



# DBL™ Application Programming Interface

Version 5.5.0.54026.PHX

February 11, 2020

All information contained in this document is proprietary to CSP, Inc. and may not be reproduced, distributed, or disseminated, in whole or in part, without the written permission of an authorized representative of CSP, Inc.

All specifications presented in this document are subject to change at any time, and without prior notice.

Myricom® and Myrinet® are registered trademarks of CSP, Inc. DBL™ is a trademark of CSP, Inc. Other trademarks appearing in this document are those of their respective owners.

©2008-2014, CSP, Inc.



# Contents

<b>1</b>	<b>DBL</b>	<b>1</b>
1.1	Introduction . . . . .	1
1.1.1	Terms and Concepts . . . . .	1
1.1.2	Example Pseudo-Code . . . . .	1
1.2	Interaction with Sockets . . . . .	3
1.3	Receive Data Buffering . . . . .	3
1.4	Direct Access to Ring Contents . . . . .	3
1.5	Batch Receive . . . . .	4
1.6	Arbitration duplicate packet elimination . . . . .	4
<b>2</b>	<b>Module Index</b>	<b>5</b>
2.1	API Reference . . . . .	5
<b>3</b>	<b>Namespace Index</b>	<b>7</b>
3.1	Namespace List . . . . .	7
<b>4</b>	<b>Data Structure Index</b>	<b>9</b>
4.1	Data Structures . . . . .	9
<b>5</b>	<b>Module Documentation</b>	<b>11</b>
5.1	API Reference . . . . .	11
5.1.1	Detailed Description . . . . .	13
5.1.2	API Reference . . . . .	13
5.1.3	Macro Definition Documentation . . . . .	13
5.1.3.1	DBL_VERSION_API . . . . .	13
5.1.4	Enumeration Type Documentation . . . . .	14
5.1.4.1	dbl_ext_recvmode . . . . .	14
5.1.4.2	dbl_filter_mode . . . . .	14
5.1.4.3	dbl_recvmode . . . . .	14
5.1.5	Function Documentation . . . . .	14
5.1.5.1	dbl_ab_get_stream . . . . .	14
5.1.5.2	dbl_ab_get_unit . . . . .	15
5.1.5.3	dbl_ab_get_unit_protocol_info . . . . .	15
5.1.5.4	dbl_ab_set_seq . . . . .	16
5.1.5.5	dbl_bind . . . . .	16
5.1.5.6	dbl_bind_addr . . . . .	17
5.1.5.7	dbl_close . . . . .	17
5.1.5.8	dbl_device_enable . . . . .	17
5.1.5.9	dbl_device_get_attrs . . . . .	18
5.1.5.10	dbl_device_handle . . . . .	18

5.1.5.11	dbl_device_set_attrs	18
5.1.5.12	dbl_ext_recvfrom	19
5.1.5.13	dbl_get_params	19
5.1.5.14	dbl_getaddress	19
5.1.5.15	dbl_getticks	20
5.1.5.16	dbl_gettime	20
5.1.5.17	dbl_init	21
5.1.5.18	dbl_mcast_block_source	21
5.1.5.19	dbl_mcast_join	21
5.1.5.20	dbl_mcast_join_source	22
5.1.5.21	dbl_mcast_leave	22
5.1.5.22	dbl_mcast_leave_source	23
5.1.5.23	dbl_mcast_unblock_source	23
5.1.5.24	dbl_open	23
5.1.5.25	dbl_open_if	24
5.1.5.26	dbl_recvfrom	24
5.1.5.27	dbl_send	25
5.1.5.28	dbl_send_connect	26
5.1.5.29	dbl_send_disconnect	26
5.1.5.30	dbl_sendto	27
5.1.5.31	dbl_set_filter_mode	27
5.1.5.32	dbl_shutdown	27
5.1.5.33	dbl_unbind	27
5.2	Flags used for dbl_open()	29
5.2.1	Detailed Description	29
5.2.2	Macro Definition Documentation	29
5.2.2.1	DBL_OPEN_DISABLED	29
5.2.2.2	DBL_OPEN_HW_TIMESTAMPING	29
5.2.2.3	DBL_OPEN_RAW_MODE	29
5.2.2.4	DBL_OPEN_THREADSAFE	29
5.3	Flags used for dbl_bind()	30
5.3.1	Detailed Description	30
5.3.2	Macro Definition Documentation	30
5.3.2.1	DBL_BIND_BROADCAST	30
5.3.2.2	DBL_BIND_DUP_TO_KERNEL	31
5.3.2.3	DBL_BIND_NO_UNICAST	31
5.3.2.4	DBL_BIND_REUSEADDR	31
5.3.2.5	DBL_BIND_TX_TIMESTAMP	31
5.3.3	Function Documentation	31
5.3.3.1	dbl_eventq_close	31
5.3.3.2	dbl_eventq_consume	31
5.3.3.3	dbl_eventq_count	32
5.3.3.4	dbl_eventq_inspect	32
5.3.3.5	dbl_eventq_open	32
5.3.3.6	dbl_eventq_peek_head	33
5.3.3.7	dbl_eventq_peek_next	33
5.3.3.8	dbl_raw_send	33
5.3.3.9	dbl_set_filter	34
5.4	Flags for dbl_send()	35
5.4.1	Detailed Description	35

5.4.2	Macro Definition Documentation	35
5.4.2.1	DBL_MCAST_LOOPBACK	35
5.4.2.2	DBL_NONBLOCK	35
5.4.2.3	MSG_WARM	35
5.4.2.4	MSG_WARM	35
5.5	Extensions	36
5.5.1	Detailed Description	37
5.5.2	Introduction to extensions	37
5.5.3	Function Documentation	37
5.5.3.1	dbl_ext_accept	37
5.5.3.2	dbl_ext_channel_type	37
5.5.3.3	dbl_ext_getchopt	38
5.5.3.4	dbl_ext_listen	38
5.5.3.5	dbl_ext_poll	38
5.5.3.6	dbl_ext_recv	39
5.5.3.7	dbl_ext_recvmsg	39
5.5.3.8	dbl_ext_send	40
5.5.3.9	dbl_ext_setchopt	40
5.6	Specific Options for dbl_ext_setchopt()	41
5.6.1	Detailed Description	41
5.6.2	Macro Definition Documentation	41
5.6.2.1	SO_TIMESTAMPING	41
<b>6</b>	<b>Namespace Documentation</b>	<b>43</b>
6.1	dbl Namespace Reference	43
6.1.1	Detailed Description	43
<b>7</b>	<b>Data Structure Documentation</b>	<b>45</b>
7.1	dbl__packet Struct Reference	45
7.2	dbl_device_attrs Struct Reference	45
7.2.1	Detailed Description	45
7.2.2	Field Documentation	46
7.2.2.1	hw_timestamping	46
7.2.2.2	recvq_filter_mode	46
7.2.2.3	recvq_size	46
7.3	dbl_ext_recv_info Struct Reference	46
7.3.1	Detailed Description	46
7.3.2	Field Documentation	46
7.3.2.1	buf	46
7.3.2.2	chan	46
7.3.2.3	chan_context	47
7.3.2.4	msg_len	47
7.3.2.5	sin_from	47
7.3.2.6	sin_to	47
7.3.2.7	timestamp	47
7.4	dbl_recv_info Struct Reference	47
7.4.1	Detailed Description	47
7.4.2	Field Documentation	48
7.4.2.1	chan	48
7.4.2.2	chan_context	48
7.4.2.3	in_buffer	48

---

7.4.2.4	msg_len	48
7.4.2.5	sin_from	48
7.4.2.6	sin_to	48
7.4.2.7	timestamp	48
7.5	dbl_ticks_ Struct Reference	48
7.6	dbl_timespec Struct Reference	49
7.7	unit_protocol_info Struct Reference	49
7.7.1	Detailed Description	49

**Index****50**

# Chapter 1

## DBL

### 1.1 Introduction

DBL provides a very low-latency interface for sending and receiving UDP datagrams or TCP packets as part of the DBL extensions. The DBL library communicates directly with the firmware on the NIC to send and receive packets, removing the overhead associated with kernel calls and the TCP/UDP stack.

#### 1.1.1 Terms and Concepts

The DBL API uses 3 different entities: "devices", "channels", and "send handles".

A device is the abstraction of a NIC, and there will generally be one device per NIC in a given process. A device is created by calling `dbl_open()`. Several channels can attach to a device.

A channel is roughly the equivalent of a socket opened on a device, with a port number specified. A channel is created by calling `dbl_bind()` on a particular device. When calling `dbl_bind` the type of the channel (e.g TCP or UDP) must be specified.

A send handle is a handle associated with a specific destination that is used to very efficiently send packets to that destination. Send handles are not necessary for sending. A send handle is created by calling `dbl_send_connect()`.

Demultiplexing of incoming data on a device is done by the user code in order to reduce overhead in the library. There is a single call, `dbl_recvfrom()` that will return the next packet available from a given device. A buffer is passed into this function, and any received data will be placed into the buffer upon return. The received packet may be intended for any channel associated with the specified device. A device allows for the mix of UDP or TCP channels.

#### 1.1.2 Example Pseudo-Code

Example use cases:

A device is opened via a call to `dbl_open()`. An interface is specified to `dbl_open` via its first argument which is a struct `in_addr`. The DBL interface whose IP address matches this address will be opened and a device handle returned.

```
dbl_init();
dbl_open(interface, flags, &dev);
```

The following pseudo-code demonstrates typical multi-port receiver. For each port on which the program wished to receive data, a

`dbl_bind()` is used to bind a port to a channel. In this example, two different ports are bound, each with a different context value. The context is returned in the `dbl_receive_info` structure filled in by `dbl_recvfrom()` and can be used to demultiplex based on the receiving channel.

```
dbl_init();
dbl_open(interface, flags, &dev);
dbl_bind(dev, port1, flags, context1, &chan1);
dbl_bind(dev, port2, flags, context2, &chan2);
while (!done) {
    dbl_recvfrom(dev, mode, buf, maxlen, &info);
    user_packet_handler(buf, info.msg_len, info.chan_context);
}
```

The basic send function is `dbl_sendto()`. The following pseudo-code demonstrates sending a packet to a destination specified by the address parameter. address is a `sockaddr_in` as used by `socket sendto()`;

```
dbl_init();
dbl_open(interface, flags, &dev);
dbl_bind(dev, port1, flags, context1, &chan1);
dbl_sendto(chan1, address, buf, buflen, flags);
```

An alternate and slightly faster way to send can be used when you have a known set of destinations to which you are sending. A "send handle" is first created using `dbl_send_connect()`. A send handle is used internally to save precomputed information for sending to that particular destination.

```
dbl_init();
dbl_open(interface, flags, &dev);
dbl_bind(dev, port1, flags, context1, &chan1);
dbl_send_connect(chan1, address, flags, ttl, &send_handle);
dbl_send(send_handle, buf, buflen, flags);
```

To receive multicast packets, a channel joins the multicast group via `dbl_mcast_join()`.

```
dbl_init();
dbl_open(interface, flags, &dev);
dbl_bind(dev, port1, flags, context1, &chan1);
dbl_mcast_join(chan1, mcast_addr, NULL);
dbl_recvfrom(dev, mode, buf, maxlen, &info);
user_packet_handler(buf, info.msg_len, info.chan_context);
```

Each channel may join many multicast groups. The example below will receive packets sent to `mcast_addr1:port1`, `mcast_addr2:port1`, `mcast_addr1:port2`, and `mcast_addr3:port2`. The packets sent to port1 will have `context = context1` and those to port2 will have `context = context2`.

```
dbl_init();
dbl_open(interface, flags, &dev);
dbl_bind(dev, port1, flags, context1, &chan1);
dbl_bind(dev, port2, flags, context2, &chan2);
dbl_mcast_join(chan1, mcast_addr1, NULL);
dbl_mcast_join(chan1, mcast_addr2, NULL);
dbl_mcast_join(chan2, mcast_addr1, NULL);
dbl_mcast_join(chan2, mcast_addr3, NULL);
dbl_recvfrom(dev, mode, buf, maxlen, &info);
user_packet_handler(buf, info.msg_len, info.chan_context);
```



## 1.2 Interaction with Sockets

Since DBL packets move straight from the NIC to the user-level library, there is generally no opportunity for these packets to be shared with other processes using the socket interface. Thus, under default conditions, if a process using the DBL API and one using the socket API both open and bind to the same address (using appropriate REUSEADDR-style flags), only the DBL process will actually receive the packets. This is because the packets are never delivered to the kernel and the DBL process has no way to know that another process is listening for the packets.

In order to allow sockets-based processes to receive packets that are being received by DBL processes, the DBL process must not only specify the DBL\_BIND\_REUSEADDR flag to `dbl_bind()`, it must also specify the DBL\_BIND\_DUP\_TO\_KERNEL flag which will cause the firmware on the NIC to duplicate each packet to the kernel UDP stack for possible delivery to any sockets-based processes wishing to receive them. Note that this duplication will happen for every packet delivered to the socket address (IP and port number) specified in the call to `dbl_bind` with the DUP\_TO\_KERNEL flag, regardless of whether there is a socket application bound to the address or not.

Specifying DBL\_BIND\_DUP\_TO\_KERNEL will add 1.8 us or less to each packet whose destination is the address specified in the `dbl_bind()` call.

## 1.3 Receive Data Buffering

There are two different places that packets are buffered in DBL. The first level of buffering is a 48k buffer onboard on the NIC. This buffer is used directly by the hardware on the NIC and is serviced independently of activity on the host.

The second level of buffering is in host memory, and is on a per-device basis, since `dbl_recvfrom` reads from a `dbl_device_t`. This is a circular buffer which defaults to 128Mb on Linux (the size of the buffer can be changed, see `recvq_size` in `dbl_device_attrs` and `dbl_device_set_attrs`). The NIC asynchronously moves data into this buffer, and the only involvement required from the host is to drain data from this buffer.

On the host buffer, each packet has its length rounded up to a multiple of 64 bytes. Since ethernet packets are a minimum of 64 bytes on lengths and there is bookkeeping data included with the packet, each packet occupies a minimum of 128 bytes of buffer space. This translates to a worst-case capacity of one million packets, or 64 megabytes of data, or roughly 64 milliseconds worth of minimum-sized packets.

There are two different counters that indicate when packets are dropped due to lack of buffering. The first counter, "Net overflow drop" indicates that packets are arriving faster than the NIC can process them. The second counter, "Receive Queue full," indicates that the user application is not draining packets from the host queue quickly enough.

## 1.4 Direct Access to Ring Contents

The DBL API allows an application to directly access the packet buffers. An application can peek at a packet without consuming it. For example, a trading application could peek at the packet to search for a specific sequence or symbol. The application can then consume packets, without overhead, until a meaningful sequence or symbol is found, which may improve overall latency.

An application can operate on ring data without copying the payload. The API has functions to get a pointer to the first and next packets in the queue and pointers to the header and data sections of a given packet. There is also a function to consume the packet using two modes: DBL\_CONSUME\_SINGLE, which consumes a single packet at the head of the queue; and DBL\_CONSUME\_ALL, which consumes all outstanding packets. See the following functions

[dbl\\_eventq\\_open](#) [dbl\\_eventq\\_close](#) [dbl\\_eventq\\_peek\\_head](#) [dbl\\_eventq\\_peek\\_next](#) [dbl\\_eventq\\_inspect](#) [dbl\\_eventq\\_consume](#)

See the example programs in the bin/tests directory and described in the DBL User Guide: `dbl_ring_access` and `dbl_eventq`

## 1.5 Batch Receive

The DBL API allows an application to determine if further packets are pending. The API also provides a function `dbl_ext_rcvfrom` that allows you to receive multiple packets in one call instead of receiving packets sequentially, one at a time `dbl_rcvfrom`, minimizing overhead. See the following function: `dbl_ext_rcvfrom`

See the example program in the bin/tests directory and described in the DBL User Guide: `dbl_batch_rcv`

## 1.6 Arbitration duplicate packet elimination

Financial exchanges often use redundant multicast UDP feeds to disseminate pricing information. Since UDP is unreliable, the redundancy is used to reduce loss of information.

The application or middleware is normally responsible for eliminating the redundant packets and providing a single more reliable stream of packets to the trading engine. This process is coined "AB Arbitration".

While the path of UDP packets from the Exchange to the Phoenix is unreliable, the path from Phoenix to application is reliable, so it is appropriate to eliminate redundant packets within the Phoenix. Phoenix can do a best-effort optimization to eliminate redundant packets.

To enable the redundant packet elimination feature, 2 environment variables are used: `DBL_AB_PROTO_FILE` - full path name of the XML file describing packet formats used by supported Exchange feeds. `DBL_AB_FEEDS_FILE` - full path name of an XML file defining the multicast addresses and UDP ports for streams of units of an Exchange feed. See the DBL User Guide for more information.

## Chapter 2

# Module Index

### 2.1 API Reference

Here is a list of all modules:

API Reference . . . . .	11
Flags used for <code>dbl_open()</code> . . . . .	29
Flags used for <code>dbl_bind()</code> . . . . .	30
Flags for <code>dbl_send()</code> . . . . .	35
Extensions . . . . .	36
Specific Options for <code>dbl_ext_setchopt()</code> . . . . .	41



## Chapter 3

# Namespace Index

### 3.1 Namespace List

Here is a list of all documented namespaces with brief descriptions:

<a href="#">dbl</a> .....	43
---------------------------	----



# Chapter 4

## Data Structure Index

### 4.1 Data Structures

Here are the data structures with brief descriptions:

<a href="#">dbl__packet</a> . . . . .	45
<a href="#">dbl_device_attrs</a> . . . . .	45
<a href="#">dbl_ext_rcv_info</a>	
Information about the packet received . . . . .	46
<a href="#">dbl_rcv_info</a>	
Information about the packet received . . . . .	47
<a href="#">dbl_ticks_</a> . . . . .	48
<a href="#">dbl_timespec</a> . . . . .	49
<a href="#">unit_protocol_info</a> . . . . .	49





# Chapter 5

## Module Documentation

### 5.1 API Reference

API Reference for DBL.

#### Data Structures

- struct [dbl\\_device\\_attrs](#)
- struct [dbl\\_ext\\_rcv\\_info](#)  
*Information about the packet received.*
- struct [dbl\\_rcv\\_info](#)  
*Information about the packet received.*
- struct [unit\\_protocol\\_info](#)

#### Modules

- Flags used for [dbl\\_open\(\)](#)
- Flags used for [dbl\\_bind\(\)](#)
- Flags for [dbl\\_send\(\)](#).

#### Macros

- `#define DBL\_VERSION\_API 0x0006`

#### Enumerations

- enum [dbl\\_filter\\_mode](#) { [DBL\\_RECV\\_FILTER\\_NORMAL](#) = 0, [DBL\\_RECV\\_FILTER\\_ALLMULTI](#) = 1, [DBL\\_RECV\\_FILTER\\_RAW](#) = 2 }
- enum [dbl\\_ext\\_rcvmode](#) { [DBL\\_EXT\\_RECV\\_DEFAULT](#) = 0, [DBL\\_EXT\\_RECV\\_NONBLOCK](#) = 1, [DBL\\_EXT\\_RECV\\_COMPLETE](#) = 2 }

- enum `dbl_recvmode` {  
`DBL_RECV_DEFAULT = 0, DBL_RECV_NONBLOCK = 1, DBL_RECV_BLOCK = 2, DBL_RECV_PEEK = 3,`  
`DBL_RECV_PEEK_MSG = 4, DBL_RECV_TX_TIMESTAMP = 6, DBL_RECV_DEFAULT_RAW = 7 }`

## Functions

- `dbl_init` (uint16\_t api\_version)  
*Initializes the dbl library.*
- `dbl_open` (const struct in\_addr \*interface\_addr, int flags, dbl\_device\_t \*dev\_out)  
*Creates an instance of a dbl\_device.*
- `dbl_open_if` (const char \*ifname, int flags, dbl\_device\_t \*dev\_out)  
*Creates an instance of a dbl\_device.*
- `dbl_device_get_attrs` (dbl\_device\_t dev, struct dbl\_device\_attrs \*attr)
- `dbl_device_set_attrs` (dbl\_device\_t dev, const struct dbl\_device\_attrs \*attr)
- `dbl_device_enable` (dbl\_device\_t dev)
- `dbl_set_filter_mode` (dbl\_device\_t dev, enum dbl\_filter\_mode mode)
- `dbl_device_handle` (dbl\_device\_t dev)  
*Returns a descriptor for use with poll() or select().*
- `dbl_close` (dbl\_device\_t dev)  
*Close a dbl device.*
- `dbl_bind` (dbl\_device\_t dev, int flags, int port, void \*context, dbl\_channel\_t \*handle\_out)  
*Create a channel on dbl device.*
- `dbl_bind_addr` (dbl\_device\_t dev, const struct in\_addr \*ipaddr, int flags, int port, void \*context, dbl\_channel\_t \*handle\_out)  
*Creates a channel, using specified ip address.*
- `dbl_unbind` (dbl\_channel\_t handle)  
*Destroys a channel.*
- `dbl_getaddress` (dbl\_channel\_t ch, struct sockaddr\_in \*sin)  
*Returns the address to which a channel is bound.*
- `dbl_getticks` (dbl\_device\_t dev, dbl\_ticks\_t \*ticks)  
*Returns the current NIC time. It reports both values, NIC ticks and time in usec since epoch.*
- `dbl_gettime` (dbl\_device\_t dev, dbl\_timespec\_t \*ts)  
*Retrieve the current NIC clock value.*
- `dbl_mcast_join` (dbl\_channel\_t ch, const struct in\_addr \*mcast\_addr, void \*unused)  
*Join a multicast group.*
- `dbl_mcast_join_source` (dbl\_channel\_t ch, const struct in\_addr \*mcast\_addr, const struct in\_addr \*src)  
*Join a multicast group on a given source address.*
- `dbl_mcast_leave` (dbl\_channel\_t ch, const struct in\_addr \*mcast\_addr)  
*Leave a multicast group.*
- `dbl_mcast_leave_source` (dbl\_channel\_t ch, const struct in\_addr \*mcast\_addr, const struct in\_addr \*src)  
*Leave a multicast group.*
- `dbl_mcast_block_source` (dbl\_channel\_t ch, const struct in\_addr \*join\_addr, const struct in\_addr \*block\_addr)  
*block sender.*

- [dbl\\_mcast\\_unblock\\_source](#) (dbl\_channel\_t ch, const struct in\_addr \*join\_addr, const struct in\_addr \*block\_addr)
  - unblock sender.*
- [dbl\\_get\\_params](#) (dbl\_device\_t dev, void \*dbl\_params)
  - Used to query global and local DBL settings.*
- [dbl\\_shutdown](#) (dbl\_device\_t dev, int how)
  - Unblock dbl\_recvfrom/dbl\_ext\_recvmsg.*
- [dbl\\_ext\\_recvfrom](#) (dbl\_device\_t dev, enum [dbl\\_ext\\_recvmode](#) mode, int \*num, struct [dbl\\_ext\\_rcv\\_info](#) \*info)
  - Receive (UDP) data.*
- [dbl\\_recvfrom](#) (dbl\_device\_t dev, enum [dbl\\_recvmode](#) mode, void \*buf, size\_t len, struct [dbl\\_rcv\\_info](#) \*info)
  - Receive data.*
- [dbl\\_send\\_connect](#) (dbl\_channel\_t chan, const struct sockaddr\_in \*dest\_sin, int flags, int ttl, dbl\_send\_t \*hsend)
  - Create a send\_handle for faster sending.*
- [dbl\\_send](#) (dbl\_send\_t sendh, const void \*buf, size\_t len, int flags)
  - Send a packet using a send handle.*
- [dbl\\_send\\_disconnect](#) (dbl\_send\_t hsend)
  - Release a send handle.*
- [dbl\\_sendto](#) (dbl\_channel\_t ch, const struct sockaddr\_in \*sin, const void \*buf, size\_t len, int flags)
  - Send a packet.*
- [dbl\\_ab\\_set\\_seq](#) (dbl\_device\_t dev, struct sockaddr\_in \*sad, int seq\_num)
  - Reset the sequence number.*
- [dbl\\_ab\\_get\\_stream](#) (dbl\_device\_t dev, int unit\_id, int abcd, struct sockaddr\_in \*sad)
  - Return a stream given unit\_id and a/b/c/d.*
- [dbl\\_ab\\_get\\_unit](#) (dbl\_device\_t dev, struct sockaddr\_in \*sad, int \*unit\_id)
  - Return the unit\_id.*
- [dbl\\_ab\\_get\\_unit\\_protocol\\_info](#) (dbl\_device\_t dev, struct sockaddr\_in \*sad, struct [unit\\_protocol\\_info](#) \*upi)
  - Return the protocol information.*

### 5.1.1 Detailed Description

API Reference for DBL.

### 5.1.2 API Reference

### 5.1.3 Macro Definition Documentation

#### 5.1.3.1 #define DBL\_VERSION\_API 0x0006

DBL API version number (16 bits) Least significant byte increases for minor backwards compatible changes in the API. Most significant byte increases for incompatible changes in the API

0x0002: Added timestamp to [dbl\\_rcv\\_info](#) 0x0003: Added TCP extensions 0x0004: Added tx timestamping 0x0005: struct [dbl\\_rcv\\_info](#) changes 0x0006: struct [dbl\\_rcv\\_info](#) should not have changed, revert

## 5.1.4 Enumeration Type Documentation

### 5.1.4.1 enum dbl\_ext\_recvmode

Specifies behavior of the dbl\_ext\_recvfrom call

#### Enumerator

**DBL\_EXT\_RECV\_DEFAULT** Busy poll forever until at least one packet has been received.

**DBL\_EXT\_RECV\_NONBLOCK** Return a packet if available, else return EAGAIN.

**DBL\_EXT\_RECV\_COMPLETE** wait until all provided buffers are filled.

### 5.1.4.2 enum dbl\_filter\_mode

Filtering modes (advanced functionality).

#### Remarks

Selecting anything but the NORMAL filter causes all other DBL devices to be deprived of data. The ALLMULTI and RAW modes cause all matching data from the underlying port to be delivered to the one endpoint.

The OS-setting of dup to kernel is honored with all filtering modes, albeit with the same performance constraints.

### 5.1.4.3 enum dbl\_recvmode

Specifies behavior of the dbl\_recvfrom call

#### Enumerator

**DBL\_RECV\_DEFAULT** Busy poll forever until a packet is received.

**DBL\_RECV\_NONBLOCK** Return a packet if available, else return EAGAIN.

**DBL\_RECV\_BLOCK** Block until a packet is available, sleep until interrupt if necessary.

**DBL\_RECV\_PEEK** Check for a packet one time, return info, or EAGAIN if no packet.

**DBL\_RECV\_PEEK\_MSG** Peek but also copy data, return info, or EAGAIN if no packet. Unsupported in the DBL TCP extensions

**DBL\_RECV\_TX\_TIMESTAMP** nonblocking recv for retrieving full payloads and associated timestamps for outgoing (tx) traffic, restricted to dbl\_ext functions only, EINVAL reported otherwise

**DBL\_RECV\_DEFAULT\_RAW** blocking recv for retrieving full headers and payloads and associated timestamps for incoming (rx) traffic.

## 5.1.5 Function Documentation

### 5.1.5.1 dbl\_ab\_get\_stream ( dbl\_device\_t dep, int unit\_id, int abcd, struct sockaddr\_in \* sad )

Return a stream given unit\_id and a/b/c/d.

Return a stream pointer given a unique A/B arbitration unit\_id and stream.

#### Parameters

<i>dev</i>	The device
<i>unit_id</i>	The unique id assigned to this feed unit when parsed from the feeds file.
<i>abcd</i>	The positional definition of the stream in the feeds file.
<i>sad</i>	The argument sad is a pointer to a sockaddr structure. This structure is filled with stream address information upon return.

**Return values**

==	0 Success
==	EPERM operation not permitted because AB is not enabled
==	EINVAL <i>unit_id</i> not found or <i>abcd</i> > number of streams for this unit

**5.1.5.2 `dbl_ab_get_unit ( dbl_device_t dev, struct sockaddr_in * sad, int * unit_id )`**

Return the *unit\_id*.

Return the unique unit ID associated with a multicast socket address.

**Parameters**

<i>dev</i>	The device
<i>sad</i>	The argument sad is a pointer to a sockaddr structure. This structure is filled with the address of any of the streams which may be using the A/B arbiter.
<i>unit_id</i>	The unique identifier will be returned in this parameter

**Return values**

==	0 Success
==	EINVAL the sockaddr pointer was NULL
==	EPERM operation not permitted because AB is not enabled

**5.1.5.3 `dbl_ab_get_unit_protocol_info ( dbl_device_t dev, struct sockaddr_in * sad, struct unit_protocol_info * upi )`**

Return the protocol information.

Return protocol information associated with the given multicast socket address.

**Parameters**

<i>dev</i>	The device
<i>sad</i>	The argument sad is a pointer to a sockaddr structure. This structure is filled with the address of any of the streams which may be using the A/B arbiter.
<i>upi</i>	Protocol information associated with provided multicast address will be returned here

**Return values**

==	0 Success
==	EINVAL the sockaddr pointer was NULL
==	EPERM operation not permitted because AB is not enabled

**5.1.5.4 dbl\_ab\_set\_seq ( dbl\_device\_t dev, struct sockaddr\_in \* sad, int seq\_num )**

Reset the sequence number.

Reset the sequence number in an A/B arbitration unit.

**Parameters**

<i>dev</i>	The device
<i>sad</i>	The argument sad is a pointer to a sockaddr structure. This structure is filled with the address of any of the streams using the A/B arbiter. When addr is NULL, the seq_num will be applied to all arbitration units.
<i>seq_num</i>	The seq_num is the next expected sequence number for a packet.

**Return values**

==	0	Success
!=	0	OS return code
==	EPERM	operation not permitted because AB is not enabled

**5.1.5.5 dbl\_bind ( dbl\_device\_t dev, int flags, int port, void \* context, dbl\_channel\_t \* handle\_out )**

Create a channel on dbl device.

Creates a channel on a specified device through which UDP datagrams or TCP streams (if using the DBL TCP extensions), may be sent and received. Any packets sent through this channel will have "port" as their source port and packets arriving on the interface addressed to "port" will be received on this channel. By default, only unicast packets, not broadcast or multicast, will be received on the channel.

**Parameters**

<i>dev</i>	A DBL device handle returned by a call to <a href="#">dbl_open()</a> .
<i>flags</i>	See <a href="#">Flags used for dbl_bind()</a> .
<i>port</i>	The port to send/receive on.
<i>context</i>	The value of context is returned on future receives on this channel.
<i>handle_out</i>	The handle to the created channel.

**Return values**

0	Success
<i>EINVAL</i>	Error in arguments
<i>EEXIST</i>	port already in use
?	Other values indicate various OS failures in the bind process

If [dbl\\_bind\(\)](#) is called multiple times on the same port on a single device, unicast packets will only be delivered to the oldest channel currently bound to the port.

**Remarks**

This function can be used in the context of DBL TCP API, with some restriction. The DBL\_BIND\_DUP\_TO\_KERNEL and DBL\_BIND\_NO\_UNICAST options are not supported.

**5.1.5.6** `dbl_bind_addr ( dbl_device_t dev, const struct in_addr * ipaddr, int flags, int port, void * context, dbl_channel_t * handle_out )`

Creates a channel, using specified ip address.

Creates a channel on a specified device, just like `dbl_bind`, except that it associates the channel with the specified address instead of the one specified in the `dbl_open` call.

The address used must correspond to an OS-level interface that maps to the same underlying Ethernet port as the interface specified in `dbl_open`. For example, this can be a VLAN interface.

**Parameters**

<i>dev</i>	A DBL device handle returned by a call to <code>dbl_open()</code> .
<i>ipaddr</i>	Specifies the IP address of the interface with which the channel created will be associated. This must be on the same underlying interface as the one used in the <code>dbl_open</code> call.
<i>flags</i>	See <a href="#">Flags used for dbl_bind()</a> .
<i>port</i>	The port to send/receive on.
<i>context</i>	The value of context is returned on future receives on this channel.
<i>handle_out</i>	The handle to the created channel.

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	Error in arguments. Specifying an address that is not on the same underlying interface as that specified with <code>dbl_open</code> will return <code>EINVAL</code> .
<i>EEXIST</i>	port already in use
<i>?</i>	Other values indicate various OS failures in the bind process

**Remarks**

DBL TCP supported

**5.1.5.7** `dbl_close ( dbl_device_t dev )`

Close a dbl device.

Terminate usage of a device returned by `dbl_open()` and free all resources associated with it.

**Parameters**

<i>dev</i>	The device handle returned from <code>dbl_open()</code> .
------------	---

**Return values**

<i>0</i>	Success
----------	---------

**5.1.5.8** `dbl_device_enable ( dbl_device_t dev )`

Function to enable a device if opened with `DBL_OPEN_DISABLED`

**Remarks**

If this call fails, the user is still responsible for calling [dbl\\_close\(\)](#) on the underlying device to free resources

**5.1.5.9 dbl\_device\_get\_attrs ( dbl\_device\_t dev, struct dbl\_device\_attrs \* attr )**

Function to retrieve device attributes.

**Parameters**

<i>dev</i>	The device handle returned from <a href="#">dbl_open()</a>
<i>attr</i>	Device attributes will be copied out.

**Remarks**

Can be used before and after calls that open and enable DBL devices.

**5.1.5.10 dbl\_device\_handle ( dbl\_device\_t dev )**

Returns a descriptor for use with poll() or select().

Returns an OS-specific file descriptor which can be passed to poll() or select() to block on receive data available. For UNIX systems, this is a file descriptor, on Windows it is a HANDLE.

**Parameters**

<i>dev</i>	The DBL device whose OS handle is needed.
------------	---

**Returns**

OS-specific handle for device

**5.1.5.11 dbl\_device\_set\_attrs ( dbl\_device\_t dev, const struct dbl\_device\_attrs \* attr )**

Function to set device attributes before a device is enabled

**Parameters**

<i>dev</i>	The device handle returned from <a href="#">dbl_open()</a> with flag DBL_OPEN_DISABLED.
<i>attr</i>	Device attributes that will be set on the device.

**Remarks**

Can't be called without having the contents of attr previously filled out by a call to [dbl\\_device\\_get\\_attrs](#). The implementation can change the size of requests to accommodate internal alignment and sizing requirements. If these sizes are changed, the new sizes are reflected during a subsequent call to [dbl\\_device\\_get\\_attrs](#).



**5.1.5.12 `dbl_ext_recvfrom ( dbl_device_t dev, enum dbl_ext_recvmode mode, int * num, struct dbl_ext_recv_info * info )`**

Receive (UDP) data.

Used to check for and read data from the channels associated with a particular `dbl_device`.

**Parameters**

<i>dev</i>	The underlying device via <code>dbl_open</code>
<i>mode</i>	See <a href="#">dbl_ext_recvmode</a>
<i>num</i>	In / Out value of given buffers pointed by struct <a href="#">dbl_ext_recv_info</a> and buffers detected via underlying receive operations.
<i>info</i>	A pointer to an array. See <a href="#">dbl_ext_recv_info</a> .

**Return values**

<i>0</i>	Success
<i>EAGAIN</i>	Returned if using mode <code>DBL_RECV_NONBLOCK</code> or <code>DBL_RECV_COMPLETE</code> when no packet is available.
<i>EINTR</i>	in case <a href="#">dbl_shutdown()</a> was called
<i>ENOBUS</i>	provided buffer was filled successfully but more data is pending.
<i>?</i>	Other codes indicate various OS failures.

**Remarks**

`dbl_ext_recvfrom()` will, by default, busy-poll checking for data available on the device. This consumes 100% of the CPU available to this single thread, but also guarantees the lowest possible latency for packet delivery.  
DBL TCP not supported.

**5.1.5.13 `dbl_get_params ( dbl_device_t dev, void * dbl_params )`**

Used to query global and local DBL settings.

Used to query global and local DBL settings

**Parameters**

<i>dev</i>	The underlying device via <code>dbl_open</code>
<i>params</i>	struct pointer

**Remarks**

DBL UDP and TCP

**Return values**

<i>0</i>	Success, parameter <code>len</code> updated with actual length
----------	--

**5.1.5.14 `dbl_getaddress ( dbl_channel_t ch, struct sockaddr_in * sin )`**

Returns the address to which a channel is bound.

Returns the address to which a channel is bound.

**Parameters**

<i>ch</i>	Specifies the channel whose bind information is required.
<i>sin</i>	sockaddr_in to which the address will be copied out.

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	Bad channel specified

**Remarks**

DBL TCP supported

**5.1.5.15 db\_l\_getticks ( db\_l\_device\_t dev, db\_l\_ticks\_t \* ticks )**

Returns the current NIC time. It reports both values, NIC ticks and time in usec since epoch.

Returns the current NIC time.

**Parameters**

<i>dev</i>	Specifies the dev channel from db_l_open
<i>ticks</i>	Specifies the db_l_ticks_t structure holding the timing information

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	Bad dev specified

**Remarks**

DBL TCP supported

Under TA , a ioctl/WSAIoctl socket call can use cmd SIO\_GETNICTIME

**5.1.5.16 db\_l\_gettime ( db\_l\_device\_t dev, db\_l\_timespec\_t \* ts )**

Retrieve the current NIC clock value.

Retrieves the current NIC clock value by reading the timestamp register on the NIC.

**Parameters**

<i>dev</i>	The dev channel from db_l_open.
<i>ts</i>	db_l_timespec_t that will hold the clock value upon return.

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	Argument error such as bad dev.

**Remarks**

Reading the timestamp register is a blocking read operation that goes over local interconnect and may have non-trivial delays on the order of 1 microsecond. This method does not adjust the clock value read from the NIC in any way.

**5.1.5.17 dbl\_init ( uint16\_t api\_version )**

Initializes the dbl library.

Initializes the dbl library.

**Parameters**

<i>api_version</i>	Must always be <code>DBL_VERSION_API</code> . This is used to ensure compatability between the application binary and the DBL library.
--------------------	--

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	Bad/incompatible version passed.

**Remarks**

`dbl_init()` must be called once at the start of any application that uses DBL.

**5.1.5.18 dbl\_mcast\_block\_source ( dbl\_channel\_t ch, const struct in\_addr \* join\_addr, const struct in\_addr \* block\_addr )**

block sender.

Indicates that the specified channel wishes to stop receiving packets from a given source and therefore block that sender  
Prerequisites : prior call to `dbl_mcast_join` on same multicast address.

**Parameters**

<i>ch</i>	Handle for the channel to leave the specified multicast group.
<i>join_addr</i>	Address of the multicast group to join.
<i>block_addr</i>	Address to block. The multicast packets will not be received from the blocked source

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	Argument error, such as address not multicast group.
<i>EADDRNOTAVAIL</i>	Not currently joined to group "address"
<i>EAGAIN</i>	internal resources temporarily unavailable, try again.
<i>?</i>	Other non-zero codes indicate various OS failures in the leave process

**5.1.5.19 dbl\_mcast\_join ( dbl\_channel\_t ch, const struct in\_addr \* mcast\_addr, void \* unused )**

Join a multicast group.

Indicates that the specified channel wishes to receive packets addressed to the multicast address specified.

**Parameters**

<i>ch</i>	Handle for the channel to add to the specified multicast group.
<i>mcast_addr</i>	Address of the multicast group to join.

**Return values**

0	Success
<i>EINVAL</i>	Argument error, such as address is not a multicast group.
?	Other values indicate various OS specific failures in the join process.

**5.1.5.20 dbl\_mcast\_join\_source ( dbl\_channel\_t ch, const struct in\_addr \* mcast\_addr, const struct in\_addr \* src )**

Join a multicast group on a given source address.

Indicates that the specified channel wishes to receive packets addressed to the multicast address specified from a specific source. For multiple sources, call this function again with the desired sources to receive from.

**Parameters**

<i>ch</i>	Handle for the channel to add to the specified multicast group.
<i>mcast_addr</i>	Address of the multicast group to join.
<i>src</i>	Address of source to receive multicast from

**Return values**

0	Success
<i>EINVAL</i>	Argument error, such as address is not a multicast group.
?	Other values indicate various OS specific failures in the join process.

**5.1.5.21 dbl\_mcast\_leave ( dbl\_channel\_t ch, const struct in\_addr \* mcast\_addr )**

Leave a multicast group.

Indicates that the specified channel wishes to stop receiving packets addressed to the multicast address specified.

**Parameters**

<i>ch</i>	Handle for the channel to leave the specified multicast group.
<i>mcast_addr</i>	Address of the multicast group to leave.

**Return values**

0	Success
<i>EINVAL</i>	Argument error, such as address not multicast group.
<i>EADDRNOTAVAIL</i>	Not currently joined to group "address"
<i>EAGAIN</i>	internal resources temporarily unavailable, try again.
?	Other non-zero codes indicate various OS failures in the leave process

**5.1.5.22 dbl\_mcast\_leave\_source ( dbl\_channel\_t ch, const struct in\_addr \* mcast\_addr, const struct in\_addr \* src )**

Leave a multicast group.

Indicates that the specified channel wishes to stop receiving packets addressed to the multicast address specified.

**Parameters**

<i>ch</i>	Handle for the channel to leave the specified multicast group.
<i>mcast_addr</i>	Address of the multicast group to leave.
<i>src</i>	Address of the source to drop

**Return values**

0	Success
<i>EINVAL</i>	Argument error, such as address not multicast group.
<i>EADDRNOTAVAIL</i>	Not currently joined to group "address"
<i>EAGAIN</i>	internal resources temporarily unavailable, try again.
?	Other non-zero codes indicate various OS failures in the leave process

**5.1.5.23 dbl\_mcast\_unblock\_source ( dbl\_channel\_t ch, const struct in\_addr \* join\_addr, const struct in\_addr \* block\_addr )**

unblock sender.

Indicates that the specified channel wishes to unblock a sender. Receiving packets will commence from the unblocked sender Prerequisites : prior call to dbl\_mcast\_join on same multicast address. Prior call to dbl\_mcast\_block\_source.

**Parameters**

<i>ch</i>	Handle for the channel to leave the specified multicast group.
<i>join_addr</i>	Address of the multicast group to join.
<i>block_addr</i>	Address to unblock. The multicast packets will again be received from the unblocked source

**Return values**

0	Success
<i>EINVAL</i>	Argument error, such as address not multicast group.
<i>EADDRNOTAVAIL</i>	Not currently joined to group "address"
<i>EAGAIN</i>	internal resources temporarily unavailable, try again.
?	Other non-zero codes indicate various OS failures in the leave process

**5.1.5.24 dbl\_open ( const struct in\_addr \* interface\_addr, int flags, dbl\_device\_t \* dev\_out )**

Creates an instance of a dbl\_device.

Creates an instance of a dbl device which can be used to subsequently open channels via [dbl\\_bind\(\)](#).

**Parameters**

<i>interface_addr</i>	Specifies the IP address of the interface with which channels created using <a href="#">dbl_bind()</a> will be associated.
-----------------------	--

<i>flags</i>	A bitmask of flags to alter open behavior. See <a href="#">Flags used for dbf_open()</a>
<i>dev_out</i>	On successful return, this is where the handle for the newly opened device will be placed.

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	bad usage. includes dbf_init not called first and bad interface_addr.
<i>ENODEV</i>	no matching IP address found on DBL-enabled NIC
<i>EAGAIN</i>	internal resources temporarily unavailable, try again.

**Remarks**

Unlike traditional sockets, a DBL channel cannot be associated with multiple network interfaces. Using the TCP extensions, dbf\_open opens an endpoint on which several channels of type UDP and TCP can be demultiplexed

**5.1.5.25 dbf\_open\_if ( const char \* ifname, int flags, dbf\_device\_t \* dev\_out )**

Creates an instance of a dbf\_device.

Like dbf\_open () except it takes an interface name instead of an ip address.

**Parameters**

<i>ifname</i>	Specifies the name of the interface with which channels created using <a href="#">dbf_bind()</a> will be associated.
<i>flags</i>	A bitmask of flags to alter open behavior. See <a href="#">Flags used for dbf_open()</a>
<i>dev_out</i>	On successful return, this is where the handle for the newly opened device will be placed.

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	bad usage. includes dbf_init not called first and bad interface_addr.
<i>EAGAIN</i>	internal resources temporarily unavailable, try again.

**Remarks**

Unlike traditional sockets, a DBL channel cannot be associated with multiple network interfaces.

**5.1.5.26 dbf\_recvfrom ( dbf\_device\_t dev, enum dbf\_recvmode mode, void \* buf, size\_t len, struct dbf\_recv\_info \* info )**

Receive data.

Used to check for and read data from the channels associated with a particular dbf\_device.

**Parameters**

<i>dev</i>	The underlying device via dbf_open
<i>mode</i>	See <a href="#">dbf_recvmode</a>
<i>buf</i>	Buffer in which to place received data.
<i>len</i>	Maximum number of bytes to write into buf.
<i>info</i>	See <a href="#">dbf_recv_info</a> .

**Return values**

<i>0</i>	Success
<i>EAGAIN</i>	Returned if using mode DBL_RECV_NONBLOCK or DBL_RECV_PEEK when no packet is available.
<i>EINTR</i>	in case <a href="#">dbl_shutdown()</a> was called
<i>?</i>	Other codes indicate various OS failures.

**Remarks**

[dbl\\_recvfrom\(\)](#) will, by default, busy-poll checking for data available on the device. This consumes 100% of the CPU available to this single thread, but also guarantees the lowest possible latency for packet delivery. A blocking mode of operation may be specified through the *recv\_mode* parameter, reducing CPU load at the expense of a few microseconds of message latency.

DBL TCP supported. Receiving a return value of 0 with a *msg\_len* of 0 means the channel is disconnected.

On endpoints with mixed channels e.g DBL and DBL extension (TCP) channels the DBL channels are prioritized to avoid packet drops

**5.1.5.27 `dbl_send ( dbl_send_t sendh, const void * buf, size_t len, int flags )`**

Send a packet using a send handle.

Sends a packet to the address associated with the specified send handle. The *send\_handle* must have been previously created by a call to [dbl\\_send\\_connect\(\)](#). If internal resources are unavailable to execute the send immediately, the send call will block until resources are available to proceed.

**Parameters**

<i>sendh</i>	Send handle specifying destination for packe.
<i>buf</i>	The data to send.
<i>len</i>	The number of bytes to send.
<i>flags</i>	See <a href="#">Flags for dbl_send()</a> ..

**Return values**

<i>0</i>	Success
<i>EAGAIN</i>	DBL_NONBLOCK specified and no resources available.
<i>EAGAIN</i>	SO_TIMESTAMPING option enabled and no TX timestamping queue exhausted. An application would need to retrieve TX packets via the DBL_RECV_MSG_ERRQUEUE at this time to be able to send again or disable the TX timestamping option.
<i>?</i>	Other codes indicate various OS failures in the send process.

**Remarks**

DBL TCP supported with no DBL flags. The function will block until all data has been transferred. For advanced handling use `dbl_ext_send` for TCP channels

**Return values**

<i>EISCONN</i>	channel already connected
----------------	---------------------------

**5.1.5.28 dbl\_send\_connect ( dbl\_channel\_t chan, const struct sockaddr\_in \* dest\_sin, int flags, int ttl, dbl\_send\_t \* hsend )**

Create a send\_handle for faster sending.

Used to create a send handle for fast sending to a remote destination.

**Parameters**

<i>chan</i>	The channel to be associated with this send handle.
<i>dest_sin</i>	Destination address of packets sent using this handle.
<i>flags</i>	Bitmask of flags to modify default send_connect operation Currently no flags are supported.
<i>ttl</i>	The value to put in the TTL field of the IP header.
<i>hsend</i>	The send_handle to be used in future calls to <a href="#">dbl_send()</a> is returned here.

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	Errors in arguments
<i>?</i>	Other codes indicate various OS failures in the send process.

**Remarks**

The returned send handle is a reference to a set of precomputed data that is needed to send a packet to a particular destination. This precomputed data is saved and cached by DBL as a matter of course through the [dbl\\_sendto\(\)](#) function, but holding a send\_handle avoids the need for a hash lookup to find the necessary information. This can take 100-200 ns off the time required to do a send.

Since `dbl_send_connect` will re-use a cached send handle to the same destination, the `ttl` parameter, if non-zero, will overwrite the `ttl` value in the cached sendhandle. This means that any future `dbl_sendto` operations to the same destination will use the new `ttl` value. This also means that if there is a need to use `dbl_sendto` with a different `ttl` than the default, it is possible to use a call to `dbl_send_connect` to change the `ttl`.

DBL TCP supported. One can use the `dbl` semantics (reuse the exact same call, besides the `ttl` value) to retrieve a send handle, or one can specify a `NULL` value for `dest_sin` to retrieve a new send handle which could be clearer in the code than keeping the `dest_sin` value.

**5.1.5.29 dbl\_send\_disconnect ( dbl\_send\_t hsend )**

Release a send handle.

Release the resources associated with a send handle.

**Parameters**

<i>hsend</i>	The send handle.
--------------	------------------

**Return values**

<i>0</i>	Success
----------	---------

**Remarks**

DBL TCP supported - in this case the connected peer will receive an EOF which will show up with a `msg` of `len 0`. The local channel is re-transitioned into the unconnected state and can be used again in `dbl_send_connect`



**5.1.5.30 db\_l\_sendto ( db\_l\_channel\_t *ch*, const struct sockaddr\_in \* *sin*, const void \* *buf*, size\_t *len*, int *flags* )**

Send a packet.

Send a packet to the address specified.

**Parameters**

<i>ch</i>	Handle for the channel to send over.
<i>sin</i>	The destination address
<i>buf</i>	The data to send.
<i>len</i>	The length of the data to send.
<i>flags</i>	See <a href="#">Flags for db_l_send()</a> .

**Return values**

0	Success
EAGAIN	DBL_NONBLOCK specified and no resources available.
?	Other codes indicate various OS failures in the send process.

**5.1.5.31 db\_l\_set\_filter\_mode ( db\_l\_device\_t *dev*, enum db\_l\_filter\_mode *mode* )**

Function to control per-port DBL filtering modes (advanced functionality).

**5.1.5.32 db\_l\_shutdown ( db\_l\_device\_t *dev*, int *how* )**

Unblock db\_l\_recvfrom/db\_l\_ext\_recvmsg.

Used to unblock a blocking db\_l\_recvfrom/db\_l\_ext\_recvmsg.

**Parameters**

<i>dev</i>	The underlying device via db_l_open
<i>how</i>	Unused for now

**Remarks**

DBL UDP and TCP

**5.1.5.33 db\_l\_unbind ( db\_l\_channel\_t *handle* )**

Destroys a channel.

Destroys a channel and releases all the resources associated with it.

**Parameters**

<i>handle</i>	The handle of the channel to unbind.
---------------	--------------------------------------

**Return values**

0	Success
---	---------

**Remarks**

DBL TCP supported

## 5.2 Flags used for `dbl_open()`

### Macros

- `#define DBL_OPEN_THREADSAFE 0x1`
- `#define DBL_OPEN_DISABLED 0x2`
- `#define DBL_OPEN_HW_TIMESTAMPING 0x4`
- `#define DBL_OPEN_RAW_MODE 0x8`

### 5.2.1 Detailed Description

### 5.2.2 Macro Definition Documentation

#### 5.2.2.1 `#define DBL_OPEN_DISABLED 0x2`

A device can be opened but separately enabled through `dbl_device_enable`. This allows users to change the size of buffers or other properties before it is enabled and ready to receive packets. By setting this flag, users are required to separately call `dbl_device_enable` after, perhaps, having changed device attributes using `dbl_device_get_attrs` followed by `dbl_device_set_attrs`.

#### 5.2.2.2 `#define DBL_OPEN_HW_TIMESTAMPING 0x4`

Request that incoming packets provide a hardware timestamp to indicate when the packet was received by the NIC. The timestamp provided is a conversion from raw NIC nanoseconds to host nanoseconds as would be returned by `gettimeofday()`. Unless HW timestamping is requested, packets will return a timestamp of 0.

Alternatively, users can enable/disable the HW timestamping once the device is opened by using `dbl_device_get_attrs` followed by `dbl_device_set_attrs`.

Note, starting with DBL v4, this is enabled by default.

#### 5.2.2.3 `#define DBL_OPEN_RAW_MODE 0x8`

Request that incoming packets not only contain the payload but also fully qualified ethernet and IP headers. Payload then starts with an offset. An endpoint needs to be opened with `DBL_OPEN_RAW_MODE` to be able to use the `DBL_RECV_RAW` `recv` flag. Note that the following flags are not available when using RAW mode: `DBL_RECV_PEEK`, `DBL_RECV_PEEK_MSG`

#### 5.2.2.4 `#define DBL_OPEN_THREADSAFE 0x1`

Used to indicate that multiple threads will be using this device, and that locking should be used internally to serialize access. Thread safety is off by default in order to improve performance for the single-threaded case.

## 5.3 Flags used for dbl\_bind()

### Macros

- #define `DBL_BIND_REUSEADDR` 0x02
- #define `DBL_BIND_DUP_TO_KERNEL` 0x04
- #define `DBL_BIND_NO_UNICAST` 0x08
- #define `DBL_BIND_BROADCAST` 0x10
- #define `DBL_BIND_TX_TIMESTAMP` 0x20

### Enumerations

- enum `dbl_consumemode` { `DBL_CONSUME_SINGLE` = 0, `DBL_CONSUME_ALL` = 1 }
- enum `dbl_filtermode` { `DBL_ADD_FILTER` = 0, `DBL_REMOVE_FILTER` = 1 }

### Functions

- `dbl_raw_send` (dbl\_device\_t dep, uint8\_t \*hdr, int hdrlen, uint32\_t chksum\_hdr, int chksum\_offset, const uint8\_t \*buffer, int paylen, int flags)  
*Sends RAW packet on a DBL endpoint opened in RAW mode.*
- `dbl_eventq_open` (dbl\_device\_t dep)  
*Open dbl\_device for eventq functions.*
- `dbl_eventq_peek_head` (dbl\_device\_t dep, `dbl_packet_t` \*pkt)  
*Peek first packet.*
- `dbl_eventq_peek_next` (dbl\_device\_t dep, `dbl_packet_t` pkt, `dbl_packet_t` \*npkt)  
*Peek next packet.*
- `dbl_eventq_inspect` (dbl\_device\_t dev, `dbl_packet_t` pkt, char \*\*hdr, char \*\*data, int \*len, uint64\_t \*timestamp)  
*Get pointers for processing.*
- `dbl_eventq_consume` (dbl\_device\_t dev, `dbl_packet_t` pkt, enum `dbl_consumemode` consume\_mode)  
*Consume the packet. Advance on receive queue.*
- `dbl_eventq_close` (dbl\_device\_t dep)  
*Close a dbl device opened by `dbl_eventq_open()`.*
- `dbl_eventq_count` (dbl\_device\_t dep)  
*Count available packets ready to receive.*
- `dbl_set_filter` (dbl\_device\_t dep, enum `dbl_filtermode` how, int ip\_proto, struct `sockaddr_in` \*dst, int flags)  
*Set filters on a dbl raw endpoint.*

#### 5.3.1 Detailed Description

#### 5.3.2 Macro Definition Documentation

##### 5.3.2.1 #define DBL\_BIND\_BROADCAST 0x10

Allows this channel to receive broadcast packets.

**5.3.2.2 #define DBL\_BIND\_DUP\_TO\_KERNEL 0x04**

Allows packets to be shared with sockets. (See [Interaction with Sockets](#))

**5.3.2.3 #define DBL\_BIND\_NO\_UNICAST 0x08**

Instructs this channel not to receive packets addressed to the unicast address.

**5.3.2.4 #define DBL\_BIND\_REUSEADDR 0x02**

Allows other [dbl\\_bind\(\)](#) and [bind\(\)](#) calls on the same port to succeed.

**5.3.2.5 #define DBL\_BIND\_TX\_TIMESTAMP 0x20**

Allows this channel to receive tx timestamped packets. Requires DBL VERSION API 0x0004 or higher

**5.3.3 Function Documentation**

**5.3.3.1 dbl\_eventq\_close ( dbl\_device\_t dep )**

Close a dbl device opened by [dbl\\_eventq\\_open\(\)](#).

Terminate usage of a device opened by [dbl\\_eventq\\_open\(\)](#)

**Parameters**

<i>dev</i>	The device handle returned from <a href="#">dbl_open()</a> .
------------	--

**Return values**

0	Success
---	---------

**5.3.3.2 dbl\_eventq\_consume ( dbl\_device\_t dev, dbl\_packet\_t pkt, enum dbl\_consumemode consume\_mode )**

Consume the packet. Advance on receive queue.

Consume the packet

**Parameters**

<i>pkt</i>	reference packet to consume.
------------	------------------------------

**Returns**

0 successful EAGAIN if packet is not head. It will be able to be consumed successful the next time if the pkt became head.

**5.3.3.3 dbl\_eventq\_count ( dbl\_device\_t dep )**

Count available packets ready to receive.

Counts available packets in recv queue but won't be reporting any packet lengths.

**Parameters**

<i>dev</i>	A DBL device handle returned by a call to <a href="#">dbl_open()</a> .
------------	--

**Returns**

number of pending packets.

**5.3.3.4 dbl\_eventq\_inspect ( dbl\_device\_t dev, dbl\_packet\_t pkt, char \*\* hdr, char \*\* data, int \* len, uint64\_t \* timestamp )**

Get pointers for processing.

Get pointers to hdr and data for given packet

**Parameters**

<i>dev</i>	A DBL device handle returned by a call to <a href="#">dbl_open()</a> .
<i>pkt</i>	reference packet for inspection
<i>hdr</i>	(Ethernet) HDRs, can be NULL
<i>data</i>	pointer to payload, can be NULL
<i>len</i>	len of payload, can be NULL
<i>timestamp</i>	return timestamp of packet, can be NULL

**Returns**

0 successful

EINVAL if params don't match

**5.3.3.5 dbl\_eventq\_open ( dbl\_device\_t dep )**

Open dbl\_device for eventq functions.

Creates an instance of a dbl device which can be used to subsequently open channels via [dbl\\_bind\(\)](#).

Opens a device for use with eventq functions

Open a dbl device for eventq functions.

**Parameters**

<i>dev</i>	The device handle returned from <a href="#">dbl_open()</a> .
------------	--

**Return values**

0	Success
---	---------

**5.3.3.6 dbl\_eventq\_peek\_head ( dbl\_device\_t dep, dbl\_packet\_t \* pkt )**

Peek first packet.

Gets pointer to head of receive queue.

**Parameters**

<i>dev</i>	A DBL device handle returned by a call to <a href="#">dbl_open()</a> .
<i>pkt</i>	for first unread packet in queue

**Returns**

0 for success, ENOBUFS for empty queue

**5.3.3.7 dbl\_eventq\_peek\_next ( dbl\_device\_t dep, dbl\_packet\_t pkt, dbl\_packet\_t \* npkt )**

Peek next packet.

Gets pointer to following packet

**Parameters**

<i>dev</i>	A DBL device handle returned by a call to <a href="#">dbl_open()</a> .
<i>pkt</i>	reference packet for returning next packet to it
<i>npkt</i>	references to next packet

**Returns**

0 for success, ENOBUFS for no more packets

**5.3.3.8 dbl\_raw\_send ( dbl\_device\_t dep, uint8\_t \* hdr, int hdrlen, uint32\_t chksum\_hdr, int chksum\_offset, const uint8\_t \* buffer, int paylen, int flags )**

Sends RAW packet on a DBL endpoint opened in RAW mode.

Sends a fully qualified packet in raw format to device. When using a DBL RAW endpoint, this function expects the provided buffer containing a fully qualified ether frame.

**Parameters**

<i>dev</i>	A DBL RAW device handle returned by a call to <a href="#">dbl_open()</a> .
<i>hdr</i>	IPV4 HDR
<i>hdrlen</i>	
<i>chksum_hdr</i>	hdr checksum
<i>chksum_offset</i>	offset
<i>buffer</i>	payload
<i>paylen</i>	length
<i>flags</i>	e.g DBL_TX_LOOPBACK or DBL_MCAST_LOOPBACK

**5.3.3.9** `dbl_set_filter ( dbl_device_t dep, enum dbl_filtermode how, int ip_proto, struct sockaddr_in * dst, int flags )`

Set filters on a dbl raw endpoint.

Sets a filter for receiving raw packets. When using a DBL RAW endpoint, this function allows for passing in the a filter so that packets go up to user space.

**Parameters**

<i>dev</i>	A DBL RAW device handle returned by a call to <a href="#">dbl_open()</a> .
<i>how</i>	add or remove filter
<i>ip_proto</i>	UDP or TCP
<i>dst</i>	Destination IP address and port
<i>flags</i>	optional flags for filter DBL_BIND_DUP_TO_KERNEL

**Return values**

<i>0</i>	Success
<i>EINVAL</i>	DBL EP not opened in RAW mode



## 5.4 Flags for `dbl_send()`.

### Macros

- `#define DBL_NONBLOCK 0x4`
- `#define MSG_WARM 0x20000`
- `#define MSG_WARM 0x20000`
- `#define DBL_MCAST_LOOPBACK 0x10`
- `#define DBL_TX_LOOPBACK 0x8`

### 5.4.1 Detailed Description

### 5.4.2 Macro Definition Documentation

#### 5.4.2.1 `#define DBL_MCAST_LOOPBACK 0x10`

Loop back mcast data to local host (default off)

#### 5.4.2.2 `#define DBL_NONBLOCK 0x4`

Return EAGAIN if send request would block for resources

#### 5.4.2.3 `#define MSG_WARM 0x20000`

Optimize send() cost: Keep TCP send path warm. Data is not put on the wire.

#### 5.4.2.4 `#define MSG_WARM 0x20000`

Optimize send() cost: Keep TCP send path warm. Data is not put on the wire.

## 5.5 Extensions

API extensions for DBL.

### Modules

- [Specific Options for `dbl\_ext\_setopt\(\)`](#).

### Macros

- `#define DBL_FUNC(type) type`
- `#define DBL_VAR(type) type`
- `#define DBL_PROTO_IS_MTCP(flags) ((flags & (1 << 7)) != 0)`
- `#define DBL_TYPE_IS_TCP(flags) ((flags & (1 << 8)) != 0)`
- `#define DBL_INITFLAGS(type, proto) (type << 8 | proto << 7)`
- `#define DBL_TCP 1`
- `#define DBL_UDP 0`
- `#define DBL_BSD 1 /* use the BSD stack */`
- `#define DBL_MYRI 0 /* use the DBL_API for UDP */`
- `#define MSG_ERRQUEUE 0x2000`
- `#define DBL_CHANNEL_FLAGS(type, proto) DBL_INITFLAGS(type, proto)`

### Enumerations

- `enum { SO_TX_TIMESTAMPING = 0x0900 }`

### Functions

- `dbl_ext_send` (`dbl_channel_t ch`, `const void *buf`, `size_t paylen`, `int flags`, `int *nbytes`)  
*send on a channel and report number of bytes sent*
- `dbl_ext_accept` (`dbl_channel_t ch`, `struct sockaddr *sad`, `int *len`, `void *rcontext`, `dbl_channel_t *rch`)  
*Accept an incoming TCP connection, returns a new channel.*
- `dbl_ext_listen` (`dbl_channel_t ch`)  
*Allow for incoming connections/channels.*
- `dbl_ext_recv` (`dbl_channel_t ch`, `enum dbl_recvmode mode`, `void *buf`, `size_t len`, `struct dbl_recv_info *info`)  
*Receive data from a specific non-DBL channel.*
- `dbl_ext_recvmsg` (`dbl_device_t dev`, `enum dbl_recvmode recv_mode`, `struct dbl_recv_info **info`, `int recvmax`)  
*Receive data from many channels from a same device.*
- `dbl_ext_poll` (`dbl_channel_t *chs`, `int nchs`, `int timeout`)  
*Report on available incoming data on DBL channels.*
- `dbl_ext_getchopt` (`dbl_channel_t ch`, `int level`, `int optname`, `void *optval`, `socklen_t *optlen`)  
*DBL channels are using the same option semantics than in traditional socket environment.*
- `dbl_ext_setopt` (`dbl_channel_t ch`, `int level`, `int optname`, `const void *optval`, `socklen_t optlen`)

*DBL channels are using the same option semantics than in traditional socket environment.*

- `dbl_ext_channel_type` (`dbl_channel_t ch`)

*On a given channel TRUE is returned if the channel is TCP.*

### 5.5.1 Detailed Description

API extensions for DBL.

### 5.5.2 Introduction to extensions

### 5.5.3 Function Documentation

#### 5.5.3.1 `dbl_ext_accept ( dbl_channel_t ch, struct sockaddr * sad, int * len, void * rcontext, dbl_channel_t * rch )`

Accept an incoming TCP connection, returns a new channel.

Accepting incoming TCP channel connection demand.

#### Parameters

<i>ch</i>	The channel (from <code>dbl_bind()</code> ) on which connections are accepted
<i>sad</i>	The argument <i>sad</i> is a pointer to a <code>sockaddr</code> structure. This structure is filled with the address of the peer socket, as known to the communications layer. When <i>addr</i> is NULL, <i>addrlen</i> is not used, and should also be NULL.
<i>len</i>	The <i>len</i> argument is a value-result argument: the caller must initialize it to contain the size (in bytes) of the structure pointed to by <i>sad</i> ; on return it will contain the actual size of the peer address.
<i>rcontext</i>	The value of <i>rcontext</i> is associated with the new channel
<i>rch</i>	The channel which can be used to communicate with the remote peer.

#### Return values

<i>0</i>	Success
<i>?</i>	Other codes indicate various OS failures.

#### 5.5.3.2 `dbl_ext_channel_type ( dbl_channel_t ch )`

On a given channel TRUE is returned if the channel is TCP.

This call returns a bool on whether a channel is TCP or not

#### Parameters

<i>ch</i>	A valid channel
-----------	-----------------

#### Return values

<i>1</i>	Channel is TCP
<i>0</i>	Otherwise

**5.5.3.3 db\_ext\_getchopt ( db\_channel\_t ch, int level, int optname, void \* optval, socklen\_t \* optlen )**

DBL channels are using the same option semantics than in traditional socket environment.

This call is used to get information on DBLTCP channel options

**Parameters**

<i>ch</i>	The channel
<i>level</i>	Level of the option (IPPROTO_IP...)
<i>optname</i>	Option's name (IP_TTL...)
<i>optval</i>	The pointer on the value
<i>optlen</i>	The pointer on the option's length

**Return values**

==	0 Success
>	0 OS return code

**Remarks**

DBL channel can not be modified or any option read. A EOPNOTSUPP return code is given back to the user in that case.

**5.5.3.4 db\_ext\_listen ( db\_channel\_t ch )**

Allow for incoming connections/channels.

Used to transition the channel into the listening state

**Parameters**

<i>ch</i>	The channel (from <a href="#">db_bind()</a> )
-----------	---

**Return values**

0	Success
?	Other codes indicate various OS failures.

**5.5.3.5 db\_ext\_poll ( db\_channel\_t \* chs, int nchs, int timeout )**

Report on available incoming data on DBL channels.

Polling function for multiplexed channels, in/out channels in chs, timeout in mseconds

**Parameters**

<i>chs</i>	An array of channels to query
<i>nchs</i>	number of entries in the array
<i>timeout</i>	a timeout in milliseconds, -1 for INFINITE

**5.5.3.6** `dbl_ext_rcv ( dbl_channel_t ch, enum dbl_rcvmode mode, void * buf, size_t len, struct dbl_rcv_info * info )`

Receive data from a specific non-DBL channel.

Used to check for and read data from a specific channel if non-DBL channel

**Parameters**

<i>ch</i>	The channel (from <a href="#">dbl_bind()</a> ) on which a packet has been received.
<i>mode</i>	See <a href="#">dbl_rcvmode</a>
<i>buf</i>	Buffer in which to place received data.
<i>len</i>	Maximum number of bytes to write into buf.
<i>info</i>	See <a href="#">dbl_rcv_info</a> .

**Return values**

<i>0</i>	Success
<i>EAGAIN</i>	Returned if using mode DBL_RECV_NONBLOCK or DBL_RECV_PEEK when no packet is available.
<i>?</i>	Other codes indicate various OS failures.

**Remarks**

Receiving a return value of 0 with a `msg_len` of 0 means the channel is disconnected.

**5.5.3.7** `dbl_ext_rcvmsg ( dbl_device_t dev, enum dbl_rcvmode rcv_mode, struct dbl_rcv_info ** info, int rcvmax )`

Receive data from many channels from a same device.

Is the extension of a `rcvfrom`, but to load a array of receive information

**Parameters**

<i>dev</i>	The device
<i>rcv_mode</i>	See <a href="#">dbl_rcvmode</a>
<i>info</i>	the array which describes in/out parameters. The important parameters are: the void * <code>unused</code> field used to provide the pointer to the buffer where the data should be copied, the <code>msg_len</code> is an input-output param, describing then len of the buffer in input, and returning the len of the message copied (see <a href="#">dbl_rcv_info</a> )
<i>rcvmax</i>	the number of message which can be loaded

**Return values**

<i>&gt;=</i>	0 number of messages to retrieve in the info array
<i>&lt;</i>	0 error should be retrieved in <code>errno</code>

**Remarks**

Receiving a `msg_len` of 0 in the receive info structure means the channel returned is disconnected.

**5.5.3.8 dbl\_ext\_send ( dbf\_channel\_t ch, const void \* buf, size\_t paylen, int flags, int \* nbytes )**

send on a channel and report number of bytes sent

send on DBL extension channel

**Parameters**

<i>ch</i>	The connected channel
<i>buf</i>	pointer to buffer
<i>paylen</i>	size to send See <a href="#">Flags for dbf_send()</a> . return the number of bytes sent

**Return values**

0	Success
?	Other codes indicate various OS failures.

**5.5.3.9 dbf\_ext\_setopt ( dbf\_channel\_t ch, int level, int optname, const void \* optval, socklen\_t optlen )**

DBL channels are using the same option semantics than in traditional socket environment.

This call is used to set information on DBLTCP channel options

**Parameters**

<i>ch</i>	The channel
<i>level</i>	Level of the option (IPPROTO_IP...)
<i>optname</i>	Option's name (IP_TTL...) and specific options, see <a href="#">Specific Options for dbf_ext_setopt()</a> .
<i>optval</i>	The pointer on the value
<i>optlen</i>	The option's type length

**Return values**

==	0 Success
>	0 OS return code

**Remarks**

DBL channel can not be modified or any option read. A EOPNOTSUPP return code is given back to the user in that case.

## 5.6 Specific Options for `dbl_ext_setopt()`.

### Macros

- `#define SO_TIMESTAMPING 0x0025`

#### 5.6.1 Detailed Description

#### 5.6.2 Macro Definition Documentation

##### 5.6.2.1 `#define SO_TIMESTAMPING 0x0025`

Enable given channel to perform tx timestamping when sending. Retrieve timestamps via `DBL_MSG_ERRQUEUE`





## Chapter 6

# Namespace Documentation

### 6.1 dbi Namespace Reference

#### 6.1.1 Detailed Description

DBL

**Author**

CSP Inc.



## Chapter 7

# Data Structure Documentation

### 7.1 `dbl__packet` Struct Reference

#### Data Fields

- struct `dbl__ep` \* **dep**
- `uintptr_t` **uevt**
- `uintptr_t` **desc\_offset**
- `uintptr_t` **data\_offset**
- `uint64_t` **length**
- `void *` **pkt**
- `void *` **hdr**
- `void *` **payload**
- `int` **paylen**
- `uint64_t` **timestamp**

### 7.2 `dbl_device_attrs` Struct Reference

#### Data Fields

- `uint32_t` [recvq\\_filter\\_mode](#)
- `uint32_t` [recvq\\_size](#)
- `uint32_t` [hw\\_timestamping](#)
- `uint32_t` **reserved\_1**

#### 7.2.1 Detailed Description

Structure for retrieving and setting device attributes when [dbl\\_open](#) is opened with [DBL\\_OPEN\\_DISABLED](#).

## 7.2.2 Field Documentation

### 7.2.2.1 `uint32_t dbl_device_attrs::hw_timestamping`

Timestamp field is filled in for [dbl\\_rcv\\_info](#)

### 7.2.2.2 `uint32_t dbl_device_attrs::rcvq_filter_mode`

DBL receive filter mode, see [dbl\\_filter\\_mode](#)

### 7.2.2.3 `uint32_t dbl_device_attrs::rcvq_size`

Host receive queue size for device

## 7.3 `dbl_ext_rcv_info` Struct Reference

Information about the packet received.

### Data Fields

- `dbl_channel_t chan`
- `void * chan_context`
- `void * buf`
- `int buflen`
- `struct sockaddr_in sin_from`
- `struct sockaddr_in sin_to`
- `uint32_t msg_len`
- `uint64_t timestamp`

### 7.3.1 Detailed Description

Information about the packet received.

Information about the packet received.

### 7.3.2 Field Documentation

#### 7.3.2.1 `void* dbl_ext_rcv_info::buf`

The buffer is to be used for data

#### 7.3.2.2 `dbl_channel_t dbl_ext_rcv_info::chan`

The channel (from [dbl\\_bind\(\)](#)) on which a packet has been received

**7.3.2.3 void\* dbl\_ext\_rcv\_info::chan\_context**

The context value passed to [dbl\\_bind\(\)](#) when a receiving channel was created.

**7.3.2.4 uint32\_t dbl\_ext\_rcv\_info::msg\_len**

The actual transmitted length of the packet.

**7.3.2.5 struct sockaddr\_in dbl\_ext\_rcv\_info::sin\_from**

Source address of the received packet

**7.3.2.6 struct sockaddr\_in dbl\_ext\_rcv\_info::sin\_to**

Destination address of the received packet. This can be used to differentiate between packets to different multicast joins on the same channel.

**7.3.2.7 uint64\_t dbl\_ext\_rcv\_info::timestamp**

Timestamp in nanosecs when the packet was received by the adapter.

## 7.4 dbl\_rcv\_info Struct Reference

Information about the packet received.

### Data Fields

- `dbl_channel_t` [chan](#)
- `void *` [chan\\_context](#)
- `void *` [in\\_buffer](#)
- `struct sockaddr_in` [sin\\_from](#)
- `struct sockaddr_in` [sin\\_to](#)
- `uint32_t` [msg\\_len](#)
- `uint64_t` [timestamp](#)

### 7.4.1 Detailed Description

Information about the packet received.

Information about the packet received.

## 7.4.2 Field Documentation

### 7.4.2.1 `dbl_channel_t dbl_rcv_info::chan`

The channel (from `dbl_bind()`) on which a packet has been received

### 7.4.2.2 `void* dbl_rcv_info::chan_context`

The context value passed to `dbl_bind()` when a receiving channel was created.

### 7.4.2.3 `void* dbl_rcv_info::in_buffer`

The `in_buffer` is used in the extension of the DBL API to provide memory references in the `dbl*rcvmsg()` function.

### 7.4.2.4 `uint32_t dbl_rcv_info::msg_len`

The actual transmitted length of the packet. This may be greater than the number of bytes received if the length parameter is less than the actual number of bytes in the packet. In the case of the DBL TCP API, `msg_len` is an in-out parameter, used to fetch messages and given back to the user to indicate the length of the received packet.

### 7.4.2.5 `struct sockaddr_in dbl_rcv_info::sin_from`

Source address of the received packet

### 7.4.2.6 `struct sockaddr_in dbl_rcv_info::sin_to`

Destination address of the received packet. This can be used to differentiate between packets to different multicast joins on the same channel.

### 7.4.2.7 `uint64_t dbl_rcv_info::timestamp`

Timestamp in nanosecs when the packet was received by the adapter. Timestamping must have been enabled through `dbl_device_set_attr`

## 7.5 `dbl_ticks` Struct Reference

### Data Fields

- `uint64_t nic_ticks`
- `uint64_t host_nsecs`
- `uint64_t host_nsecs_delay`

## 7.6 `dbl_timespec` Struct Reference

### Data Fields

- long `tv_sec`
- long `tv_nsec`

## 7.7 `unit_protocol_info` Struct Reference

### Data Fields

- uint8\_t `header_offset`
- uint8\_t `header_length`
- uint8\_t `header_big_endian`
- uint8\_t `sequence_offset`
- uint8\_t `sequence_length`
- uint8\_t `sequence_big_endian`
- uint8\_t `message_count_offset`
- uint8\_t `message_count_length`
- uint8\_t `message_count_big_endian`

### 7.7.1 Detailed Description

This structure is used to transfer internal unit protocol information to the user when `dbl_ab_get_unit_protocol_info` is called.

# Index

## API Reference

- DBL\_EXT\_RECV\_COMPLETE, [14](#)
- DBL\_EXT\_RECV\_DEFAULT, [14](#)
- DBL\_EXT\_RECV\_NONBLOCK, [14](#)
- DBL\_RECV\_BLOCK, [14](#)
- DBL\_RECV\_DEFAULT, [14](#)
- DBL\_RECV\_DEFAULT\_RAW, [14](#)
- DBL\_RECV\_NONBLOCK, [14](#)
- DBL\_RECV\_PEEK, [14](#)
- DBL\_RECV\_PEEK\_MSG, [14](#)
- DBL\_RECV\_TX\_TIMESTAMP, [14](#)

## API Reference, [11](#)

- DBL\_VERSION\_API, [13](#)
- dbl\_ab\_get\_stream, [14](#)
- dbl\_ab\_get\_unit, [15](#)
- dbl\_ab\_get\_unit\_protocol\_info, [15](#)
- dbl\_ab\_set\_seq, [15](#)
- dbl\_bind, [16](#)
- dbl\_bind\_addr, [16](#)
- dbl\_close, [17](#)
- dbl\_device\_enable, [17](#)
- dbl\_device\_get\_attrs, [18](#)
- dbl\_device\_handle, [18](#)
- dbl\_device\_set\_attrs, [18](#)
- dbl\_ext\_recvfrom, [18](#)
- dbl\_ext\_recvmode, [14](#)
- dbl\_filter\_mode, [14](#)
- dbl\_get\_params, [19](#)
- dbl\_getaddress, [19](#)
- dbl\_getticks, [20](#)
- dbl\_gettime, [20](#)
- dbl\_init, [21](#)
- dbl\_mcast\_block\_source, [21](#)
- dbl\_mcast\_join, [21](#)
- dbl\_mcast\_join\_source, [22](#)
- dbl\_mcast\_leave, [22](#)
- dbl\_mcast\_leave\_source, [22](#)
- dbl\_mcast\_unblock\_source, [23](#)
- dbl\_open, [23](#)
- dbl\_open\_if, [24](#)
- dbl\_recvfrom, [24](#)

- dbl\_recvmode, [14](#)
- dbl\_send, [25](#)
- dbl\_send\_connect, [25](#)
- dbl\_send\_disconnect, [26](#)
- dbl\_sendto, [26](#)
- dbl\_set\_filter\_mode, [27](#)
- dbl\_shutdown, [27](#)
- dbl\_unbind, [27](#)

## buf

- dbl\_ext\_recv\_info, [46](#)

## chan

- dbl\_ext\_recv\_info, [46](#)
- dbl\_recv\_info, [48](#)

## chan\_context

- dbl\_ext\_recv\_info, [46](#)
- dbl\_recv\_info, [48](#)

## DBL\_EXT\_RECV\_COMPLETE

- API Reference, [14](#)

## DBL\_EXT\_RECV\_DEFAULT

- API Reference, [14](#)

## DBL\_EXT\_RECV\_NONBLOCK

- API Reference, [14](#)

## DBL\_RECV\_BLOCK

- API Reference, [14](#)

## DBL\_RECV\_DEFAULT

- API Reference, [14](#)

## DBL\_RECV\_DEFAULT\_RAW

- API Reference, [14](#)

## DBL\_RECV\_NONBLOCK

- API Reference, [14](#)

## DBL\_RECV\_PEEK

- API Reference, [14](#)

## DBL\_RECV\_PEEK\_MSG

- API Reference, [14](#)

## DBL\_RECV\_TX\_TIMESTAMP

- API Reference, [14](#)

## DBL\_BIND\_BROADCAST

- Flags used for `dbl_bind()`, [30](#)

## DBL\_BIND\_REUSEADDR



Flags used for `dbl_bind()`, 31

`DBL_MCAST_LOOPBACK`  
Flags for `dbl_send()`, 35

`DBL_NONBLOCK`  
Flags for `dbl_send()`, 35

`DBL_OPEN_DISABLED`  
Flags used for `dbl_open()`, 29

`DBL_OPEN_RAW_MODE`  
Flags used for `dbl_open()`, 29

`DBL_VERSION_API`  
API Reference, 13

`dbl`, 43

`dbl__packet`, 45

`dbl_ab_get_stream`  
API Reference, 14

`dbl_ab_get_unit`  
API Reference, 15

`dbl_ab_get_unit_protocol_info`  
API Reference, 15

`dbl_ab_set_seq`  
API Reference, 15

`dbl_bind`  
API Reference, 16

`dbl_bind_addr`  
API Reference, 16

`dbl_close`  
API Reference, 17

`dbl_device_attrs`, 45  
hw\_timestamping, 46  
recvq\_filter\_mode, 46  
recvq\_size, 46

`dbl_device_enable`  
API Reference, 17

`dbl_device_get_attrs`  
API Reference, 18

`dbl_device_handle`  
API Reference, 18

`dbl_device_set_attrs`  
API Reference, 18

`dbl_eventq_close`  
Flags used for `dbl_bind()`, 31

`dbl_eventq_consume`  
Flags used for `dbl_bind()`, 31

`dbl_eventq_count`  
Flags used for `dbl_bind()`, 31

`dbl_eventq_inspect`  
Flags used for `dbl_bind()`, 32

`dbl_eventq_open`  
Flags used for `dbl_bind()`, 32

`dbl_eventq_peek_head`  
Flags used for `dbl_bind()`, 32

`dbl_eventq_peek_next`  
Flags used for `dbl_bind()`, 33

`dbl_ext_accept`  
Extensions, 37

`dbl_ext_channel_type`  
Extensions, 37

`dbl_ext_getchopt`  
Extensions, 37

`dbl_ext_listen`  
Extensions, 38

`dbl_ext_poll`  
Extensions, 38

`dbl_ext_recv`  
Extensions, 38

`dbl_ext_recv_info`, 46  
buf, 46  
chan, 46  
chan\_context, 46  
msg\_len, 47  
sin\_from, 47  
sin\_to, 47  
timestamp, 47

`dbl_ext_recvfrom`  
API Reference, 18

`dbl_ext_recvmode`  
API Reference, 14

`dbl_ext_recvmsg`  
Extensions, 39

`dbl_ext_send`  
Extensions, 39

`dbl_ext_setchopt`  
Extensions, 40

`dbl_filter_mode`  
API Reference, 14

`dbl_get_params`  
API Reference, 19

`dbl_getaddress`  
API Reference, 19

`dbl_getticks`  
API Reference, 20

`dbl_gettime`  
API Reference, 20

`dbl_init`  
API Reference, 21

`dbl_mcast_block_source`  
API Reference, 21

`dbl_mcast_join`  
API Reference, 21

`dbl_mcast_join_source`

- API Reference, 22
- dbl\_mcast\_leave
  - API Reference, 22
- dbl\_mcast\_leave\_source
  - API Reference, 22
- dbl\_mcast\_unblock\_source
  - API Reference, 23
- dbl\_open
  - API Reference, 23
- dbl\_open\_if
  - API Reference, 24
- dbl\_raw\_send
  - Flags used for dbl\_bind(), 33
- dbl\_rcv\_info, 47
  - chan, 48
  - chan\_context, 48
  - in\_buffer, 48
  - msg\_len, 48
  - sin\_from, 48
  - sin\_to, 48
  - timestamp, 48
- dbl\_rcvfrom
  - API Reference, 24
- dbl\_rcvmode
  - API Reference, 14
- dbl\_send
  - API Reference, 25
- dbl\_send\_connect
  - API Reference, 25
- dbl\_send\_disconnect
  - API Reference, 26
- dbl\_sendto
  - API Reference, 26
- dbl\_set\_filter
  - Flags used for dbl\_bind(), 33
- dbl\_set\_filter\_mode
  - API Reference, 27
- dbl\_shutdown
  - API Reference, 27
- dbl\_ticks\_, 48
- dbl\_timespec, 49
- dbl\_unbind
  - API Reference, 27
- Extensions, 36
  - dbl\_ext\_accept, 37
  - dbl\_ext\_channel\_type, 37
  - dbl\_ext\_getchopt, 37
  - dbl\_ext\_listen, 38
  - dbl\_ext\_poll, 38
  - dbl\_ext\_rcv, 38
  - dbl\_ext\_rcvmsg, 39
  - dbl\_ext\_send, 39
  - dbl\_ext\_setchopt, 40
- Flags for dbl\_send(), 35
  - DBL\_NONBLOCK, 35
  - MSG\_WARM, 35
- Flags used for dbl\_bind(), 30
  - dbl\_eventq\_close, 31
  - dbl\_eventq\_consume, 31
  - dbl\_eventq\_count, 31
  - dbl\_eventq\_inspect, 32
  - dbl\_eventq\_open, 32
  - dbl\_eventq\_peek\_head, 32
  - dbl\_eventq\_peek\_next, 33
  - dbl\_raw\_send, 33
  - dbl\_set\_filter, 33
- Flags used for dbl\_open(), 29
  - DBL\_OPEN\_DISABLED, 29
- hw\_timestamping
  - dbl\_device\_attrs, 46
- in\_buffer
  - dbl\_rcv\_info, 48
- MSG\_WARM
  - Flags for dbl\_send(), 35
- msg\_len
  - dbl\_ext\_rcv\_info, 47
  - dbl\_rcv\_info, 48
- rcvq\_filter\_mode
  - dbl\_device\_attrs, 46
- rcvq\_size
  - dbl\_device\_attrs, 46
- SO\_TIMESTAMPING
  - Specific Options for dbl\_ext\_setchopt(), 41
- sin\_from
  - dbl\_ext\_rcv\_info, 47
  - dbl\_rcv\_info, 48
- sin\_to
  - dbl\_ext\_rcv\_info, 47
  - dbl\_rcv\_info, 48
- Specific Options for dbl\_ext\_setchopt(), 41
  - SO\_TIMESTAMPING, 41
- timestamp
  - dbl\_ext\_rcv\_info, 47
  - dbl\_rcv\_info, 48
- unit\_protocol\_info, 49