

# XORP Libxorp Library Overview

## Version 1.2

XORP Project  
International Computer Science Institute  
Berkeley, CA 94704, USA  
<http://www.xorp.org/>  
[feedback@xorp.org](mailto:feedback@xorp.org)

March 8, 2006

## 1 Introduction

The *libxorp* library contains a set of classes for basic XORP functionality such as IP addresses and subnets, timers, event loops, etc. It is used by virtually every other XORP component, and its main purpose is to simplify the implementation of those components.

### 1.1 Overview

Currently, the libxorp library contains the following classes and components (in alphabetical order):

- *asnum.hh*: *class AsNum*: A class for storing an AS number used by protocols such as BGP.
- *asyncio.hh*: *class AsyncFileReader*, *class AsyncFileWriter*: Asynchronous file transfer classes.
- *buffer.hh*: *class Buffer*: A class for storing buffered data.
- *buffered\_asyncio.hh*: *class BufferedAsyncReader*: Buffered asynchronous file reader class.
- *c\_format.hh*: *c\_format()*: A macro that creates a C++ string from a C-style printf(3)-formatted string.
- *callback.hh*, *callback\_debug.hh*, *callback\_nodebug.hh*, *safe\_callback\_obj.hh*: Callback mechanism.
- *clock.hh*: *class SystemClock*: A class for providing the interface to obtain the system clock.
- *config\_param.hh*: *template class ConfigParam*: A class for storing a configuration parameter.
- *debug.h*: Provides facility for generating debug messages.
- *ether\_compat.h*: Ethernet manipulation compatibility functions.
- *eventloop.hh*: *class EventLoop*: Event loop class for coordinated operations between timers and I/O operations on file descriptors.

- *exceptions.hh*: Standard XORP C++ exceptions.
- *heap.hh*: *class Heap*: Provides Heap data structure.
- *ipnet.hh*, *ipv4net.hh*, *ipv6net.hh*, *ipvxnet.hh*: *class IPv4Net*, *class IPv6Net*, *class IPvXNet*: Implementation of classes for basic subnet addresses (for IPv4, IPv6 and dual IPv4/6 address family respectively).
- *ipv4.hh*, *ipv6.hh*, *ipvx.hh*: *class IPv4*, *class IPv6*, *class IPvX*: Implementation of classes for basic IP addresses (for IPv4, IPv6 and dual IPv4/6 address family respectively).
- *mac.hh*: *class Mac*, *class EtherMac*: Containers for MAC types.
- *nexthop.hh*: Classes that contain routing next-hop information.
- *popen.hh*: The interface for the local implementation of *popen(2)* and *pclose(2)*.
- *profile.hh*: Implementation of a mechanism for event profiling.
- *range.hh*: *class U32Range*, *class IPv4Range*, *class IPv4Range*: A set of classes that implement linear ranges (e.g., integers or addresses).
- *ref\_ptr.hh*: *template class ref\_ptr*: Reference counted pointer class.
- *ref\_trie.hh*: Implementation of a trie to support route lookups. Based on *trie.hh*, but with reference-counted storage supporting delayed deletion.
- *run\_command.hh*: *class RunCommand*: A class for running an external command.
- *safe\_callback\_obj.hh*, *class CallbackSafeObject*: Implementation of a base class for objects that are callback safe.
- *selector.hh*: I/O multiplexing interface.
- *service.hh*: Provides base for asynchronous service classes.
- *status\_codes.h*: Process states status codes used by processes when reporting their operational status to the router manager.
- *time\_slice.hh*: *class TimeSlice*: A class for computing whether some processing is taking too long.
- *timer.hh*, *class XorpTimer*: XORP timer facility.
- *timespent.hh*: *class TimeSpent*: A class used for debugging purpose to find code that has taken too long to execute.
- *timeval.hh*: *class TimeVal*: A class for storing time values (similar to *struct timeval*).
- *tlv.hh*, *class Tlv*: Facility for reading and writing TLV (Type-Lenght-Value) records.
- *token.hh*: Token related definitions.
- *transactions.hh*: Facility for transaction operations.

- *trie.hh*: Implementation of a trie to support route lookups.
- *utility.h*: Contains various mini-utilities (mostly compiler-related helpers).
- *utils.hh*: Contains various utilities (e.g., to delete a list or array of pointers and the objects pointed to).
- *vif.hh*: *class Vif*, *class VifAddr*: Virtual interface and virtual interface address classes.
- *xlog.h*: Provides facility for log messages generation.
- *xorp.h*: The XORP main include file that should be included by all XORP C and C++ files.
- *xorpdf.hh*: *class XorpFd*: Implementation of a wrapper class used to encapsulate a file descriptor.

Each of the components is described in Section 2.

## 2 Components Description

This section contains a brief description of each of the components of the *libxorp* library. This description is for informative purpose only. The source code for each component is the ultimate source for programming reference, and implementation details.

### 2.1 *asnum.hh*

This file contains *class AsNum*: a class for storing an AS number used by protocols such as BGP. This class can be used to store an AS number that can be either 16 or 32 bits. Originally, the AS numbers were defined as 16-bit unsigned numbers. Later the “extended” AS numbers were introduced, which are unsigned 32-bit numbers.

### 2.2 *asyncio.hh*

This file contains asynchronous file transfer classes. These utilize XORP EventLoop and its SelectorList to read or write files asynchronously. The user creates an *AsyncFile{Reader,Writer}* and adds a buffer for reading or writing with *add\_buffer()*. A callback provided with each buffer is called every time I/O happens on the buffer. Reading or writing only begins when *start()* is called, and normally continues until there are no buffers left.

From the developer’s point of view, the following classes are of interest: *class AsyncFileReader*, *class AsyncFileWriter*.

### 2.3 *buffer.hh*

This file contains *class Buffer*: a class for conveniently storing and accessing buffered data. Currently it has limited applicability.

### 2.4 *buffered\_asyncio.hh*

This file contains *class BufferedAsyncReader*: a class for buffered asynchronous file reading. Unlike the *class AsyncFileReader* (see Section 2.2) which delivers data whenever it is available, the *class BufferedAsyncReader* buffers the data and delivers it whenever its size has reached a predefined threshold.

## 2.5 `c_format.hh`

This file contains `c_format()`: a macro that creates a C++ string from a C-style `printf(3)`-formatted string. It takes the same arguments as `printf(3)`, but `%n` is illegal and will cause `abort` to be called.

In practice, `c_format()` is a nasty macro, but by doing this we can check the compile time arguments are sane and the run time arguments.

## 2.6 `callback.hh`, `callback_debug.hh`, `callback_nodebug.hh`, `safe_callback_obj.hh`

These files contain an implementation of a callback mechanism. XORP is an asynchronous programming environment and as a result there are many places where callbacks are useful. Callbacks are typically invoked to signify the completion or advancement of an asynchronous operation.

XORP provides a generic and flexible callback interface that utilizes overloaded templated functions for generating callbacks in conjunction with many small templated classes. Whilst this makes the syntax a little unpleasant, it provides a great deal of flexibility.

`XorpCallback` objects are objects created by the `callback()` function which returns a reference pointer to a newly created callback object. The callback is invoked by calling the `dispatch()` method on that object.

There are two versions of the callback mechanism: debug and non-debug version. The debug version includes additional information with each callback (*e.g.*, file name and line number where the callback was invoked), records callback tracing events, etc, but creates additional overhead to the system. Non-debug callbacks are used by default; the debug callbacks can be enabled by defining `DEBUG_CALLBACK` before including `callback.hh`, or by running `./configure --enable-callback-debug` before compiling XORP.

For more details on the callback mechanism, and for usage examples, see the beginning of `callback_debug.hh` or `callback_nodebug.hh`. Note that these files are auto-generated by `callback-gen.py` (a Python script), therefore they should never be edited.

## 2.7 `clock.hh`

This file contains the implementation of `class SystemClock` that provides the interface for obtaining the system clock.

## 2.8 `config_param.hh`

This file contains the implementation of `template class ConfigParam`: a class for storing a configuration parameter.

This class can be used to store the value of a configuration parameter. Such parameter has a current and a default value. The `ConfigParam` class has the facility to add a callback that is invoked whenever the value of the configuration parameter is changed.

## 2.9 `debug.h`

This file provides facility for debug messages generation. More specifically, it defines the `debug_msg()`, the macro responsible for generating debug messages. It takes the same arguments as `printf(3)`. For example:

```
debug_msg("The number is %d\n", 5);
```

For more details see the comments inside that file.

## 2.10 ether\_compat.h

This file contains Ethernet-related manipulation compatibility functions. For example, it includes the appropriate system files, and declares functions `ether_aton()` and `ether_ntoa()` (implemented locally in `ether_compat.c`) if the system is missing the corresponding `ether_aton(3)` and `ether_ntoa(3)`.

## 2.11 eventloop.hh

This file defines *class EventLoop*. It is used to co-ordinate interactions between a `TimerList` and a `SelectorList` for XORP processes. All `XorpTimer` and `select` operations should be co-ordinated through this interface.

## 2.12 exceptions.hh

This file contains *class XorpException*: a base class for XORP C++ exceptions. It contains also all standard XORP C++ exceptions. An example of such exception is *class InvalidFamily* which is thrown if the address family is invalid (for example, by an IPvX constructor when invoked with an invalid address family).

## 2.13 heap.hh

This file contains *class Heap*. The `Heap` class is used by the `TimerList` class as it's priority queue for timers. This implementation supports removal of arbitrary objects from the heap, even if they are not located at the top.

## 2.14 ipnet.hh, ipv4net.hh, ipv6net.hh, ipvxnet.hh

These files contain the declaration of the following classes: *class IPv4Net*, *class IPv6Net*, *class IPvXNet*, which are classes for basic subnet addresses (for IPv4, IPv6 and dual IPv4/6 address family respectively). `IPvXNet` can be used to store a subnet address that has either IPv4 or IPv6 address family.

Most of the implementation is contained in file *ipnet.hh*, which contains a *template class IPNet*. The `IPv4Net`, `IPv6Net`, and `IPvXNet` classes are derived from that template.

## 2.15 ipv4.hh, ipv6.hh, ipvx.hh

These files contain the declaration for the following classes: *class IPv4*, *class IPv6*, *class IPvX*, which are classes for basic IP addresses (for IPv4, IPv6 and dual IPv4/6 address family respectively). `IPvX` can be used to store an address that has either IPv4 or IPv6 address family.

## 2.16 mac.hh

This file declares the following classes: *class Mac*, *class EtherMac*. The first class is a generic container for any type of MAC. The second class is a container for Ethernet MAC address.

## 2.17 nexthop.hh

This file declares a number of classes that can be used to contain routing next-hop information. For example, *class NextHop* is the generic class for holding information about routing next hops. `NextHops` can be of many types, including immediate neighbors, remote routers (with IBGP), discard interfaces, encapsulation

endpoints, etc. NextHop itself doesn't really do anything useful, except to provide a generic handle for the specialized subclasses. The specialized subclasses are:

- IPPeerNextHop is for next hops that are local peers.
- IPEncapsNextHop is for “next hops” that are non-local, and require encapsulation to reach. An example is the PIM Register Encapsulation.
- IPEXternalNextHop An IP nexthop that is not an intermediate neighbor.
- DiscardNextHop is a discard interface.

## 2.18 popen.hh

This file contains the interface for the local implementation of *popen(2)* and *pclose(2)*. The corresponding local names are *popen2()* and *pclose2()* respectively. Unlike the system's *popen(2)*, the local *popen2()* implementation allows the user to specify the streams where the *stdout* and *stderr* of the command will be redirected to.

## 2.19 profile.hh

This file implements the mechanism for event profiling. The developer can add profiling entries at various places of the program. Each profiling entry has a name and it can be enabled or disabled. In addition, a number of strings can be added to each profiling entry, and those strings can be read at some later stage.

## 2.20 range.hh

This file implements the following classes: *class U32Range*, *class IPv4Range*, *class IPv4Range*.

Those classes implement linear ranges X..Y (*e.g.*, for integers or addresses). A linear range is defined by its low and high inclusive boundaries. It is the user's responsibility to ensure that the condition ( $\text{low} \leq \text{high}$ ) always holds.

## 2.21 ref\_ptr.hh

This file declares *template class ref\_ptr*: reference counted pointer class.

The *ref\_ptr* class is a strong reference class. It maintains a count of how many references to an object exist and releases the memory associated with the object when the reference count reaches zero. The reference pointer can be dereferenced like an ordinary pointer to call methods on the reference counted object.

At the time of writing the only supported memory management is through the *new* and *delete* operators. At a future date, this class should support the STL allocator classes or an equivalent to provide greater flexibility.

## 2.22 ref\_trie.hh

This file implements a trie to support route lookups. The implementation is template-based, and is based on the code in *trie.hh*. From developer's point of view, templates *RefTrie*, *RefTrieNode*, *RefTriePreOrderIterator*, and *RefTriePostOrderIterator* are the most important. Those templates should be invoked with two classes, the basetype “A” for the search Key (which is a subnet,  $\text{IPNet} \langle A \rangle$ ), and the Payload.

RefTrie differs from Trie (and its associated classes) in that the RefTrieNode includes a reference count of how many RefTrieIterators are pointing at it. If a RefTrieNode is deleted, but has a non-zero reference count, deletion will be delayed until the reference count becomes zero. In this way, additions and deletions to the RefTrie cannot cause a RefTriePreOrderIterator or RefTriePostOrderIterator to reference invalid memory, although a deletion and subsequent addition can cause the payload data referenced by an iterator to change.

## 2.23 run\_command.hh

This file implements *class RunCommand* which provides the mechanism for running an external command. In addition to the command name and its arguments, the developer can specify three callbacks:

- *stdout\_cb*: the callback to call when there is data on the standard output.
- *stderr\_cb*: the callback to call when there is data on the standard error.
- *done\_cb*: the callback to call when the command is completed.

## 2.24 safe\_callback\_obj.hh

This file declares *class CallbackSafeObject*. Objects that wish to be callback safe should be derived from this class. When a CallbackSafeObject is destructed it informs all the callbacks that refer to it that this is the case and invalidates (sets to null) the object they point to.

## 2.25 selector.hh

This file contains the I/O multiplexing interface. The particular class of interest is *class SelectorList*.

A SelectorList provides an entity where callbacks for pending I/O operations on file descriptors may be registered. The callbacks are invoked when one of the select methods is called and I/O is pending on the particular descriptors.

## 2.26 service.hh

This declares *class ServiceBase*. A service is a class that can be started and stopped and would typically involve some asynchronous processing to transition between states. The base class provides a state model and methods for transitioning between states. Mandatory transition methods, like start and stop, are abstract in the base class.

## 2.27 status\_codes.h

This file contains the enumerated *ProcessStatus* codes that a XORP process should report to the XORP router manager (*rtrmgr*) [1]. The file itself contains a detailed explanation of the process states (valid transaction between states, triggering events, actions, etc).

## 2.28 time\_slice.hh

This file declares *class TimeSlice*. This class can be used to compute whether some processing is taking too long time to complete. It is up to the program that uses TimeSlice to check whether the processing is taking too long, and suspend processing of that task if necessary.

## 2.29 timer.hh

This file declares the XORP timer facility. The only class of interest from a developer's point of view is *class XorpTimer*.

## 2.30 timespent.hh

This file declares and implements *class TimeSpent*. This class is used for debugging purposes to find code that has taken too long to execute.

## 2.31 timeval.hh

This file contains the implementation of *class TimeVal* for storing time values (similar to *struct timeval*). *TimeVal* implements the appropriate constructors and numerous helper methods (*e.g.*, Less-Than and Addition operators, etc).

## 2.32 tlv.hh

This file contains the implementation of *class Tlv* for reading and writing TLV (Type-Length-Value) records from/to a file. The records are stored in network byte order format.

## 2.33 token.hh

This file contains various token-related definitions. A token is a sequence of symbols separated from other tokens by some pre-defined symbols. In this implementation, the separators are the `is_space(3)` and `'—'` characters. The facilities in that file are to copy tokens, removing them from a token line, etc. Currently, this file is used only by the CLI, therefore in the future it may be moved to the CLI itself.

## 2.34 transactions.hh

This file contains the facility for transactions. A transaction consists of a sequence of transaction operations, each of which is a command. The *TransactionManager* class provides a front-end for creating, dispatching, and destroying transactions.

## 2.35 trie.hh

This file implements a trie to support route lookups. The implementation is template-based. From a developer's point of view, templates *Trie*, *TrieNode*, *TriePreOrderIterator*, and *TriePostOrderIterator* are the most important. Those templates should be invoked with two classes, the basetype "A" for the search Key (which is a subnet, `IPNet<A>`), and the Payload.

## 2.36 utility.h

This file contains various mini-utilities. Those utilities are mostly compiler-related helpers; *e.g.*, compile-time assertion, `UNUSED(var)` macro to suppress warnings about unused function arguments, etc.

## 2.37 `utils.hh`

This file contains various helper utilities. Currently, the only two utilities are template functions to delete a list or array of pointers and the objects pointed to.

## 2.38 `vif.hh`

This file declares the following classes: *class Vif*, *class VifAddr*.

Class *Vif* holds information about a virtual interface. A *Vif* may represent a physical interface, or may represent more abstract entities such as the Discard interface, or a VLAN on a physical interface. *VifAddr* holds information about an address of a virtual interface. A virtual interface may have more than one *VifAddr*.

## 2.39 `xlog.h`

This file provides facility for log messages generation, similar to `syslog`. The log messages may be output to multiple output streams simultaneously. Below is a description of how to use the log utility.

- The `xlog` utility assumes that `XORP_MODULE_NAME` is defined (per module). To do so, you must have in your directory a file like “`foo.module.h`”, and inside it should contain something like:

```
#define XORP_MODULE_NAME "BGP"
```

This file then has to be included by each `*.c` and `*.cc` file, and **MUST** be the first of the included local files.

- Before using the `xlog` utility, a program **MUST** initialize it first (think of this as the `xlog` constructor):

```
int xlog_init(const char *process_name, const char *preamble_message);
```

Further, if a program tries to use `xlog` without initializing it first, the program will exit.

- To add output streams, you **MUST** use one of the following (or both):

```
int xlog_add_output(FILE* fp);
int xlog_add_default_output(void);
```

- To change the verbosity of all `xlog` messages, use:

```
xlog_set_verbose(xlog_verbose_t verbose_level);
```

where “`verbose_level`” is one of the following (`XLOG_VERBOSE_MAX` excluded):

```
typedef enum {
    XLOG_VERBOSE_LOW = 0,          /* 0 */
    XLOG_VERBOSE_MEDIUM,         /* 1 */
    XLOG_VERBOSE_HIGH,           /* 2 */
    XLOG_VERBOSE_MAX
} xlog_verbose_t;
```

Default value is `XLOG_VERBOSE_LOW` (least details). Larger value for “verbose\_level” adds more details to the preamble message (e.g., file name, line number, etc, about the place where the log message was initiated).

Note that the verbosity level of message type `XLOG_LEVEL_FATAL` (see below) cannot be changed and is always set to the most verbose level (`XLOG_VERBOSE_HIGH`).

- To change the verbosity of a particular message type, use:

```
void xlog_level_set_verbose(xlog_level_t log_level,
xlog_verbose_t verbose_level);
```

where “log\_level” is one of the following (`XLOG_LEVEL_MAX` excluded):

```
typedef enum {
    XLOG_LEVEL_FATAL = 0,          /* 0 */
    XLOG_LEVEL_ERROR,           /* 1 */
    XLOG_LEVEL_WARNING,         /* 2 */
    XLOG_LEVEL_INFO,            /* 3 */
    XLOG_LEVEL_TRACE,           /* 4 */
    XLOG_LEVEL_MAX
} xlog_level_t;
```

Note that the verbosity level of message type `XLOG_LEVEL_FATAL` cannot be changed and is always set to the most verbose level (`XLOG_VERBOSE_HIGH`).

- To start the xlog utility, you **MUST** use:

```
int xlog_start(void);
```

- To enable or disable a particular message type, use:

```
int xlog_enable(xlog_level_t log_level);
int xlog_disable(xlog_level_t log_level);
```

By default, all levels are enabled. Note that `XLOG_LEVEL_FATAL` cannot be disabled.

- To stop the logging, use:

```
int xlog_stop(void);
```

Later you can restart it again by `xlog_start()`

- To gracefully exit the xlog utility, use

```
int xlog_exit(void);
```

(think of this as the xlog destructor).

Below is an example of using the XLOG facility:

```
int
main(int argc, char *argv[])
{
    //
    // Initialize and start xlog
    //
    xlog_init(argv[0], NULL);
    xlog_set_verbose(XLOG_VERBOSE_LOW); // Least verbose messages
    // Increase verbosity of the error messages
    xlog_level_set_verbose(XLOG_LEVEL_ERROR, XLOG_VERBOSE_HIGH);
    xlog_add_default_output();
    xlog_start();

    // Do something

    //
    // Gracefully stop and exit xlog
    //
    xlog_stop();
    xlog_exit();

    exit (0);
}
```

Typically, a developer would use the macros described below to print a message, add an assert statement, place a marker, etc. If a macro accepts a message to print, the format of the message is same as printf(3). The only difference is that the xlog utility automatically adds '\n', (i.e. end-of-line) at the end of each string specified by format:

- `XLOG_FATAL(const char *format, ...)`  
Write a FATAL message to the xlog output streams and abort the program.
- `XLOG_ERROR(const char *format, ...)`  
Write an ERROR message to the xlog output streams.
- `XLOG_WARNING(const char *format, ...)`  
Write a WARNING message to the xlog output streams.
- `XLOG_INFO(const char *format, ...)`  
Write an INFO message to the xlog output streams.
- `XLOG_TRACE(int cond_boolean, const char *format, ...)`  
Write a TRACE message to the xlog output stream, but only if `cond_boolean` is not 0.

- `XLOG_ASSERT(assertion)`  
The XORP replacement for `assert(3)`, except that it cannot be conditionally disabled and logs error messages through the standard `xlog` mechanism. It calls `XLOG_FATAL( )` if the assertion fails.
- `XLOG_UNREACHABLE( )`  
A marker that can be used to indicate code that should never be executed.
- `XLOG_UNFINISHED( )`  
A marker that can be used to indicate code that is not yet implemented and hence should not be run.

## 2.40 `xorp.h`

This is the XORP main include file that should be included by all XORP C and C++ files. This file itself includes a number of frequently used system header files, defines several commonly used values, etc.

## 2.41 `xorpdf.hh`

This file contains the implementation of *class XorpFd* used to encapsulate a file descriptor.

It exists because of fundamental differences between UNIX and Windows in terms of how the two families of operating systems deal with file descriptors; in most flavours of UNIX, all file descriptors are created equal, and may be represented using an 'int' type which is usually 32 bits wide. In Windows, sockets are of type `SOCKET`, which is a typedef alias of `u_int`; whereas all other system objects are of type `HANDLE`, which in turn is a typedef alias of 'void \*'.

## A Modification History

- December 11, 2002: Version 0.1 completed.
- March 10, 2003: Updated to match XORP version 0.2 release code; add information about `RefTrie`; cleanup.
- June 9, 2003: Updated to match XORP version 0.3 release code.
- August 28, 2003: Updated to match XORP version 0.4 release code.
- November 6, 2003: Updated to match XORP version 0.5 release code.
- July 8, 2004: Updated to match XORP version 1.0 release code.
- April 13, 2005: Updated to match XORP version 1.1 release code: added information for *buffered\_asyncio.hh*, *clock.hh*, *popen.hh*, *profile.hh* and *run\_command.hh*.
- March 8, 2006: Updated to match XORP version 1.2 release code: added information for *range.hh*, *tlv.hh* and *xorpdf.hh*.

## References

- [1] XORP Router Manager Process (rtrmgr). XORP technical document. <http://www.xorp.org/>.