstraintExp` object, defines the type of *header*, *property name*, and *condition* used to test the value of the property.
- The `add_filter()` method is used to add the filter to the proxy, and *not* the `set_filter()` method which is used for an log-based filter

```

try
{
    // set filter on proxy to only accept events where
    // the "number" property is an odd number (value of 1)
    FilterFactory filterFactory = log.default_filter_factory();
    filter = filterFactory.create_filter(default_grammar.value);

    EventType type = new EventType("software", "example");
    EventType[] types = {type};
    ConstraintExp exp =
        new ConstraintExp(types, ".$filterable_data(number) == 1");

    ConstraintExp[] expressions = {exp};
    filter.add_constraints(expressions);
    pushSupplierProxy.add_filter(filter);
}
catch (org.omg.CosNotifyFilter.InvalidGrammar ex)
{
    System.out.println("InvalidGrammar");
    System.exit(1);
}
catch (org.omg.CosNotifyFilter.InvalidConstraint ex)
{
    System.out.println("InvalidConstraint");
    System.exit(1);
}

```

The client is now ready to receive events after having configured and added the filter to the proxy.

## Receiving the Structured Events

The NotifyLog consumer receives events similarly to the previous consumer example that used event-style events, noting the differences in the proxy type and version of connection methods.

```
// connect to the proxy
pushSupplierProxy.connect_structured_push_consumer(consumer);
```

The client also ensures that the forwarding state for the log is enabled.

```
// ensure that log is enabled to forward records
ForwardingState state = log.get_forwarding_state();
if (state == ForwardingState.off)
{
    log.set_forwarding_state(ForwardingState.on);
}

// thread for receiving events
// events are received via push_structured_event(),
// the PushConsumerOperations callback method
System.out.println("Will receive structured events for " +
    duration / 60000 + " minutes from now.");
synchronized(this)
{
    try
    {
        this.wait(duration);
    }
    catch(java.lang.InterruptedException ex)
    {
        return;
    }
}

System.out.println("Finished getting structured events.");
```

The *push\_structured\_event* callback method is responsible for processing any events which are received from the supplier via the log's notify channel. The example implementation extracts the data values held in the filterable data part of received events, as well as extracting the data held in the *remainder\_of\_body*. The only events which should be received by *push\_structured\_event* will be those whose *number* property has a value of 1.

```
// StructuredPushConsumerOperations callback methods
public void disconnect_structured_push_consumer()
{
    System.out.println ("Disconnected by proxy");
} // end disconnect_push_consumer
```

```
public void push_structured_event(
    org.omg.CosNotification.StructuredEvent event)
{
    Property[] filterable_data;
    org.omg.CORBA.Any remainder_of_body = orb.create_any();

    filterable_data = event.filterable_data;
    int len = filterable_data.length ;

    System.out.println("\nNotifyConsumer received event containing:");
    for( int j = 0 ; j < len ; j++ )
    {
        if(filterable_data[j].name.equals("time"))
        {
            System.out.println("    filterable_data[" + j + "] " +
                "time: " +
                new
                Date((long)filterable_data[j].value.extract_ulonglong()).
```



```

        toString());
    } else
    {
        System.out.println("    filterable_data[" + j + "]" + " +
            filterable_data[j].name + ": " +
            filterable_data[j].value);
    }
}

remainder_of_body = event.remainder_of_body ;
System.out.println("    remainder_of_body = " + remainder_of_body);
} // end push

```

```

public void offer_change(EventType[] added, EventType[] removed)
{
    System.out.println ("Consumer offer changed");
} // end offer_change

```

## Retrieving Log Records

The client *retrieves* the log records as structured events, in addition to *receiving* the events sent through the log's channel. The following example shows how the structured event records are retrieved and information extracted from them. The value of the filterable body's (*filterable\_data*) *number* property should be 0, since the filter for the log's record store was set to only accept events containing this value in the *number* property.

```

// array for retrieved records
LogRecord[] records;
IteratorHolder iter = new IteratorHolder();

// establish current time and convert ms to 100 ns
GregorianCalendar calendar = new GregorianCalendar();
long basetime = - calendar.getGregorianChange().getTime();
long currentTime = (System.currentTimeMillis () + basetime) * 10000;

System.out.println(
    "\nRetrieving the latest 10 records backwards from the current
time.");
records = log.retrieve(currentTime, -10, iter);
if (records.length < 1 )
{
    System.out.println ("No records retrieved.\n");
}
else
{
    // display information for each retrieved record
    System.out.println ("\n-----" +
        "Retrieved records listed below.");
    for(int i = 0; i < records.length; i++)
    {
        org.omg.CORBA.Any info = orb.create_any();
        StructuredEvent se;
        Property[] filterable_data;
        org.omg.CORBA.Any remainder_of_body = orb.create_any();

        se = StructuredEventHelper.extract(records[i].info );

        filterable_data = se.filterable_data;
        int len = filterable_data.length ;

        System.out.println("\nrecord id: " +records[i].id);

        for( int j = 0 ; j < len ; j++ )
        {
            if(filterable_data[j].name.equals("time"))
            {
                System.out.println("    filterable_data[" + j + "]" + " +
                    "time: " +

```

```

        new
Date((long)filterable_data[j].value.extract_ulonglong()).
    toString());
    } else
    {
        System.out.println("    filterable_data[" + j + "] " +
            filterable_data[j].name + ", " +
            filterable_data[j].value);
    }
}
remainder_of_body = se.remainder_of_body ;
System.out.println( "    remainder_of_body = " +
    remainder_of_body);
}
System.out.println ("-----");
}
}

```

## Cleanup

When the client is finished receiving events, it destroys the proxy's filter, disconnects from (and destroys) the proxy, then notifies the ORB it is unable to accept requests.

```

public void cleanup()
{
    filter.destroy();
    pushSupplierProxy.disconnect_structured_push_supplier();
    ObjectAdapter.shutdown();
}

```

## Consumer Client Exceptions

The example code shown above must either throw or catch the exceptions listed below.

### ORB

*org.omg.CORBA.ORBPackage.InvalidName*

### NotifyLog

*org.omg.CosEventChannelAdmin.AlreadyConnected*  
*org.omg.CosNotifyChannelAdmin.AdminLimitExceeded*  
*org.omg.CosEventChannelAdmin.TypeError*

### Naming Service

*org.omg.CosNaming.NamingContextPackage.InvalidName*  
*org.omg.CosNaming.NamingContextPackage.NotFound*  
*org.omg.CosNaming.NamingContextPackage.CannotProceed*

## Monitoring Log Generated Events

The following example monitor client shows how events which have been generated by a log and its log factory can be received and monitored.

The client is similar to a standard NotifyLog consumer client, but with the following notable differences:

- An *event* channel, as opposed to a *notify* channel, is used to transmit and receive the events (note the differences of their relative import statements - see ["Import Statements" on page 51](#) and ["Import Statements" on page 57](#)): generated events are transmitted as Event-style events (Anys) using an event channel.
- The client connects to the log *factory's* proxy (instead of the log): all generated events, whether from the factory or the log, are transmitted by the factory that created the log.

- The client does not need to obtain a reference to the log object, but to the factory that creates it. Therefore, there is no need to use the Naming Service or other mechanism to locate the log object. In fact, a client may likely be instantiated and active before the log is created: this will occur if the client intends to monitor log creation events generated by the log factory.

## Import Statements

The program uses the services and associated APIs listed below. Applications or classes that use these services or APIs must include the import statements shown below.

### OpenFusion Orb Adaptor

```
com.prismt.orb.ObjectAdapter
```

### NotifyLog and for Time values

```
org.omg.DsLogNotification.*
```

```
org.omg.DsLogAdmin.*
```

```
org.omg.DsNotifyLogAdmin.*
```

```
org.omg.CosNotification.EventType;
```

```
import org.omg.CosNotifyComm.*;
```

```
import org.omg.CosNotifyChannelAdmin.*
```

```
java.util.GregorianCalendar
```

```
java.util.Date
```

## Monitor Client Initialisation

The client implements `PushConsumerOperations` and is initialised similarly to previous `NotifyLog` client, but note that a reference is obtained to the log factory and not to a log.

```
orb = ObjectAdapter.init(args);

// instantiate this client as a transient CORBA object
ObjectAdapter.createTransient(this);

// obtain reference to log factory
obj = orb.resolve_initial_references("NotifyLogFactory");

NotifyLogFactory factory = NotifyLogFactoryHelper.narrow(obj);
if(factory == null)
{
    System.out.println("\nLog factory not available\n");
    System.exit(1);
}

// establish this object as a push consumer
PushConsumer consumer =
    PushConsumerHelper.narrow(ObjectAdapter.getObject(this));

// obtain the log factory's push supplier proxy and connect to it
org.omg.CORBA.IntHolder id = new org.omg.CORBA.IntHolder();
ProxySupplier proxy =
    factory.obtain_notification_push_supplier(ClientType.ANY_EVENT,
id);
pushSupplierProxy = ProxyPushSupplierHelper.narrow(proxy);
pushSupplierProxy.connect_any_push_consumer(consumer);

// set base time for date and time conversion from TimeT
basetime = - new
GregorianCalendar().getGregorianChange().getTime();
```

## Receiving the Generated Events

The `push()` callback method receives the generated events. The example implementation of this method, shown below, extracts and displays the property values for the various generated events.

```
// StructuredPushConsumerOperations callback methods
public void disconnect_structured_push_consumer()
{
    System.out.println ("Disconnected by proxy");
} // end disconnect_push_consumer
```

An important aspect of retrieving generated events is that it may not be possible to determine the property *type* for a particular event in advance, since certain events can contain different types depending on the situation which generated the event. For example, `AttributeValueChange` can contain different types, e.g. a `long` or `omg.org::CORBA::Object`, depending on which attribute has changed, as shown in the output generated by the above code:

```
AttributeValueChange
  id:          10
  time:        Wed Oct 30 13:35:40 GMT 2002
  type:        8
  value type:  long

AttributeValueChange
  id:          10
  time:        Wed Oct 30 13:35:41 GMT 2002
  type:        6
  value type:  omg.org::CORBA::Object
```

## Monitor Client Exceptions

The example code shown above must either throw or catch the exceptions listed below.

### ORB

`org.omg.CORBA.ORBPackage.InvalidName`

### NotifyLog

`org.omg.CosEventChannelAdmin.AlreadyConnected`

`org.omg.CosNotifyChannelAdmin.AdminLimitExceeded`

`org.omg.CosEventChannelAdmin.TypeError`

# Supplemental Information

## Exceptions

A brief description of the Log Service exceptions is provided in [Table 3](#) below. These exceptions are defined in the *DsLogAdmin* module.

*i* However, there are additional exceptions for those log types which inherit from the Event and Notification Services. These exceptions are described in the **OpenFusion Notification Service Guide**.

**Table 3 Log Service Exceptions**

Exception	Reason
<i>InvalidParam</i>	Invalid parameter supplied.
<i>InvalidThreshold</i>	Invalid threshold value supplied when setting an alarm threshold. Valid values are expressed as percentages between 0 and 100, non-inclusive.
<i>InvalidTime</i>	The <i>TimeInterval's</i> or <i>Time24Interval's</i> <i>stop</i> time is set earlier than the current time, or when a <i>Time24</i> structure contains invalid values.
<i>InvalidTimeInterval</i>	The <i>start</i> time is set later than the <i>stop</i> time.
<i>InvalidMask</i>	The days bit field of a <i>WeekMaskItem</i> contains an illegal value.
<i>LogIdAlreadyExists</i>	Raised when creating a new log and a log with the same id already exists.
<i>InvalidGrammar</i>	The grammar specified is not supported.
<i>InvalidConstraint</i>	The supplied constraint does not conform to the specified grammar and is invalid.
<i>LogFull</i>	An attempt is made to store a record when the record store is full and the log full behaviour is set to <i>halt</i> .
<i>LogOffDuty</i>	An attempt is made to store a record when the log is off duty.
<i>LogLocked</i>	An attempt is made to store a record when the log is locked.
<i>LogDisabled</i>	An attempt is made to store a record when the log is disabled.
<i>InvalidRecordId</i>	An attempt is made to retrieve or delete a log record that does not exist.
<i>InvalidAttribute</i>	The log attribute is invalid.
<i>InvalidLogFullAction</i>	The specified log full action is not supported.
<i>UnsupportedQoS</i>	The specified QoS property is not supported or invalid.



# Part II

## Configuration and Management

### **In this part**

This part contains the following:

<a href="#">Log Service Configuration</a>	<a href="#">page 63</a>
<a href="#">Log Service Manager</a>	<a href="#">page 71</a>





# Log Service Configuration

## Overview

The configuration of Singleton properties specific to the Log Service is described in this section. These properties appear in the Administration Manager, a graphical user interface (GUI) based administration tool included with the OpenFusion Graphical Tools.

The Administration Manager can be used to set the Singleton properties. These properties can also be set programatically, generally as described in the service description sections.

Details for configuring Persistence, Logging, CORBA, Java and System properties for the Log Service are described in the **System Guide**.

## Common Properties

Instances of some common properties are used by a number of different OpenFusion Log Service' interfaces and services. Settings for these property instances appear in the Administration Manager's Object Hierarchy for the service's Singleton node. This small group of properties is included in this section in order to facilitate configuration of the Service while using the Administration Manager. These properties include:

- IOR Name Service Entry
- IOR URL
- IOR File Name
- Resolve Name
- IOR Name Service

## LogFactorySingleton Configuration

This section lists the properties for each of the Log Factory Singletons:

- **BasicLogFactorySingleton**
- **NotifyLogFactorySingleton**

Each Singleton must be configured separately but they have identical properties as defined in this section.

## CORBA Properties

### IOR Name Service Entry

The Naming Service entry for the Singleton.

<b>Property Name</b>	Object.Name
<b>Property Type</b>	FIXED
<b>Data Type</b>	STRING
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## IOR URL

The **IOR URL** property specifies the location of an Interoperable Object Reference (IOR) for the Service, using the Universal Resource Locator (URL) format. This information is used when a client attempts to resolve a reference to the Service. Some examples are:

```
file:/usr/users/openfusion/servers/NameService.ior
http://www.microfocus.com/of/servers/NameService.ior
corbaloc::server.microfocus.com/NameService
```

OpenFusion supports URLs in *Corbaloc*, *Corbaname*, *file*, *FTP* and *HTTP URL* formats, although some ORBs do not support all of these mechanisms. Consult your ORB documentation for specific details.

<b>Property Name</b>	IOR.URL
<b>Property Type</b>	FIXED
<b>Data Type</b>	URL
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## IOR File Name

The **IOR File Name** option specifies the name and location of the IOR file for the Singleton. If this property is not set, the IOR file name will be:

```
<INSTALL>/domains/<domain>/<node>/<service>/<singleton>/
<singleton>.ior
```

where **<INSTALL>** is the OpenFusion installation path. See the **System Guide** for details of the `domains` directory structure.

<b>Property Name</b>	IOR.File
<b>Property Type</b>	FIXED
<b>Data Type</b>	FILE
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## Resolve Name

The ORB Service resolution name used to resolve calls to the Singleton.

<b>Property Name</b>	ResolveName
<b>Property Type</b>	FIXED
<b>Data Type</b>	STRING
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	YES

## IOR Name Service

The name of the Naming Service which will be used to resolve the Singleton object.

<b>Property Name</b>	IOR.Server
<b>Property Type</b>	FIXED
<b>Data Type</b>	STRING
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## Persistence Properties

### Enable Write Ahead Log

When the write-ahead log is enabled, information that is normally written to the underlying database is written to a log file instead. When the log file reaches a specific size (defined by the **Write Ahead Log Maximum Size** property), the database is updated and the log file is reused. The location of the log file is defined by the **Write Ahead Log Directory** property and must be held locally on the machine running the Service.

The write-ahead log may increase performance when persistent events are required, particularly when events are being delivered quickly.

The write-ahead log is enabled when this property is set `TRUE` (checked). By default, the write-ahead log is disabled. It can be safely enabled unless your system is not set up to allow files to be written to the local machine.

<b>Property Name</b>	DB.WAL
<b>Property Type</b>	STATIC
<b>Data Type</b>	BOOLEAN
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

### Write Ahead Log Directory

The directory used to contain write-ahead log files. This directory must be local to the host running the Service. The default location is:

```
<INSTALL>/domains/<domain>/<node>/LogService/data
```

where `<INSTALL>` is the OpenFusion installation path. See the **System Guide** for details of the `domains` directory structure.

<b>Property Name</b>	DB.WAL.Dir
<b>Property Type</b>	STATIC
<b>Data Type</b>	DIRECTORY
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	YES

## Write Ahead Log Maximum Size

The maximum number of entries that can be stored in the write-ahead log before flushing (writing to the underlying database) takes place.

<b>Property Name</b>	DB.WAL.MaxSize
<b>Property Type</b>	STATIC
<b>Data Type</b>	INTEGER
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## Database Plugin Class

The database plugin class used for persistent storage of logs. If a custom log store plugin is used, the class should be specified here.

If this property is left blank, the default OpenFusion database plugin implementations will be used.

<b>Property Name</b>	DB.Plugin
<b>Property Type</b>	STATIC
<b>Data Type</b>	STRING
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## General Properties

### Enable Event Queuing

Indicates if a queue should be used before the events have been sent to persistent storage. This property can accelerate the speed of saving events when using slow storage.

<b>Property Name</b>	Queue.Enable
<b>Property Type</b>	STATIC
<b>Data Type</b>	BOOLEAN
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

### File Queuing Directory

Indicates the directory where the queue is located.

<b>Property Name</b>	Queue.Dir
<b>Property Type</b>	STATIC
<b>Data Type</b>	DIRECTORY
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## Initial Number of Logs

The specific properties for each Singleton are:

- **Initial Number of Basic Logs**
- **Initial Number of Notify Logs**

This property determines the initial number of logs (of the appropriate type) created when the Service is run for the first time.

<b>Property Name</b>	InitialLogs
<b>Property Type</b>	STATIC
<b>Data Type</b>	INTEGER
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## Maximum Number of Queue Files

Indicates how many files can be used to save the queued events where the *maximum number of queue file* multiplied by the *maximum queue file size* is the maximum size that the queue can use.

<b>Property Name</b>	Queue.MaxFiles
<b>Property Type</b>	STATIC
<b>Data Type</b>	INTEGER
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## Maximum Queue File Size

The maximum size of the file which the queued events are saved to.

<b>Property Name</b>	Queue.MaxFileSize
<b>Property Type</b>	STATIC
<b>Data Type</b>	INTEGER
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## Max Records Returned From Query

Sets the maximum number of records that will be in the RecordList returned by a query operation. The default value is 100.

If the query returns more records than the maximum set for the RecordList, the additional records can be obtained using the Iterator.

<b>Property Name</b>	MaxRecordsReturnedFromQuery
<b>Property Type</b>	STATIC
<b>Data Type</b>	INTEGER
<b>Accessibility</b>	READ-WRITE
<b>Mandatory</b>	NO

## Order query results

When set to `TRUE` (checked), the records returned by a query will be ordered by time, earliest records first. By default, this property is `FALSE` and query results are unordered.

Only the records returned in the `RecordList` will be ordered in this way. Additional records returned by the `Iterator` may or may not be ordered. (The **Max records returned from query** property, above, sets how many records will be returned in the `RecordList`.)



Ordering is performed in memory so if a large number of records are returned, a correspondingly large amount of memory will be used.

<b>Property Name</b>	<code>OrderQueryResults</code>
<b>Property Type</b>	<code>STATIC</code>
<b>Data Type</b>	<code>BOOLEAN</code>
<b>Accessibility</b>	<code>READ-WRITE</code>
<b>Mandatory</b>	<code>OPTIONAL</code>

## ProcessSingleton Configuration

### IOR Name Service Entry

The Naming Service entry for the Singleton.

<b>Property Name</b>	<code>Object.Name</code>
<b>Property Type</b>	<code>FIXED</code>
<b>Data Type</b>	<code>STRING</code>
<b>Accessibility</b>	<code>READ/WRITE</code>
<b>Mandatory</b>	<code>NO</code>

### IOR URL

The **IOR URL** property specifies the location of an Interoperable Object Reference (IOR) for the Service, using the Universal Resource Locator (URL) format. This information is used when a client attempts to resolve a reference to the Service. Currently only *http* and *file* URLs are supported, for example:

```
file:/usr/users/openfusion/ProcessSingleton.iior
```

```
http://www.abigbank.com/openfusion/ProcessSingleton.iior
```

<b>Property Name</b>	<code>IOR.URL</code>
<b>Property Type</b>	<code>FIXED</code>
<b>Data Type</b>	<code>URL</code>
<b>Accessibility</b>	<code>READ/WRITE</code>
<b>Mandatory</b>	<code>NO</code>

## IOR File Name

The **IOR File Name** option specifies the name and location of the IOR file for the Singleton. If this property is not set, the IOR file name will be:

```
<INSTALL>/domains/<domain>/<node>/<service>/<singleton>/  
<singleton>.ior
```

where <INSTALL> is the OpenFusion installation path. See the **System Guide** for details of the domains directory structure.

<b>Property Name</b>	IOR.File
<b>Property Type</b>	FIXED
<b>Data Type</b>	FILE
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO

## IOR Name Service

The name of the Naming Service which will be used to resolve the Singleton object.

<b>Property Name</b>	IOR.Server
<b>Property Type</b>	FIXED
<b>Data Type</b>	STRING
<b>Accessibility</b>	READ/WRITE
<b>Mandatory</b>	NO





# Log Service Manager

## Overview

The following tasks can be carried out from the Log Service Manager:

- Create Log objects.
- Configure Log objects.
- Run queries that return specified Log records.
- View, edit, and delete Log records.
- Add and remove attributes belonging to a Log record.

## Using the Log Service Manager

To start the Log Service Manager, right-click either the (**BasicLogFactorySingleton**, or **NotifyLogFactorySingleton** in the **Object Hierarchy** and select **Log Manager** from the pop-up menu. See [“Administration Manager”](#) for details.

The Log Service Manager is shown in [Figure 7](#).

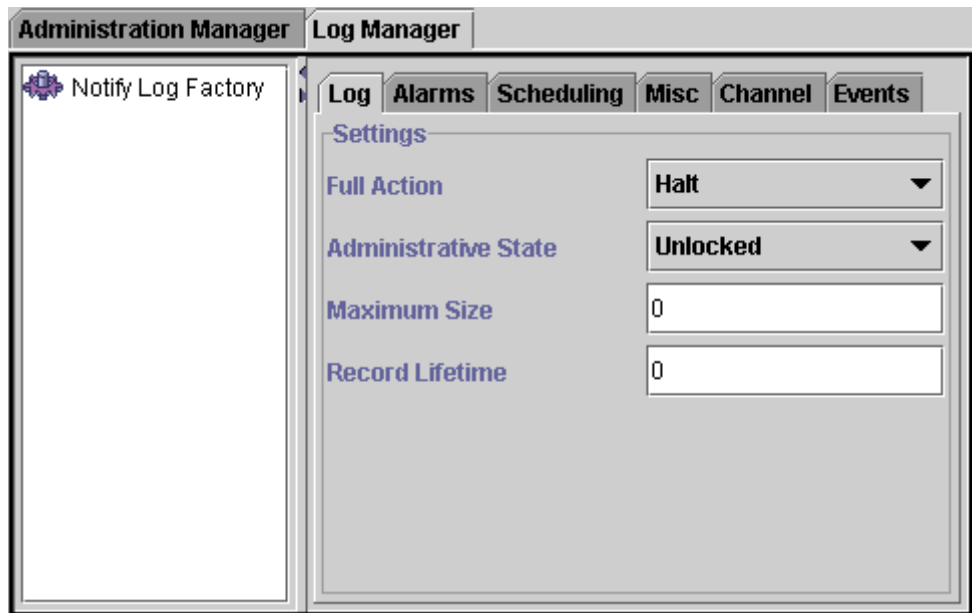




Figure 7 Log Service Manager

Nodes in the Log Service Manager are identified by different icons. These icons are shown in [Table 4](#).

**Table 4 Log Service Object Icons**

Icon	Node
	<p><b>Log Factory</b></p> <p>The root node of the Log Service hierarchy. The description next to this icon indicates which Log Factory is being browsed.</p>
	<p><b>Log</b></p> <p>A representation of a Log object.</p>

## Log Object Settings

The Log Manager has six tabbed panels with settings configuring Log objects and capturing Log Service events. Not all Log Managers have all six panels. These panels are described in the following sections. All settings must be correctly configured before a new Log object can be created (see [“Create a New Log”](#)).

### Log

The **Log** panel is used by all Log Factories. The following properties must be set on the **Log** panel:

#### Full Action

The action to be taken when a log is full. Available actions are:

- **Halt** - No further log records will be written to the persistent log record store. The administrative state of the log will be changed to Locked.
- **Wrap** - The oldest logs are deleted to make space for the new ones.

#### Administrative State

This can be either:

- **Locked** - The Log Service will not write any log records to the persistent store.
- **Unlocked** - The Log Service can write log records to the persistent store.

#### Maximum Size

The maximum size of the log, in bytes, kilobytes, or Megabytes.

To specify a size in kilobytes, append `kb` to the number. To specify a size in Megabytes, append `Mb` to the number. For example: `100Kb`, `100Mb`. If no suffix is appended, a value in bytes is assumed. Note that case is not significant, so `KB`, `kb`, and `Kb` are all valid suffixes.

A value of `zero` means that there is no maximum limit to the log size.

Regardless of what units are specified for the property, they will be converted to bytes and displayed as such.

#### Record Lifetime

The length of time, in seconds, minutes, hours, or days, that a log record will exist in the log store before it is automatically deleted.

To specify a lifetime in minutes, append `m` to the number. To specify a lifetime in hours, append `h` to the number. To specify a lifetime in days, append `d` to the number. If no suffix is appended, a value in seconds is

assumed. Note that case is not significant, so `D` and `d` are both valid suffixes.

A value of `zero` means that log records never expire.

Regardless of what units are specified for the property, they will be converted to seconds and displayed as such.

## Alarms

The **Alarms** panel is used by all Log Factories.

The **Alarms** panel allows unique threshold values to be added to logs. An alarm threshold must be an integer value between 1 and 100. If several thresholds are entered, each must be a different value.

To add a threshold, enter the value in the **Enter value** field and click the **Add** button.

To remove a threshold, select the threshold in the list and click the **Remove** button.

To set the range of thresholds as the defaults for all logs, click the **Set As Defaults** button. All values are stored together as a set of defaults.

To re-use previously set default values, click the **Use Defaults** button. Note that it is possible to restore previous defaults and then add or delete thresholds in this instance without altering the default set (although changes could be saved as new defaults).

## Scheduling

The **Scheduling** panel is used by all Log Factories. These settings allow a schedule to be set up for starting and stopping logging activity.

### Current setting

The current setting can be:

- **Default**
- **Custom**

If the **Default** setting is selected, the default **Start** time of `Immediately` and **Stop** time of `Never` are used, and these settings cannot be changed. To manually configure the schedule, select the **Custom** setting.

### Start and Stop

If the **Custom** setting is selected, the **Start** and **Stop** times must be manually configured. Enter the required times and click the **Apply** button.

### Week Mask

For fine-grained scheduling, a week mask can be set up. This allows different start and stop times to be set for different days of the week.

## Miscellaneous

These settings are used only by the Notify Log Factory.

### Forwarding State

Select whether log forwarding is **On** or **Off**:

- **On**: The log will forward incoming events from all suppliers to all consumers currently connected.
- **Off**: The log will not forward incoming events from suppliers to any consumers.

## Filter Constraints

Filter constraints are only available for Notify and TypedNotify Logs.

To enter a filter constraint, type the filter expression into the **Value** field and click the **Add** button.

To delete a filter constraint, select the entry in the list and click the **Remove** button.

To edit a filter constraint, select the constraint in the list so that it is displayed in the **Value** field. Edit the constraint value as required, then click the **Replace** button. The new constraint replaces the previous one in the list. If **Add** is clicked instead of **Replace**, the modified constraint is added to the list as a new constraint and the existing constraint will remain in the list.

Filter constraints added here will be applied to any future logs that are created with the Log Service Manager. The filter will not be applied to any logs created programmatically or to any logs which were created in the Log Service Manager *before* the filter was added.

To set a filter constraint on an existing log, select the log in the tree view. Select the **Settings** tab and the **Misc** tab, and add the filter constraint as described above. This filter will apply to the selected log only.

## Channel

These settings are used only by the Notify Log Factory. This panel allows Quality of Service (QoS) and Administrative settings to be set for the log. The following properties can be set:

### QoS Settings

- AcknowledgeLevel
- AcknowledgeMode
- ConnectionReliability
- DeadLetterDrop
- DiscardPolicy
- DisconnectCallback
- EventReliability
- EventIDSupported
- LazyAcknowledgeInterval
- MaximumBatchSize
- MaxEventsPerConsumer
- MaxInactivityInterval
- MaxReconnectEvents
- MultipleReceiversSupported
- OrderPolicy
- PacingInterval
- Priority
- ReconnectInterval
- StartTimeSupported
- StopTimeSupported
- ThreadIdleTime
- ThreadPoolSize
- Timeout

### Administrative Properties

- MaxQueueLength
- MaxConsumers
- MaxSuppliers
- RejectNewEvents

## Events

This panel is used only by the Notify Log Factory and displays information about the events stored in the Log.

The following details are displayed:

- **Event**
- **Time**
- **Log**
- **Details**

The types of events generated by Log Factories and Logs are as follows:

- **StateChange**: Generated when a Log's admin or operational state is changed.
- **ObjectCreation**: Generated when a Log is created.
- **ObjectDeletion**: Generated when a Log is deleted.
- **ThresholdAlarm**: Generated when a Log threshold has been exceeded.
- **AttributeValueChange**: Generated when any of the following Log attributes are changed:
  - Capacity alarm threshold
  - Log full action
  - Max Log size
  - Start time
  - Stop time
  - Week mask
  - Filter
  - Max record life
  - Quality of Service (QoS)

## Create a New Log

To create a new Log object, right-click on the Log Factory root node and select **Create** from the pop-up menu.

This will create a new Log object as a child of the root node, but only if all properties have been properly configured. (See "[Log Object Settings](#)".)

The **Create With ID** option does the same thing as the **Create** option but prompts for a Log ID. A valid unique ID must be entered for the Log object to be created.

The Log node in the Log Factory hierarchy displays the Log's unique number and description. To view information about a log, select the log in the hierarchy. Details are displayed in the right-hand pane, as shown in [Figure 8](#).

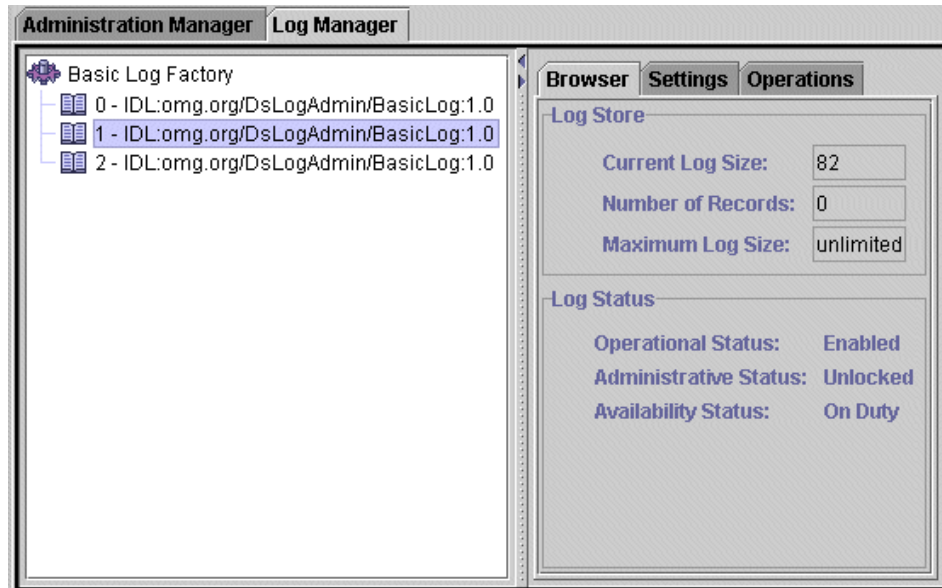


Figure 8 Viewing Log Details

## Copy Logs

To make a copy of the Log object, right-click on the Log object in the hierarchy and select **Copy** from the pop-up menu.

The new Log object is added to the Log Factory hierarchy and retains all the configuration settings of the original.

The **Copy with ID** option does the same thing as the **Copy** option but prompts for a Log ID. A valid unique ID must be entered for the copy of the Log object to be created. The ID cannot be the same as the original Log.

## Destroy Logs

To remove a Log object from the Log Factory hierarchy and from the Log Service, right-click the Log object and select **Destroy** from the pop-up menu.

## Browser

The **Browser** panel displays information about the current status of the Log Service. The following details are displayed:

- **Current Log Size.**
- **Number of Records.**
- **Maximum Log Size.**
- The percentage of the allocated Log **Capacity** currently in use (also summarised in a pie chart).
- **Operational Status** (Enabled or Disabled).
- **Administrative Status** (Locked or Unlocked).
- **Availability Status.**
- Number of **Connected Consumers.**
- Number of **Connected Suppliers.**

For details of these properties, see [“Log Object Settings”](#).

## Settings

The **Settings** panel allows Log object properties to be changed. The panel contains properties on five tabbed panels, corresponding to the properties that can be set when a new Log is created, as follows:

- **Log** (page 72).
- **Alarms** (page 73).
- **Scheduling** (page 73).
- **Misc** (page 73).
- **Channel** (page 74).

Changes to a property in a text box are **not** applied to the Log object until focus is moved to another property *in the same panel*. Navigating to another panel without doing this will cause the change to be lost. (This does not apply to properties which are set with a check box or drop-down list.)

## Operations

The **Operations** panel is used to retrieve and display selected Log records.

Five different retrieval actions can be performed on the **Operations** panel. Select which action to perform from the **Operation** drop-down list:

- **Query**
- **Retrieve**
- **Match**
- **Delete**
- **Set Attributes**

### Query

The **Query** operation selects and displays records that match a query string.

- 1 Use the **Grammar** drop-down list to select the query language, `EXTENDED_TCL` (ETCL). Refer to the Notification Service documentation for details of ETCL.

The database plugin used for Log persistence may restrict the choice of grammar.

- 2 Enter the query string in the **Constraint** field.
- 3 Click the **execute** button to run the query. Each returned record is displayed in the **Record Details** table.
- 4 Select a record to display the record's attributes in the **Attribute Details** table.

Records can be individually deleted by selecting them and using the **Delete** button. Deleted records are removed from the Log Service.

Attributes can be added or deleted from the record using the **Add** and **Remove** buttons.

### Retrieve

The **Retrieve** operation selects and displays a specified number of log records.

- Enter a positive number in the **Records to return** field to return records starting at the specified time.
- Enter a negative number in the **Records to return** field to return records immediately prior to the specified time.

For example, with 30 March 2001 12:00 as the date:

- **Records to return** - 50 will retrieve the first 50 records logged after 12:00 on 30 March 2001.
- **Records to return** -50 will retrieve the last 50 records logged before 12:00 on 30 March 2001.

Each returned record is displayed in the **Record Details** table. Records can be individually deleted by selecting them and using the **Delete** button. Deleted records are removed from the Log Service.

Select a record to display the record's attributes in the **Attribute Details** table. Attributes can be added or deleted from the record using the **Add** and **Remove** buttons.

## Match

The **Match** operation is similar to the **Query** operation but instead of returning a list of Log messages, it returns a count of messages.

## Delete

The **Delete** operation deletes all records that match a query string.

- 1 Use the **Grammar** drop-down list to select the query language, `EXTENDED_TCL` (ETCL). Refer to the Notification Service documentation for details of ETCL.

The database plugin used for Log persistence may restrict the choice of grammar.

- 2 Enter the query string in the **Constraint** field.
- 3 Click the **execute** button to delete all matching Log records.

A pop-up message box will report the number of records deleted by the query.



**Warning:** deleted Log records **can not** be recovered.

## Set Attributes

The **Set Attributes** operation allows a batch update of attributes for all Log records which match a query string.

- 1 Click the **Add** button to add an attribute to the **Attribute Details** table.
- 2 Enter a **Name**, **Value**, and **Type** for the attribute.
- 3 Repeat steps 1 and 2 to add more attributes to the table.
- 4 Use the **Grammar** drop-down list to select the query language, `EXTENDED_TCL` (ETCL). Refer to the Notification Service documentation for details of ETCL.

The database plugin used for Log persistence may restrict the choice of grammar.

- 5 Enter the query string in the **Constraint** field.
- 6 Click the **execute** button to run the query.

Each Log record retrieved by the query is automatically updated to include the attributes added to the **Attribute Details** table.

A pop-up message box will report the number of records updated by the query.



**Warning:** this change **can not** be undone.



# Part III

## Appendices

### In this part

This part contains the following:

<a href="#">Using the Naming Service to Locate a Log</a>	page 81
<a href="#">Using with JacORB</a>	page 83



# Using the Naming Service to Locate a Log

Distributed CORBA objects can be readily located using the Naming Service. The Naming Service associates a name with an object. This name can then be used to efficiently locate the object. The association between an object and a name is called a *name binding*.

The Naming Service stores name bindings in a hierarchical arrangement called a *naming context*. As well as organising name bindings, contexts are used to manage the bindings, including their creation and destruction. The top-most context is called the *root context* and may contain *child* contexts, as well as name bindings. The entire arrangement of root context, child contexts, and name bindings is called a *naming graph*.

## Creating a Name Binding

In order to create a name binding, a reference to the root context (or child context) where it will be stored must be obtained. The name itself is stored in a *NameComponent* object, which is then added to the context. The following code, taken from “[Locating the Log](#)”, shows the basic approach to binding a log object to a name.

**i** The ***OpenFusion Naming Service Guide*** provides complete details on using contexts and name bindings.

The object to be bound and an `int` holder are declared. The initialisation of the `orb` object is not shown. The name binding is then created as described in the steps given below

```
.
org.omg.CORBA.ORB orb = null;
BasicLog log = null;
org.omg.CORBA.IntHolder logId = null;

// bind the log object to a name with Naming Service
org.omg.CORBA.Object namingService = null;
NamingContextExt rootContext = null;
NameComponent basicLog[];

namingService = orb.resolve_initial_references ("NameService");
rootContext = NamingContextExtHelper.narrow(namingService);
basicLog = new NameComponent[1];

basicLog[0] = new NameComponent("BasicLog", "log");
rootContext.rebind(basicLog, log);

System.out.println("BasicLog created.");
```

- 1 Declare a CORBA object which will be resolved to the Naming Service, a root context object (of type *NamingContextExt*), and an array of type *NameComponent*. The *NameComponent* array is used to contain the name which will be used to identify the log object.
- 2 Resolve the Naming Service object using the standard CORBA `resolve_initial_references()` method, then narrow the Naming Service object to a root context using *NamingContextExtHelper.narrow()*. Initialise the *NameComponent* array, *basicLog*, to hold the name identifying the log.

- 3 Assign a new name component to `basicLog`, using two `String` parameters defining its `id` and `type`. The `id` and `type` can be any desired string, or even contain `NULL` values, provided `NULL` is not used for both `id` and `type`. Also, the `id-type` combination must be unique within the context. Finally, bind the log object to the name component using root context's `bind()` or `rebind()` methods. The `bind` method will throw an exception if the name component is already bound to an object; `rebind` will replace an existing binding or create a new binding if one does not exist. An object can be bound to more than one name component, but one name component can not be bound to more than one object.
- 4 A name binding can be removed using the root context's `unbind()` method. If a context is empty (that is, it does not contain any name bindings or child contexts), it can be destroyed using the context's `destroy()` method.

*i*

Removing a name binding does **not** destroy the object that the name is bound to.

The log object can now be located, from clients located anywhere in the system, using the context's `resolve()` method described below.

## Obtaining the Object with Resolve

Clients can now obtain a reference to the log by using the `resolve()` method. This is referred to as *resolving* the object.

An object can be resolved by simply passing a `NameComponent` object (containing the log's `id` and `kind` values) to the context's `resolve()` method, then narrowing the object to the `BasicLog` type (`resolve()` returns an `Object` type).

```
// obtain reference to the orb
orb = ORBAdapter.init(args);

// obtain reference to log using Naming Service
obj = orb.resolve_initial_references("NameService");
rootContext = NamingContextExtHelper.narrow(obj);
basicLog = new NameComponent[1];
basicLog[0] = new NameComponent("BasicLog", "log");
obj = rootContext.resolve(basicLog);
log = BasicLogHelper.narrow(obj);
```

# Using with JacORB

If the OpenFusion Log Service's NotifyLog is used with the NotifyLog with *JacORB*, then JacORB's *compactTypecodes* property must be turned off by setting it to 0 (zero). The *compactTypecodes* property is located in JacORB's *jacorb.properties* file. The default location of the *jacorb.properties* file is the *install/classes* directory.

The property should be changed from:

```
jacorb.compactTypecodes=2
```

to:

```
jacorb.compactTypecodes=0
```



# Index

## A

- AcknowledgeLevel (QoS property) 74
- AcknowledgeMode (QoS property) 74
- add\_constraints 43, 44
- add\_filter 44, 53
- Administration Manager 4
- Administrative Properties 74
- Administrative State 10, 11, 18, 72
  - locked 11, 18
  - Set 18
  - unlocked 11, 18
- Administrative State (log setting) 72
- Alarms 73
- Architecture 5, 7
- arg\_list 10
- attr\_list 9, 23
- Attribute Change Events 13
- AttributeValueChange 10, 11, 13, 75
- Availability Status 12, 22

## B

- BasicLog 5, 6, 16, 17, 20
  - Creating 17
  - Creating a Client 20
- BasicLog (including Interval) 20
- BasicLog (including Write and Query Related) 24
- BasicLogFactory 8, 18
- BasicLogFactorySingleton 63, 71
  - Configuration 63
- Benefits 3
- bind 17, 82
- Browser 76

## C

- Callback Operations 29, 34, 36
  - Push 58
  - push\_structured\_event 54
- Capacity 10
- Channel 9, 74
- child context 81
- Cleanup 39, 56
- Client
  - Connecting to an Event Channel 32
  - Event Consumer 36
  - Event Supplier 31
  - Terminating 33
- Clients 16
- ClientType 47
- Concepts 5
- Configuration 17, 26, 40, 49
- ConnectionReliability (QoS property) 74
- Constraint Expression 23, 43, 44
- Context 19
- copy 5, 10
- Copy Logs 76
- copy\_with\_id 10
- CORBA Properties 63

- CosEvent 29
- Create 17, 25, 27
- create\_filter 43
- Current setting 73
- Current setting (Log Scheduling) 73

## D

- Data, Sending 21
- Database Plugin Class (property) 66
  - Basic Log Factory 66
- DB.Plugin (property) 66
- DB.WAL (property) 65
- DB.WAL.Dir (property) 65
- DB.WAL.MaxSize (property) 66
- DeadLetterDrop (QoS property) 74
- Default
  - Values 17, 37
- Default Grammar 43
- Delete 78
- Dependencies (on Other Services) 5
- Dependencies on Other Services 5
- Destroy 5, 10
- destroy 27, 33, 43, 45, 82
- Destroy Logs 76
- DiscardPolicy (QoS property) 74
- disconnect\_push\_supplier 29, 31
- DisconnectCallback (QoS property) 74
- documentation
  - .pdf format ix
  - updates on the web ix
- Duration 12, 20

## E

- Enable Event Queueing (property) 66
- Enable Write Ahead Log (property) 65
- Event
  - Filtering 53
- Event Channel 5, 9, 29, 31, 32, 41
- Event Header 48
- Event Supplier 29
- EventIDSupported (QoS property) 74
- EventLog 6, 25, 30, 39
  - Creating 25
  - Creating a Consumer Client 34
  - Creating a Supplier Client 46
  - Duration 28
  - Time Values 35
- EventLogFactory 8, 26
- EventReliability (QoS property) 74
- Events 75
  - Creating and Sending 49
  - Filtering 42
  - Forwarding 13, 37
  - Receiving 37, 54
  - Sending 33
  - Working with 25
- EventType 44, 53
- Exceptions 19, 24, 28, 34, 39, 45, 51, 56, 58, 59
- EXTENDED\_TCL

SeeExtended Trader Constraint  
Language

## F

Factories  
  Log 5, 7  
Factory Id 11  
Features 14  
  Comparison 14  
File Queueing directory (property) 66  
Filter 42  
  add constraints 43, 44  
  Add Event Filter 44, 53  
  Constraints 74  
  Create 43  
  Creation 43  
  Location 43  
  Removal and Destruction 45  
  Set Record Filter 44, 53  
filter  
  get 44  
  remove 45  
Filter Factory 53  
filterable\_body 48  
filterable\_data 55  
Filtering 6, 11, 13, 42, 53  
Fixed Header 44, 48  
Forwarding 5, 11  
Forwarding State 13, 37, 73  
Full Action 72  
Full Action (property) 72

## G

General Properties 66  
Generated Events 5, 6, 10, 56, 58  
  Monitoring 56  
  Receiving 58  
generated events  
  AttributeValueChange 13  
  ObjectCreation 10  
  ObjectDeletion 10  
  ProcessingErrorAlarm 10  
  StateChange 10  
  ThresholdAlarm 10, 12  
get\_availability\_status 41  
get\_filter 44  
Grammar 23, 43  
Gregorian Calander 18, 23, 27, 37

## H

halt 11, 26, 41  
Header 44, 48

## I

id 9, 11, 23, 24  
info 9, 23  
Inheritance  
  Hierarchy 6, 8  
Initial Number of Logs (property) 67  
InitialLogs (property) 67  
Instrumentation 4

interface\_id 10  
IOR File Name 69  
IOR File Name (property) 69  
IOR Name Service 69  
IOR Name Service (property) 65, 69  
IOR Name Service Entry 68  
IOR Name Service Entry (property) 63, 68  
IOR URL 68  
IOR URL (property) 68  
IOR.File (property) 64, 69  
IOR.Server (property) 65  
IOR.URL (property) 64, 68

## L

LazyAcknowledgeInterval (QoS  
  property) 74  
Lifecycle Operations 10  
Lifetime 12, 72  
Locating the Log 19, 28  
locked 11, 18  
Log 5, 9, 72  
  BasicLog 5, 6  
  Configuration 17, 26, 40, 49  
  Create 17, 25, 27  
  EventLog 6, 25  
  Factories 5, 7  
  Generated Events 5, 6, 10, 56, 58  
  Locating 19, 28  
  Management 11  
  Maximum Log Size 11, 17  
  NotifyLog 5, 6  
  TypedEventLog 6  
  TypedNotifyLog 6  
  Types 5, 6, 7, 8  
    Features 14  
Log Capacity Alarm Threshold 12  
Log Clients 16  
Log Creation 16, 75  
Log Creation and Configuration 17, 26, 40  
Log Duration 12, 20  
Log Factories 7  
Log Factory 72  
Log Full Actions 11, 27  
  halt 11, 26, 41  
  wrap 11  
Log Generated Events 10  
  Monitoring 56  
Log Management 11  
Log Networks 9  
Log Object 5  
Log Object Settings 72  
Log Record 9, 12, 17, 23, 24, 41  
  filtering 13  
  Lifetime 12  
  retrieve 23, 24, 30, 36, 37, 55  
Log Records  
  Retrieving 37, 55  
Log Scheduling 12  
Log Service  
  Configuration 63  
  Exceptions 59



- Manager 71
  - Supplemental Info 59
  - Using Specific Features 15
- Log Service Manager
  - Using 71
- Log Types 5
- LogFactorySingleton Configuration 63
- LogMgr 7
- LogRecord 9

## M

- Main Filter Components 43
- Management 11
  - Additional Operations 13
- Management Operations
  - Common 11
- Match 78
- Max records returned from query (property) 67
- MaxConsumers (admin property) 74
- MaxEventsPerConsumer (QoS property) 74
- Maximum Log Size 11, 17
- Maximum number of queue file (property) 67
- Maximum queue file size (property) 67
- Maximum Size 72
- Maximum Size (log setting) 72
- MaximumBatchSize (QoS property) 74
- MaxInactivityInterval (QoS property) 74
- MaxQueueLength (admin property) 74
- MaxReconnectEvents (QoS property) 74
- MaxRecordsReturnedFromQuery (property) 67
- MaxSuppliers (admin property) 74
- MultipleReceiversSupported (QoS property) 74

## N

- Name Binding 81
  - Creating 81
- NameComponent 21, 28
- Naming Graph 81
- Naming Service 20, 56, 81
  - bind 17, 82
  - Child Context 81
  - Context 19
  - Name Binding 81
  - Naming Graph 81
  - Resolve 81, 82
  - resolve 19
  - Root Context 81
- Notification Service
  - Introduction 3
  - Service Dependencies 5
- NotifyLog 5, 6, 16, 40, 45, 46, 51, 56, 58
  - Creating 39
  - Creating a Consumer Client 51
  - Creating a Supplier Client 29
  - Structured Events 39
  - Time Values 52, 57

- NotifyLogFactory 8, 26
- NotifyLogFactorySingleton 63, 71

## O

- Object.Name (property) 63, 68
- ObjectAdapter 29, 48
- ObjectCreation 10, 75
- ObjectDeletion 10, 75
- obtain\_push\_supplier 36
- OMG
  - Standard Features 4
- OpenFusion
  - Enhancements 4
  - Extensions 16
  - Features 4
  - Orb Adaptor 17, 20, 25, 30, 51, 57
- operation\_name 10
- Operations 77
- ORB 17
- Orb Abstraction Layer 17
- orb abstraction layer 29, 48
- Orb Adaptor 17, 20, 25, 30, 51, 57
- ORBAdapter
  - Also see ObjectAdapter
- Order query results (property) 68
- OrderPolicy (QoS property) 74
- OrderQueryResults (property) 68

## P

- PacingInterval (QoS property) 74
- PDF documentation ix
- Persistence Properties 65
- Priority (QoS property) 74
- ProcessingErrorAlarm 10
- ProcessSingleton Configuration 68
- Properties
  - Common 63
  - General 66
- Proxies 32
- Proxy
  - See Proxies
- Proxy Filter, Setting a 53
- Pull Model 31, 32, 36
- Push 31, 33, 35, 36, 58
- Push Model 31, 32, 36
- push\_structured\_event 54
- PushSupplierOperations 29, 34

## Q

- QoS
  - See Quality of Service
- QoS Settings 74
  - Log Factories 74
- QoSFlush 14, 26
- QoSNone 14, 26
- QoSReliability 14, 26
- Quality of Service 13, 14, 47
  - QoSFlush 14, 26
  - QoSNone 14, 26
  - QoSReliability 14, 26
- Query 4, 20, 23, 24, 77

Querying the Log Records 23

## R

Read and Write Settings 11  
Read-Only Settings 11  
Receiving Events 37, 54  
Receiving Generated Events 58  
ReconnectInterval (QoS property) 74  
Record  
    See Log Record  
record  
    See Log Record  
Record Lifetime 72  
Record Lifetime (log setting) 72  
Registering with the Naming Service 45  
RejectNewEvents (admin property) 74  
remainder\_of\_body 48, 49, 54  
remove\_filter 45  
Repository Id 10  
Resolve 81, 82  
resolve 19  
ResolveName (property) 64  
Retrieve 23, 24, 30, 36, 37, 55, 77  
Retrieving Log Records 37, 55  
root context 81

## S

Scheduling 12, 73  
Sending Data 21  
Sending Events 33  
Set Attributes 78  
set\_administrative\_state 18  
set\_filter 44, 53  
set\_interval 17, 18, 23  
set\_week\_mask 23, 27  
Setting, Current 73  
Settings 77  
    Log Object 72  
Singletons  
    BasicLogFactorySingleton 63  
SNMP 4  
Start and Stop 73  
Start and Stop (log scheduling) 73  
Starting Log Service Manager 71  
StartTimeSupported (QoS property) 74  
StateChange 10, 75  
StopTimeSupported (QoS property) 74  
Structured Event 6, 43, 49  
    header 44, 48  
Structured Events 39, 48  
Structured Push Supplier 47  
StructuredProxyPushSupplier 52

## T

ThreadIdleTime (QoS property) 74  
ThreadPoolSize (QoS property) 74  
ThresholdAlarm 10, 12, 75  
Time 9, 18, 23  
TimeInterval 18, 27  
Timeout (QoS property) 74  
Timer 20, 30, 46

TimerTask 20, 30, 46  
TypedEventLog 6  
TypedEventLogFactory 8  
TypedLogRecord 9  
    arg\_list 10  
    interface\_id 10  
    operation\_name 10  
TypedNotifyLog 6  
TypedNotifyLogFactory 8  
Types  
    Log 5, 6, 7, 8

## U

unlocked 11, 18  
Using the Log Service Manager 71

## V

Values  
    Default 17, 37  
Variable Header 48

## W

Week Mask 27, 73  
Week Mask (Log Scheduling) 73  
WeekMask 27  
wrap 11  
Write 21, 22, 23  
Write Ahead Log 65  
Write Ahead Log Directory (property) 65  
Write Ahead Log Maximum Size  
    (property) 66  
Write Data 21, 22, 23  
write\_records 23