

Xenomai RTDM skin API  
2.6.5

Generated by Doxygen 1.8.5

Thu Jul 7 2016 13:25:19



# Contents

<b>1</b>	<b>Deprecated List</b>	<b>1</b>
<b>2</b>	<b>Module Index</b>	<b>3</b>
2.1	Modules . . . . .	3
<b>3</b>	<b>Data Structure Index</b>	<b>5</b>
3.1	Data Structures . . . . .	5
<b>4</b>	<b>File Index</b>	<b>7</b>
4.1	File List . . . . .	7
<b>5</b>	<b>Module Documentation</b>	<b>9</b>
5.1	CAN Devices . . . . .	9
5.1.1	Detailed Description . . . . .	15
5.1.2	Macro Definition Documentation . . . . .	17
5.1.2.1	CAN_CTRLMODE_3_SAMPLES . . . . .	17
5.1.2.2	CAN_CTRLMODE_LISTENONLY . . . . .	18
5.1.2.3	CAN_CTRLMODE_LOOPBACK . . . . .	18
5.1.2.4	CAN_ERR_LOSTARB_UNSPEC . . . . .	18
5.1.2.5	CAN_RAW_ERR_FILTER . . . . .	18
5.1.2.6	CAN_RAW_FILTER . . . . .	19
5.1.2.7	CAN_RAW_LOOPBACK . . . . .	19
5.1.2.8	CAN_RAW_RECV_OWN_MSGS . . . . .	20
5.1.2.9	RTCAN_RTIOC_RCV_TIMEOUT . . . . .	20
5.1.2.10	RTCAN_RTIOC_SND_TIMEOUT . . . . .	20
5.1.2.11	RTCAN_RTIOC_TAKE_TIMESTAMP . . . . .	21
5.1.2.12	SIOCGCANBAUDRATE . . . . .	22
5.1.2.13	SIOCGCANCTRLMODE . . . . .	22
5.1.2.14	SIOCGCANCUSTOMBITTIME . . . . .	23
5.1.2.15	SIOCGCANSTATE . . . . .	23
5.1.2.16	SIOCGIFINDEX . . . . .	24
5.1.2.17	SIOCSCANBAUDRATE . . . . .	24
5.1.2.18	SIOCSCANCTRLMODE . . . . .	25

5.1.2.19	SIOCSCANCUSTOMBITTIME	26
5.1.2.20	SIOCSCANMODE	27
5.1.2.21	SOL_CAN_RAW	27
5.1.3	Typedef Documentation	28
5.1.3.1	can_filter_t	28
5.1.3.2	can_frame_t	28
5.1.4	Enumeration Type Documentation	28
5.1.4.1	CAN_BITTIME_TYPE	28
5.1.4.2	CAN_MODE	28
5.1.4.3	CAN_STATE	29
5.2	Real-Time Driver Model	30
5.2.1	Detailed Description	30
5.2.2	Macro Definition Documentation	30
5.2.2.1	RTDM_TIMEOUT_INFINITE	30
5.2.2.2	RTDM_TIMEOUT_NONE	30
5.2.3	Typedef Documentation	31
5.2.3.1	nanosecs_abs_t	31
5.2.3.2	nanosecs_rel_t	31
5.3	User API	32
5.3.1	Detailed Description	33
5.3.2	Function Documentation	33
5.3.2.1	rt_dev_accept	33
5.3.2.2	rt_dev_bind	33
5.3.2.3	rt_dev_close	34
5.3.2.4	rt_dev_connect	35
5.3.2.5	rt_dev_getpeername	35
5.3.2.6	rt_dev_getsockname	36
5.3.2.7	rt_dev_getsockopt	36
5.3.2.8	rt_dev_ioctl	37
5.3.2.9	rt_dev_listen	37
5.3.2.10	rt_dev_open	38
5.3.2.11	rt_dev_read	39
5.3.2.12	rt_dev_recv	39
5.3.2.13	rt_dev_recvfrom	40
5.3.2.14	rt_dev_recvmsg	40
5.3.2.15	rt_dev_send	41
5.3.2.16	rt_dev_sendmsg	41
5.3.2.17	rt_dev_sendto	42
5.3.2.18	rt_dev_setsockopt	42
5.3.2.19	rt_dev_shutdown	43

5.3.2.20	rt_dev_socket	43
5.3.2.21	rt_dev_write	44
5.4	Real-time IPC protocols	45
5.4.1	Detailed Description	47
5.4.2	Macro Definition Documentation	47
5.4.2.1	BUFP_BUFSZ	47
5.4.2.2	BUFP_LABEL	48
5.4.2.3	IDDP_LABEL	48
5.4.2.4	IDDP_POOLSZ	49
5.4.2.5	SO_RCVTIMEO	50
5.4.2.6	SO_SNDTIMEO	50
5.4.2.7	XDDP_BUFSZ	50
5.4.2.8	XDDP_EVTDOWN	51
5.4.2.9	XDDP_EVTIN	51
5.4.2.10	XDDP_EVTNOBUF	51
5.4.2.11	XDDP_EVTOUT	51
5.4.2.12	XDDP_LABEL	52
5.4.2.13	XDDP_MONITOR	52
5.4.2.14	XDDP_POOLSZ	53
5.4.3	Enumeration Type Documentation	54
5.4.3.1	anonymous enum	54
5.4.4	Function Documentation	54
5.4.4.1	bind__AF_RTIPC	54
5.4.4.2	close__AF_RTIPC	55
5.4.4.3	connect__AF_RTIPC	56
5.4.4.4	getpeername__AF_RTIPC	57
5.4.4.5	getsockname__AF_RTIPC	57
5.4.4.6	getsockopt__AF_RTIPC	58
5.4.4.7	recvmsg__AF_RTIPC	58
5.4.4.8	sendmsg__AF_RTIPC	58
5.4.4.9	setsockopt__AF_RTIPC	59
5.4.4.10	socket__AF_RTIPC	59
5.5	Serial Devices	61
5.5.1	Detailed Description	65
5.5.2	Macro Definition Documentation	66
5.5.2.1	RTSER_RTIOC_BREAK_CTL	66
5.5.2.2	RTSER_RTIOC_GET_CONFIG	67
5.5.2.3	RTSER_RTIOC_GET_CONTROL	67
5.5.2.4	RTSER_RTIOC_GET_STATUS	68
5.5.2.5	RTSER_RTIOC_SET_CONFIG	68

5.5.2.6	RTSER_RTIOC_SET_CONTROL	69
5.5.2.7	RTSER_RTIOC_WAIT_EVENT	69
5.6	Testing Devices	71
5.6.1	Detailed Description	72
5.7	Inter-Driver API	73
5.7.1	Detailed Description	74
5.7.2	Function Documentation	74
5.7.2.1	rtdm_accept	74
5.7.2.2	rtdm_bind	74
5.7.2.3	rtdm_close	74
5.7.2.4	rtdm_connect	74
5.7.2.5	rtdm_context_get	75
5.7.2.6	rtdm_context_lock	76
5.7.2.7	rtdm_context_put	76
5.7.2.8	rtdm_context_unlock	77
5.7.2.9	rtdm_getpeername	77
5.7.2.10	rtdm_getsockname	78
5.7.2.11	rtdm_getsockopt	78
5.7.2.12	rtdm_ioctl	78
5.7.2.13	rtdm_listen	78
5.7.2.14	rtdm_open	78
5.7.2.15	rtdm_read	78
5.7.2.16	rtdm_recv	79
5.7.2.17	rtdm_recvfrom	79
5.7.2.18	rtdm_recvmsg	79
5.7.2.19	rtdm_select_bind	79
5.7.2.20	rtdm_send	80
5.7.2.21	rtdm_sendmsg	80
5.7.2.22	rtdm_sendto	80
5.7.2.23	rtdm_setsockopt	81
5.7.2.24	rtdm_shutdown	81
5.7.2.25	rtdm_socket	81
5.7.2.26	rtdm_write	81
5.8	Device Registration Services	82
5.8.1	Detailed Description	84
5.8.2	Macro Definition Documentation	84
5.8.2.1	RTDM_CLOSING	84
5.8.2.2	RTDM_CREATED_IN_NRT	84
5.8.2.3	RTDM_DEVICE_TYPE_MASK	84
5.8.2.4	RTDM_EXCLUSIVE	84

5.8.2.5	RTDM_NAMED_DEVICE . . . . .	84
5.8.2.6	RTDM_PROTOCOL_DEVICE . . . . .	84
5.8.3	Typedef Documentation . . . . .	84
5.8.3.1	rtdm_close_handler_t . . . . .	84
5.8.3.2	rtdm_ioctl_handler_t . . . . .	85
5.8.3.3	rtdm_open_handler_t . . . . .	85
5.8.3.4	rtdm_read_handler_t . . . . .	85
5.8.3.5	rtdm_recvmmsg_handler_t . . . . .	86
5.8.3.6	rtdm_select_bind_handler_t . . . . .	86
5.8.3.7	rtdm_sendmsg_handler_t . . . . .	87
5.8.3.8	rtdm_socket_handler_t . . . . .	87
5.8.3.9	rtdm_write_handler_t . . . . .	88
5.8.4	Function Documentation . . . . .	89
5.8.4.1	rtdm_context_to_private . . . . .	89
5.8.4.2	rtdm_dev_register . . . . .	89
5.8.4.3	rtdm_dev_unregister . . . . .	90
5.8.4.4	rtdm_private_to_context . . . . .	90
5.9	Driver Development API . . . . .	91
5.9.1	Detailed Description . . . . .	91
5.10	Clock Services . . . . .	92
5.10.1	Detailed Description . . . . .	92
5.10.2	Function Documentation . . . . .	92
5.10.2.1	rtdm_clock_read . . . . .	92
5.10.2.2	rtdm_clock_read_monotonic . . . . .	92
5.11	Task Services . . . . .	94
5.11.1	Detailed Description . . . . .	95
5.11.2	Typedef Documentation . . . . .	95
5.11.2.1	rtdm_task_proc_t . . . . .	95
5.11.3	Function Documentation . . . . .	95
5.11.3.1	rtdm_task_busy_sleep . . . . .	95
5.11.3.2	rtdm_task_current . . . . .	95
5.11.3.3	rtdm_task_destroy . . . . .	96
5.11.3.4	rtdm_task_init . . . . .	97
5.11.3.5	rtdm_task_join_nrt . . . . .	97
5.11.3.6	rtdm_task_set_period . . . . .	98
5.11.3.7	rtdm_task_set_priority . . . . .	98
5.11.3.8	rtdm_task_sleep . . . . .	99
5.11.3.9	rtdm_task_sleep_abs . . . . .	99
5.11.3.10	rtdm_task_sleep_until . . . . .	100
5.11.3.11	rtdm_task_unblock . . . . .	100

5.11.3.12	rtdm_task_wait_period	100
5.12	Timer Services	102
5.12.1	Detailed Description	102
5.12.2	Typedef Documentation	102
5.12.2.1	rtdm_timer_handler_t	102
5.12.3	Enumeration Type Documentation	102
5.12.3.1	rtdm_timer_mode	102
5.12.4	Function Documentation	103
5.12.4.1	rtdm_timer_destroy	103
5.12.4.2	rtdm_timer_init	103
5.12.4.3	rtdm_timer_start	103
5.12.4.4	rtdm_timer_start_in_handler	104
5.12.4.5	rtdm_timer_stop	104
5.12.4.6	rtdm_timer_stop_in_handler	105
5.13	Synchronisation Services	106
5.13.1	Detailed Description	108
5.13.2	Macro Definition Documentation	108
5.13.2.1	RTDM_EXECUTE_ATOMICALY	108
5.13.2.2	rtdm_lock_get	109
5.13.2.3	rtdm_lock_get_irqsave	109
5.13.2.4	rtdm_lock_init	110
5.13.2.5	rtdm_lock_irqrestore	110
5.13.2.6	rtdm_lock_irqsave	110
5.13.2.7	rtdm_lock_put	111
5.13.2.8	rtdm_lock_put_irqrestore	111
5.13.3	Enumeration Type Documentation	112
5.13.3.1	rtdm_selecttype	112
5.13.4	Function Documentation	112
5.13.4.1	EXPORT_SYMBOL_GPL	112
5.13.4.2	EXPORT_SYMBOL_GPL	113
5.13.4.3	EXPORT_SYMBOL_GPL	113
5.13.4.4	EXPORT_SYMBOL_GPL	113
5.13.4.5	EXPORT_SYMBOL_GPL	114
5.13.4.6	EXPORT_SYMBOL_GPL	114
5.13.4.7	EXPORT_SYMBOL_GPL	114
5.13.4.8	EXPORT_SYMBOL_GPL	115
5.13.4.9	EXPORT_SYMBOL_GPL	115
5.13.4.10	EXPORT_SYMBOL_GPL	115
5.13.4.11	EXPORT_SYMBOL_GPL	116
5.13.4.12	EXPORT_SYMBOL_GPL	116



5.13.4.13	EXPORT_SYMBOL_GPL	116
5.13.4.14	EXPORT_SYMBOL_GPL	117
5.13.4.15	EXPORT_SYMBOL_GPL	117
5.13.4.16	rtdm_event_clear	117
5.13.4.17	rtdm_event_destroy	118
5.13.4.18	rtdm_event_init	118
5.13.4.19	rtdm_event_pulse	118
5.13.4.20	rtdm_event_select_bind	119
5.13.4.21	rtdm_event_signal	119
5.13.4.22	rtdm_event_timedwait	120
5.13.4.23	rtdm_event_wait	121
5.13.4.24	rtdm_mutex_destroy	122
5.13.4.25	rtdm_mutex_init	122
5.13.4.26	rtdm_mutex_lock	123
5.13.4.27	rtdm_mutex_timedlock	123
5.13.4.28	rtdm_mutex_unlock	124
5.13.4.29	rtdm_select_bind	124
5.13.4.30	rtdm_sem_destroy	125
5.13.4.31	rtdm_sem_down	125
5.13.4.32	rtdm_sem_init	126
5.13.4.33	rtdm_sem_select_bind	126
5.13.4.34	rtdm_sem_timeddown	126
5.13.4.35	rtdm_sem_up	127
5.13.4.36	rtdm_toseq_init	128
5.14	Interrupt Management Services	129
5.14.1	Detailed Description	129
5.14.2	Macro Definition Documentation	129
5.14.2.1	rtdm_irq_get_arg	129
5.14.3	Typedef Documentation	130
5.14.3.1	rtdm_irq_handler_t	130
5.14.4	Function Documentation	130
5.14.4.1	rtdm_irq_disable	130
5.14.4.2	rtdm_irq_enable	131
5.14.4.3	rtdm_irq_free	131
5.14.4.4	rtdm_irq_request	132
5.15	Non-Real-Time Signalling Services	133
5.15.1	Detailed Description	133
5.15.2	Typedef Documentation	133
5.15.2.1	rtdm_nrtsig_handler_t	133
5.15.3	Function Documentation	133

5.15.3.1	rtm_nrtsig_destroy	133
5.15.3.2	rtm_nrtsig_init	134
5.15.3.3	rtm_nrtsig_pend	134
5.16	Utility Services	135
5.16.1	Detailed Description	135
5.16.2	Function Documentation	135
5.16.2.1	rtm_copy_from_user	135
5.16.2.2	rtm_copy_to_user	136
5.16.2.3	rtm_free	137
5.16.2.4	rtm_in_rt_context	137
5.16.2.5	rtm_iomap_to_user	137
5.16.2.6	rtm_malloc	138
5.16.2.7	rtm_mmap_to_user	139
5.16.2.8	rtm_munmap	140
5.16.2.9	rtm_printk	140
5.16.2.10	rtm_printk_ratelimited	141
5.16.2.11	rtm_ratelimit	141
5.16.2.12	rtm_read_user_ok	142
5.16.2.13	rtm_rt_capable	142
5.16.2.14	rtm_rw_user_ok	143
5.16.2.15	rtm_safe_copy_from_user	143
5.16.2.16	rtm_safe_copy_to_user	144
5.16.2.17	rtm_strncpy_from_user	145
5.17	Device Profiles	146
5.17.1	Detailed Description	147
5.17.2	Macro Definition Documentation	147
5.17.2.1	RTIOC_DEVICE_INFO	147
5.17.2.2	RTIOC_PURGE	147
<b>6</b>	<b>Data Structure Documentation</b>	<b>149</b>
6.1	can_bittime Struct Reference	149
6.1.1	Detailed Description	149
6.2	can_bittime_btr Struct Reference	149
6.2.1	Detailed Description	150
6.3	can_bittime_std Struct Reference	150
6.3.1	Detailed Description	150
6.4	can_filter Struct Reference	150
6.4.1	Detailed Description	151
6.4.2	Field Documentation	151
6.4.2.1	can_id	151

6.4.2.2	can_mask	151
6.5	can_frame Struct Reference	151
6.5.1	Detailed Description	152
6.5.2	Field Documentation	152
6.5.2.1	can_id	152
6.6	rtdm_dev_context Struct Reference	152
6.6.1	Detailed Description	153
6.7	rtdm_device Struct Reference	153
6.7.1	Detailed Description	154
6.7.2	Field Documentation	154
6.7.2.1	open_rt	154
6.7.2.2	socket_rt	155
6.8	rtdm_device_info Struct Reference	155
6.8.1	Detailed Description	155
6.9	rtdm_operations Struct Reference	155
6.9.1	Detailed Description	156
6.9.2	Field Documentation	156
6.9.2.1	close_rt	156
6.10	rtipc_port_label Struct Reference	156
6.10.1	Detailed Description	157
6.10.2	Field Documentation	157
6.10.2.1	label	157
6.11	rtser_config Struct Reference	157
6.11.1	Detailed Description	158
6.12	rtser_event Struct Reference	158
6.12.1	Detailed Description	158
6.13	rtser_status Struct Reference	158
6.13.1	Detailed Description	159
6.14	sockaddr_can Struct Reference	159
6.14.1	Detailed Description	159
6.14.2	Field Documentation	159
6.14.2.1	can_ifindex	159
6.15	sockaddr_ipc Struct Reference	159
6.15.1	Detailed Description	160
6.15.2	Field Documentation	160
6.15.2.1	sipc_port	160
<b>7</b>	<b>File Documentation</b>	<b>161</b>
7.1	include/rtdm/rtdmcan.h File Reference	161
7.1.1	Detailed Description	166

7.2	include/rtdm/rtdm.h File Reference	167
7.2.1	Detailed Description	168
7.3	include/rtdm/rtdm_driver.h File Reference	169
7.3.1	Detailed Description	174
7.4	include/rtdm/rtipc.h File Reference	174
7.4.1	Detailed Description	176
7.5	include/rtdm/rtserial.h File Reference	176
7.5.1	Detailed Description	180
7.6	include/rtdm/rttesting.h File Reference	180
7.6.1	Detailed Description	181
7.7	ksrc/skins/rtdm/device.c File Reference	181
7.7.1	Detailed Description	182
7.8	ksrc/skins/rtdm/drvlib.c File Reference	182
7.8.1	Detailed Description	186
7.9	ksrc/skins/rtdm/module.c File Reference	186
7.9.1	Detailed Description	186
7.10	ksrc/skins/rtdm/core.c File Reference	186
7.10.1	Detailed Description	189
<b>8</b>	<b>Example Documentation</b>	<b>191</b>
8.1	bufp-label.c	191
8.2	bufp-readwrite.c	193
8.3	cross-link.c	196
8.4	iddp-label.c	200
8.5	iddp-sendrecv.c	202
8.6	rtcan_rtt.c	205
8.7	rtcanconfig.c	210
8.8	rtcanrecv.c	213
8.9	rtcansend.c	217
8.10	xddp-echo.c	221
8.11	xddp-label.c	224
8.12	xddp-stream.c	228

# Chapter 1

## Deprecated List

**Global [rtm\\_device::open\\_rt](#)**

Only use non-real-time open handler in new drivers.

**Global [rtm\\_device::socket\\_rt](#)**

Only use non-real-time socket creation handler in new drivers.

**Global [rtm\\_operations::close\\_rt](#)**

Only use non-real-time close handler in new drivers.

**Global [rtm\\_task\\_sleep\\_until \(nanosecs\\_abs\\_t wakeup\\_time\)](#)**

Use [rtm\\_task\\_sleep\\_abs](#) instead!



# Chapter 2

## Module Index

### 2.1 Modules

Here is a list of all modules:

Real-Time Driver Model . . . . .	30
User API . . . . .	32
Driver Development API . . . . .	91
Inter-Driver API . . . . .	73
Device Registration Services . . . . .	82
Synchronisation Services . . . . .	106
Clock Services . . . . .	92
Task Services . . . . .	94
Timer Services . . . . .	102
Synchronisation Services . . . . .	106
Interrupt Management Services . . . . .	129
Non-Real-Time Signalling Services . . . . .	133
Utility Services . . . . .	135
Device Profiles . . . . .	146
CAN Devices . . . . .	9
Real-time IPC protocols . . . . .	45
Serial Devices . . . . .	61
Testing Devices . . . . .	71





# Chapter 3

## Data Structure Index

### 3.1 Data Structures

Here are the data structures with brief descriptions:

<a href="#">can_bittime</a>	Custom CAN bit-time definition . . . . .	149
<a href="#">can_bittime_btr</a>	Hardware-specific BTR bit-times . . . . .	149
<a href="#">can_bittime_std</a>	Standard bit-time parameters according to Bosch . . . . .	150
<a href="#">can_filter</a>	Filter for reception of CAN messages . . . . .	150
<a href="#">can_frame</a>	Raw CAN frame . . . . .	151
<a href="#">rtdm_dev_context</a>	Device context . . . . .	152
<a href="#">rtdm_device</a>	RTDM device . . . . .	153
<a href="#">rtdm_device_info</a>	Device information . . . . .	155
<a href="#">rtdm_operations</a>	Device operations . . . . .	155
<a href="#">rtipc_port_label</a>	Port label information structure . . . . .	156
<a href="#">rtser_config</a>	Serial device configuration . . . . .	157
<a href="#">rtser_event</a>	Additional information about serial device events . . . . .	158
<a href="#">rtser_status</a>	Serial device status . . . . .	158
<a href="#">sockaddr_can</a>	Socket address structure for the CAN address family . . . . .	159
<a href="#">sockaddr_ipc</a>	Socket address structure for the RTIPC address family . . . . .	159



# Chapter 4

## File Index

### 4.1 File List

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

<a href="#">include/rtdm/<b>rtcan.h</b></a>	Real-Time Driver Model for RT-Socket-CAN, CAN device profile header . . . . .	161
<a href="#">include/rtdm/<b>rtdm.h</b></a>	Real-Time Driver Model for Xenomai, user API header . . . . .	167
<a href="#">include/rtdm/<b>rtdm_driver.h</b></a>	Real-Time Driver Model for Xenomai, driver API header . . . . .	169
<a href="#">include/rtdm/<b>rtipc.h</b></a>	This file is part of the Xenomai project . . . . .	174
<a href="#">include/rtdm/<b>rtserial.h</b></a>	Real-Time Driver Model for Xenomai, serial device profile header . . . . .	176
<a href="#">include/rtdm/<b>rttesting.h</b></a>	Real-Time Driver Model for Xenomai, testing device profile header . . . . .	180
<a href="#">include/rtdm/<b>syscall.h</b></a>	. . . . .	??
<a href="#">ksrc/skins/rtdm/<b>core.c</b></a>	Real-Time Driver Model for Xenomai, device operation multiplexing . . . . .	186
<a href="#">ksrc/skins/rtdm/<b>device.c</b></a>	Real-Time Driver Model for Xenomai, device management . . . . .	181
<a href="#">ksrc/skins/rtdm/<b>drvlib.c</b></a>	Real-Time Driver Model for Xenomai, driver library . . . . .	182
<a href="#">ksrc/skins/rtdm/<b>internal.h</b></a>	. . . . .	??
<a href="#">ksrc/skins/rtdm/<b>module.c</b></a>	Real-Time Driver Model for Xenomai . . . . .	186



# Chapter 5

## Module Documentation

### 5.1 CAN Devices

#### Files

- file [rtcan.h](#)  
*Real-Time Driver Model for RT-Socket-CAN, CAN device profile header.*

#### Data Structures

- struct [can\\_bittime\\_std](#)  
*Standard bit-time parameters according to Bosch.*
- struct [can\\_bittime\\_btr](#)  
*Hardware-specific BTR bit-times.*
- struct [can\\_bittime](#)  
*Custom CAN bit-time definition.*
- struct [can\\_filter](#)  
*Filter for reception of CAN messages.*
- struct [sockaddr\\_can](#)  
*Socket address structure for the CAN address family.*
- struct [can\\_frame](#)  
*Raw CAN frame.*

#### Macros

- `#define AF_CAN 29`  
*CAN address family.*
- `#define PF_CAN AF_CAN`  
*CAN protocol family.*
- `#define SOL_CAN_RAW 103`  
*CAN socket levels.*

#### Typedefs

- typedef uint32\_t [can\\_id\\_t](#)  
*Type of CAN id (see [CAN\\_xxx\\_MASK](#) and [CAN\\_xxx\\_FLAG](#))*

- typedef `can_id_t` `can_err_mask_t`  
*Type of CAN error mask.*
- typedef `uint32_t` `can_baudrate_t`  
*Baudrate definition in bits per second.*
- typedef enum `CAN_BITTIME_TYPE` `can_bittime_type_t`  
*See `CAN_BITTIME_TYPE`.*
- typedef enum `CAN_MODE` `can_mode_t`  
*See `CAN_MODE`.*
- typedef `int` `can_ctrlmode_t`  
*See `CAN_CTRLMODE`.*
- typedef enum `CAN_STATE` `can_state_t`  
*See `CAN_STATE`.*
- typedef struct `can_filter` `can_filter_t`  
*Filter for reception of CAN messages.*
- typedef struct `can_frame` `can_frame_t`  
*Raw CAN frame.*

## Enumerations

- enum `CAN_BITTIME_TYPE` { `CAN_BITTIME_STD`, `CAN_BITTIME_BTR` }  
*Supported CAN bit-time types.*

## CAN ID masks

### Bit masks for masking CAN IDs

- #define `CAN_EFF_MASK` `0x1FFFFFFF`  
*Bit mask for extended CAN IDs.*
- #define `CAN_SFF_MASK` `0x000007FF`  
*Bit mask for standard CAN IDs.*

## CAN ID flags

### Flags within a CAN ID indicating special CAN frame attributes

- #define `CAN_EFF_FLAG` `0x80000000`  
*Extended frame.*
- #define `CAN_RTR_FLAG` `0x40000000`  
*Remote transmission frame.*
- #define `CAN_ERR_FLAG` `0x20000000`  
*Error frame (see [Errors](#)), not valid in struct `can_filter`.*
- #define `CAN_INV_FILTER` `CAN_ERR_FLAG`  
*Invert CAN filter definition, only valid in struct `can_filter`.*

## Particular CAN protocols

### Possible protocols for the PF\_CAN protocol family

Currently only the RAW protocol is supported.

- #define `CAN_RAW` `1`  
*Raw protocol of PF\_CAN, applicable to socket type `SOCK_RAW`.*

## CAN operation modes

Modes into which CAN controllers can be set

- enum `CAN_MODE` { `CAN_MODE_STOP` = 0, `CAN_MODE_START`, `CAN_MODE_SLEEP` }

## CAN controller modes

Special CAN controllers modes, which can be or'ed together.

Note

These modes are hardware-dependent. Please consult the hardware manual of the CAN controller for more detailed information.

- #define `CAN_CTRLMODE_LISTENONLY` 0x1
- #define `CAN_CTRLMODE_LOOPBACK` 0x2
- #define `CAN_CTRLMODE_3_SAMPLES` 0x4

## CAN controller states

States a CAN controller can be in.

- enum `CAN_STATE` {  
    `CAN_STATE_ERROR_ACTIVE` = 0, `CAN_STATE_ERROR_WARNING` = 1, `CAN_STATE_ERROR_PASSIVE`  
    = 2, `CAN_STATE_BUS_OFF`,  
    `CAN_STATE_SCANNING_BAUDRATE`, `CAN_STATE_STOPPED`, `CAN_STATE_SLEEPING` }

## Timestamp switches

Arguments to pass to `RTCAN_RTIOC_TAKE_TIMESTAMP`

- #define `RTCAN_TAKE_NO_TIMESTAMPS` 0  
    *Switch off taking timestamps.*
- #define `RTCAN_TAKE_TIMESTAMPS` 1  
    *Do take timestamps.*

## RAW socket options

Setting and getting CAN RAW socket options.

- #define `CAN_RAW_FILTER` 0x1  
    *CAN filter definition.*
- #define `CAN_RAW_ERR_FILTER` 0x2  
    *CAN error mask.*
- #define `CAN_RAW_LOOPBACK` 0x3  
    *CAN TX loopback.*
- #define `CAN_RAW_RECV_OWN_MSGS` 0x4  
    *CAN receive own messages.*

## IOCTLs

### CAN device IOCTLs

- #define [SIOCGIFINDEX](#) defined\_by\_kernel\_header\_file  
*Get CAN interface index by name.*
- #define [SIOCSCANBAUDRATE](#) \_IOW(RTIOC\_TYPE\_CAN, 0x01, struct ifreq)  
*Set baud rate.*
- #define [SIOCGCANBAUDRATE](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x02, struct ifreq)  
*Get baud rate.*
- #define [SIOCSCANCUSTOMBITTIME](#) \_IOW(RTIOC\_TYPE\_CAN, 0x03, struct ifreq)  
*Set custom bit time parameter.*
- #define [SIOCGCANCUSTOMBITTIME](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x04, struct ifreq)  
*Get custom bit-time parameters.*
- #define [SIOCSCANMODE](#) \_IOW(RTIOC\_TYPE\_CAN, 0x05, struct ifreq)  
*Set operation mode of CAN controller.*
- #define [SIOCGCANSTATE](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x06, struct ifreq)  
*Get current state of CAN controller.*
- #define [SIOCSCANCTRLMODE](#) \_IOW(RTIOC\_TYPE\_CAN, 0x07, struct ifreq)  
*Set special controller modes.*
- #define [SIOCGCANCTRLMODE](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x08, struct ifreq)  
*Get special controller modes.*
- #define [RTCAN\\_RTIOC\\_TAKE\\_TIMESTAMP](#) \_IOW(RTIOC\_TYPE\_CAN, 0x09, int)  
*Enable or disable storing a high precision timestamp upon reception of a CAN frame.*
- #define [RTCAN\\_RTIOC\\_RCV\\_TIMEOUT](#) \_IOW(RTIOC\_TYPE\_CAN, 0x0A, nanosecs\_rel\_t)  
*Specify a reception timeout for a socket.*
- #define [RTCAN\\_RTIOC\\_SND\\_TIMEOUT](#) \_IOW(RTIOC\_TYPE\_CAN, 0x0B, nanosecs\_rel\_t)  
*Specify a transmission timeout for a socket.*

### Error mask

Error class (mask) in `can_id` field of struct [can\\_frame](#) to be used with [CAN\\_RAW\\_ERR\\_FILTER](#).

**Note:** Error reporting is hardware dependent and most CAN controllers report less detailed error conditions than the SJA1000.

**Note:** In case of a bus-off error condition ([CAN\\_ERR\\_BUSOFF](#)), the CAN controller is **not** restarted automatically. It is the application's responsibility to react appropriately, e.g. calling [CAN\\_MODE\\_START](#).

**Note:** Bus error interrupts ([CAN\\_ERR\\_BUSERROR](#)) are enabled when an application is calling a [Recv](#) function on a socket listening on bus errors (using [CAN\\_RAW\\_ERR\\_FILTER](#)). After one bus error has occurred, the interrupt will be disabled to allow the application time for error processing and to efficiently avoid bus error interrupt flooding.

- #define [CAN\\_ERR\\_TX\\_TIMEOUT](#) 0x00000001U  
*TX timeout (netdevice driver)*
- #define [CAN\\_ERR\\_LOSTARB](#) 0x00000002U  
*Lost arbitration (see [data\[0\]](#))*
- #define [CAN\\_ERR\\_CRTL](#) 0x00000004U  
*Controller problems (see [data\[1\]](#))*
- #define [CAN\\_ERR\\_PROT](#) 0x00000008U  
*Protocol violations (see [data\[2\]](#), [data\[3\]](#))*
- #define [CAN\\_ERR\\_TRX](#) 0x00000010U



Transceiver status (see [data\[4\]](#))

- #define [CAN\\_ERR\\_ACK](#) 0x00000020U  
*Received no ACK on transmission.*
- #define [CAN\\_ERR\\_BUSOFF](#) 0x00000040U  
*Bus off.*
- #define [CAN\\_ERR\\_BUSERROR](#) 0x00000080U  
*Bus error (may flood!)*
- #define [CAN\\_ERR\\_RESTARTED](#) 0x00000100U  
*Controller restarted.*
- #define [CAN\\_ERR\\_MASK](#) 0x1FFFFFFFU  
*Omit EFF, RTR, ERR flags.*

## Arbitration lost error

Error in the `data[0]` field of struct `can_frame`.

- #define [CAN\\_ERR\\_LOSTARB\\_UNSPEC](#) 0x00  
*unspecified*

## Controller problems

Error in the `data[1]` field of struct `can_frame`.

- #define [CAN\\_ERR\\_CRTL\\_UNSPEC](#) 0x00  
*unspecified*
- #define [CAN\\_ERR\\_CRTL\\_RX\\_OVERFLOW](#) 0x01  
*RX buffer overflow.*
- #define [CAN\\_ERR\\_CRTL\\_TX\\_OVERFLOW](#) 0x02  
*TX buffer overflow.*
- #define [CAN\\_ERR\\_CRTL\\_RX\\_WARNING](#) 0x04  
*reached warning level for RX errors*
- #define [CAN\\_ERR\\_CRTL\\_TX\\_WARNING](#) 0x08  
*reached warning level for TX errors*
- #define [CAN\\_ERR\\_CRTL\\_RX\\_PASSIVE](#) 0x10  
*reached passive level for RX errors*
- #define [CAN\\_ERR\\_CRTL\\_TX\\_PASSIVE](#) 0x20  
*reached passive level for TX errors*

## Protocol error type

Error in the `data[2]` field of struct `can_frame`.

- #define [CAN\\_ERR\\_PROT\\_UNSPEC](#) 0x00  
*unspecified*
- #define [CAN\\_ERR\\_PROT\\_BIT](#) 0x01  
*single bit error*
- #define [CAN\\_ERR\\_PROT\\_FORM](#) 0x02  
*frame format error*
- #define [CAN\\_ERR\\_PROT\\_STUFF](#) 0x04

- *bit stuffing error*
- #define CAN\_ERR\_PROT\_BIT0 0x08  
*unable to send dominant bit*
- #define CAN\_ERR\_PROT\_BIT1 0x10  
*unable to send recessive bit*
- #define CAN\_ERR\_PROT\_OVERLOAD 0x20  
*bus overload*
- #define CAN\_ERR\_PROT\_ACTIVE 0x40  
*active error announcement*
- #define CAN\_ERR\_PROT\_TX 0x80  
*error occurred on transmission*

## Protocol error location

Error in the data[4] field of struct `can_frame`.

- #define CAN\_ERR\_PROT\_LOC\_UNSPEC 0x00  
*unspecified*
- #define CAN\_ERR\_PROT\_LOC\_SOF 0x03  
*start of frame*
- #define CAN\_ERR\_PROT\_LOC\_ID28\_21 0x02  
*ID bits 28 - 21 (SFF: 10 - 3)*
- #define CAN\_ERR\_PROT\_LOC\_ID20\_18 0x06  
*ID bits 20 - 18 (SFF: 2 - 0)*
- #define CAN\_ERR\_PROT\_LOC\_SRTR 0x04  
*substitute RTR (SFF: RTR)*
- #define CAN\_ERR\_PROT\_LOC\_IDE 0x05  
*identifier extension*
- #define CAN\_ERR\_PROT\_LOC\_ID17\_13 0x07  
*ID bits 17-13.*
- #define CAN\_ERR\_PROT\_LOC\_ID12\_05 0x0F  
*ID bits 12-5.*
- #define CAN\_ERR\_PROT\_LOC\_ID04\_00 0x0E  
*ID bits 4-0.*
- #define CAN\_ERR\_PROT\_LOC\_RTR 0x0C  
*RTR.*
- #define CAN\_ERR\_PROT\_LOC\_RES1 0x0D  
*reserved bit 1*
- #define CAN\_ERR\_PROT\_LOC\_RES0 0x09  
*reserved bit 0*
- #define CAN\_ERR\_PROT\_LOC\_DLC 0x0B  
*data length code*
- #define CAN\_ERR\_PROT\_LOC\_DATA 0x0A  
*data section*
- #define CAN\_ERR\_PROT\_LOC\_CRC\_SEQ 0x08  
*CRC sequence.*
- #define CAN\_ERR\_PROT\_LOC\_CRC\_DEL 0x18  
*CRC delimiter.*
- #define CAN\_ERR\_PROT\_LOC\_ACK 0x19  
*ACK slot.*

- #define `CAN_ERR_PROT_LOC_ACK_DEL` 0x1B  
*ACK delimiter.*
- #define `CAN_ERR_PROT_LOC_EOF` 0x1A  
*end of frame*
- #define `CAN_ERR_PROT_LOC_INTERM` 0x12  
*intermission*
- #define `CAN_ERR_TRX_UNSPEC` 0x00  
*0000 0000*
- #define `CAN_ERR_TRX_CANH_NO_WIRE` 0x04  
*0000 0100*
- #define `CAN_ERR_TRX_CANH_SHORT_TO_BAT` 0x05  
*0000 0101*
- #define `CAN_ERR_TRX_CANH_SHORT_TO_VCC` 0x06  
*0000 0110*
- #define `CAN_ERR_TRX_CANH_SHORT_TO_GND` 0x07  
*0000 0111*
- #define `CAN_ERR_TRX_CANL_NO_WIRE` 0x40  
*0100 0000*
- #define `CAN_ERR_TRX_CANL_SHORT_TO_BAT` 0x50  
*0101 0000*
- #define `CAN_ERR_TRX_CANL_SHORT_TO_VCC` 0x60  
*0110 0000*
- #define `CAN_ERR_TRX_CANL_SHORT_TO_GND` 0x70  
*0111 0000*
- #define `CAN_ERR_TRX_CANL_SHORT_TO_CANH` 0x80  
*1000 0000*

### 5.1.1 Detailed Description

This is the common interface a RTDM-compliant CAN device has to provide. Feel free to report bugs and comments on this profile to the "Socketcan" mailing list ([Socketcan-core@lists.berlios.de](mailto:Socketcan-core@lists.berlios.de)) or directly to the authors ([wg@grandegger.com](mailto:wg@grandegger.com) or [Sebastian.Smolorz@stud.uni-hannover.de](mailto:Sebastian.Smolorz@stud.uni-hannover.de)).

**Profile Revision:** 2

#### Device Characteristics

**Device Flags:** RTDM\_PROTOCOL\_DEVICE  
**Protocol Family:** PF\_CAN  
**Socket Type:** SOCK\_RAW  
**Device Class:** RTDM\_CLASS\_CAN

#### Supported Operations

##### Socket

Environments: non-RT (RT optional, deprecated)

Specific return values:

- -EPROTONOSUPPORT (Protocol is not supported by the driver. See [CAN protocols](#) for possible protocols.)

##### Close

Blocking calls to any of the [Send](#) or [Receive](#) functions will be unblocked when the socket is closed and return with an error.

Environments: non-RT (RT optional, deprecated)

Specific return values: none

## IOCTL

Mandatory Environments: see [below](#)

Specific return values: see [below](#)

### Bind

Binds a socket to one or all CAN devices (see struct [sockaddr\\_can](#)). If a filter list has been defined with `setsockopt` (see [Socketopts](#)), it will be used upon reception of CAN frames to decide whether the bound socket will receive a frame. If no filter has been defined, the socket will receive **all** CAN frames on the specified interface(s).

Binding to special interface index 0 will make the socket receive CAN frames from all CAN interfaces. Binding to an interface index is also relevant for the [Send](#) functions because they will transmit a message over the interface the socket is bound to when no socket address is given to them.

Environments: non-RT (RT optional)

Specific return values:

- -EFAULT (It was not possible to access user space memory area at the specified address.)
- -ENOMEM (Not enough memory to fulfill the operation)
- -EINVAL (Invalid address family, or invalid length of address structure)
- -ENODEV (Invalid CAN interface index)
- -ENOSPC (No enough space for filter list)
- -EBADF (Socket is about to be closed)
- -EAGAIN (Too many receivers. Old binding (if any) is still active. Close some sockets and try again.)

### Setsockopt, Getsockopt

These functions allow to set and get various socket options. Currently, only CAN raw sockets are supported.

Supported Levels and Options:

- Level **SOL\_CAN\_RAW** : CAN RAW protocol (see [CAN\\_RAW](#))
  - Option [CAN\\_RAW\\_FILTER](#) : CAN filter list
  - Option [CAN\\_RAW\\_ERR\\_FILTER](#) : CAN error mask
  - Option [CAN\\_RAW\\_LOOPBACK](#) : CAN TX loopback to local sockets

Environments: non-RT (RT optional)

Specific return values: see links to options above.

### Recv, Recvfrom, Recvmsg

These functions receive CAN messages from a socket. Only one message per call can be received, so only one buffer with the correct length must be passed. For `SOCK_RAW`, this is the size of struct [can\\_frame](#).

Unlike a call to one of the [Send](#) functions, a Recv function will not return with an error if an interface is down (due to bus-off or setting of stop mode) or in sleep mode. Moreover, in such a case there may still be some CAN messages in the socket buffer which could be read out successfully.

It is possible to receive a high precision timestamp with every CAN message. The condition is a former instruction to the socket via [RTCAN\\_RTIOC\\_TAKE\\_TIMESTAMP](#). The timestamp will be copied to the `msg_control` buffer of struct `user_msghdr` if it points to a valid memory location with size of [nanosecs\\_abs\\_t](#). If this is a NULL pointer the timestamp will be discarded silently.

**Note:** A `msg_controllen` of 0 upon completion of the function call indicates that no timestamp is available for that message.

Supported Flags [in]:

- `MSG_DONTWAIT` (By setting this flag the operation will only succeed if it would not block, i.e. if there is a message in the socket buffer. This flag takes precedence over a timeout specified by [RTCAN\\_RTIOC\\_RCV\\_TIMEOUT](#).)
- `MSG_PEEK` (Receive a message but leave it in the socket buffer. The next receive operation will get that message again.)

Supported Flags [out]: none

Environments: RT (non-RT optional)

Specific return values:

- Non-negative value (Indicating the successful reception of a CAN message. For `SOCK_RAW`, this is the size of struct `can_frame` regardless of the actual size of the payload.)
- `-EFAULT` (It was not possible to access user space memory area at one of the specified addresses.)
- `-EINVAL` (Unsupported flag detected, or invalid length of socket address buffer, or invalid length of message control buffer)
- `-EMSGSIZE` (Zero or more than one `iovec` buffer passed, or buffer too small)
- `-EAGAIN` (No data available in non-blocking mode)
- `-EBADF` (Socket was closed.)
- `-EINTR` (Operation was interrupted explicitly or by signal.)
- `-ETIMEDOUT` (Timeout)

### Send, Sendto, Sendmsg

These functions send out CAN messages. Only one message per call can be transmitted, so only one buffer with the correct length must be passed. For `SOCK_RAW`, this is the size of struct `can_frame`. The following only applies to `SOCK_RAW`: If a socket address of struct `sockaddr_can` is given, only `can_ifindex` is used. It is also possible to omit the socket address. Then the interface the socket is bound to will be used for sending messages.

If an interface goes down (due to bus-off or setting of stop mode) all senders that were blocked on this interface will be woken up.

Supported Flags:

- `MSG_DONTWAIT` (By setting this flag the transmit operation will only succeed if it would not block. This flag takes precedence over a timeout specified by `RTCAN_RTIOC_SND_TIMEOUT`.)

Environments: RT (non-RT optional)

Specific return values:

- Non-negative value equal to given buffer size (Indicating the successful completion of the function call. See also note.)
- `-EOPNOTSUPP` (`MSG_OOB` flag is not supported.)
- `-EINVAL` (Unsupported flag detected *or*: Invalid length of socket address *or*: Invalid address family *or*: Data length code of CAN frame not between 0 and 15 *or*: CAN standard frame has got an ID not between 0 and 2031)
- `-EMSGSIZE` (Zero or more than one buffer passed or invalid size of buffer)
- `-EFAULT` (It was not possible to access user space memory area at one of the specified addresses.)
- `-ENXIO` (Invalid CAN interface index - 0 is not allowed here - or socket not bound or rather bound to all interfaces.)
- `-ENETDOWN` (Controller is bus-off or in stopped state.)
- `-ECOMM` (Controller is sleeping)
- `-EAGAIN` (Cannot transmit without blocking but a non-blocking call was requested.)
- `-EINTR` (Operation was interrupted explicitly or by signal)
- `-EBADF` (Socket was closed.)
- `-ETIMEDOUT` (Timeout)

**Note:** A successful completion of the function call does not implicate a successful transmission of the message.

## 5.1.2 Macro Definition Documentation

### 5.1.2.1 `#define CAN_CTRLMODE_3_SAMPLES 0x4`

Triple sampling mode

In this mode the CAN controller uses Triple sampling.

### 5.1.2.2 #define CAN\_CTRLMODE\_LISTENONLY 0x1

#### Listen-Only mode

In this mode the CAN controller would give no acknowledge to the CAN-bus, even if a message is received successfully and messages would not be transmitted. This mode might be useful for bus-monitoring, hot-plugging or throughput analysis.

Examples:

[rtcanconfig.c](#).

### 5.1.2.3 #define CAN\_CTRLMODE\_LOOPBACK 0x2

#### Loopback mode

In this mode the CAN controller does an internal loop-back, a message is transmitted and simultaneously received. That mode can be used for self test operation.

Examples:

[rtcanconfig.c](#).

### 5.1.2.4 #define CAN\_ERR\_LOSTARB\_UNSPEC 0x00

unspecified

else bit number in bitstream

### 5.1.2.5 #define CAN\_RAW\_ERR\_FILTER 0x2

#### CAN error mask.

A CAN error mask (see [Errors](#)) can be set with `setsockopt`. This mask is then used to decide if error frames are delivered to this socket in case of error conditions. The error frames are marked with the `CAN_ERR_FLAG` of `CAN_xxx_FLAG` and must be handled by the application properly. A detailed description of the errors can be found in the `can_id` and the `data` fields of struct `can_frame` (see [Errors](#) for further details).

Parameters

<code>in</code>	<code>level</code>	<b>SOL_CAN_RAW</b>
<code>in</code>	<code>optname</code>	<b>CAN_RAW_ERR_FILTER</b>
<code>in</code>	<code>optval</code>	Pointer to error mask of type <code>can_err_mask_t</code> .
<code>in</code>	<code>optlen</code>	Size of error mask: <code>sizeof(can_err_mask_t)</code> .

Environments: non-RT (RT optional)

Specific return values:

- -EFAULT (It was not possible to access user space memory area at the specified address.)
- -EINVAL (Invalid length "optlen")

Examples:

[rtcanrecv.c](#).

### 5.1.2.6 #define CAN\_RAW\_FILTER 0x1

CAN filter definition.

A CAN raw filter list with elements of struct [can\\_filter](#) can be installed with `setsockopt`. This list is used upon reception of CAN frames to decide whether the bound socket will receive a frame. An empty filter list can also be defined using `optlen = 0`, which is recommended for write-only sockets.

If the socket was already bound with [Bind](#), the old filter list gets replaced with the new one. Be aware that already received, but not read out CAN frames may stay in the socket buffer.

Parameters

<i>in</i>	<i>level</i>	<b>SOL_CAN_RAW</b>
<i>in</i>	<i>optname</i>	<b>CAN_RAW_FILTER</b>
<i>in</i>	<i>optval</i>	Pointer to array of struct <a href="#">can_filter</a> .
<i>in</i>	<i>optlen</i>	Size of filter list: <code>count * sizeof( struct can_filter)</code> . Environments: non-RT (RT optional) Specific return values: <ul style="list-style-type: none"><li>• -EFAULT (It was not possible to access user space memory area at the specified address.)</li><li>• -ENOMEM (Not enough memory to fulfill the operation)</li><li>• -EINVAL (Invalid length "optlen")</li><li>• -ENOSPC (No space to store filter list, check RT-Socket-CAN kernel parameters)</li></ul>

Examples:

[rtcan\\_rtt.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

### 5.1.2.7 #define CAN\_RAW\_LOOPBACK 0x3

CAN TX loopback.

The TX loopback to other local sockets can be selected with this `setsockopt`.

Note

The TX loopback feature must be enabled in the kernel and then the loopback to other local TX sockets is enabled by default.

Parameters

<i>in</i>	<i>level</i>	<b>SOL_CAN_RAW</b>
<i>in</i>	<i>optname</i>	<b>CAN_RAW_LOOPBACK</b>
<i>in</i>	<i>optval</i>	Pointer to integer value.
<i>in</i>	<i>optlen</i>	Size of int: <code>sizeof(int)</code> .

Environments: non-RT (RT optional)

Specific return values:

- -EFAULT (It was not possible to access user space memory area at the specified address.)
- -EINVAL (Invalid length "optlen")
- -EOPNOTSUPP (not supported, check RT-Socket-CAN kernel parameters).

Examples:

[rtcansend.c](#).

5.1.2.8 `#define CAN_RAW_RECV_OWN_MSGS 0x4`

CAN receive own messages.

Not supported by RT-Socket-CAN, but defined for compatibility with Socket-CAN.

5.1.2.9 `#define RTCAN_RTIOC_RCV_TIMEOUT _IOW(RTIOC_TYPE_CAN, 0x0A, nanosecs_rel_t)`

Specify a reception timeout for a socket.

Defines a timeout for all receive operations via a socket which will take effect when one of the [receive functions](#) is called without the `MSG_DONTWAIT` flag set.

The default value for a newly created socket is an infinite timeout.

Note

The setting of the timeout value is not done atomically to avoid locks. Please set the value before receiving messages from the socket.

Parameters

<code>in</code>	<code>arg</code>	Pointer to <a href="#">nanosecs_rel_t</a> variable. The value is interpreted as relative timeout in nanoseconds in case of a positive value. See <a href="#">Timeouts</a> for special timeouts.
-----------------	------------------	---

Returns

0 on success, otherwise:

- `-EFAULT`: It was not possible to access user space memory area at the specified address.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

Examples:

[rtcanrecv.c](#).

5.1.2.10 `#define RTCAN_RTIOC_SND_TIMEOUT _IOW(RTIOC_TYPE_CAN, 0x0B, nanosecs_rel_t)`

Specify a transmission timeout for a socket.

Defines a timeout for all send operations via a socket which will take effect when one of the [send functions](#) is called without the `MSG_DONTWAIT` flag set.

The default value for a newly created socket is an infinite timeout.



## Note

The setting of the timeout value is not done atomically to avoid locks. Please set the value before sending messages to the socket.

## Parameters

<code>in</code>	<code>arg</code>	Pointer to <a href="#">nanosecs_rel_t</a> variable. The value is interpreted as relative timeout in nanoseconds in case of a positive value. See <a href="#">Timeouts</a> for special timeouts.
-----------------	------------------	---

## Returns

0 on success, otherwise:

- `-EFAULT`: It was not possible to access user space memory area at the specified address.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

## Examples:

[rtcansend.c](#).

5.1.2.11 `#define RTCAN_RTIOC_TAKE_TIMESTAMP _IOW(RTIOC_TYPE_CAN, 0x09, int)`

Enable or disable storing a high precision timestamp upon reception of a CAN frame.

A newly created socket takes no timestamps by default.

## Parameters

<code>in</code>	<code>arg</code>	int variable, see <a href="#">Timestamp switches</a>
-----------------	------------------	--

## Returns

0 on success.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

## Note

Activating taking timestamps only has an effect on newly received CAN messages from the bus. Frames that already are in the socket buffer do not have timestamps if it was deactivated before. See [Receive](#) for more details.

Rescheduling: never.

Examples:

[rtcanrecv.c](#).

5.1.2.12 `#define SIOCGCANBAUDRATE _IOWR(RTIOC_TYPE_CAN, 0x02, struct ifreq)`

Get baud rate.

Parameters

<code>in,out</code>	<code>arg</code>	Pointer to interface request structure buffer (struct ifreq from linux/if.h). <code>ifr_name</code> must hold a valid CAN interface name, <code>ifr_ifru</code> will be filled with an instance of <a href="#">can_baudrate_t</a> .
---------------------	------------------	---

Returns

0 on success, otherwise:

- -EFAULT: It was not possible to access user space memory area at the specified address.
- -ENODEV: No device with specified name exists.
- -EINVAL: No baud rate was set yet.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.1.2.13 `#define SIOCGCANCTRLMODE _IOWR(RTIOC_TYPE_CAN, 0x08, struct ifreq)`

Get special controller modes.

Parameters

<code>in</code>	<code>arg</code>	Pointer to interface request structure buffer (struct ifreq from linux/if.h). <code>ifr_name</code> must hold a valid CAN interface name, <code>ifr_ifru</code> must be filled with an instance of <a href="#">can_ctrlmode_t</a> .
-----------------	------------------	---

Returns

0 on success, otherwise:

- -EFAULT: It was not possible to access user space memory area at the specified address.
- -ENODEV: No device with specified name exists.
- -EINVAL: No baud rate was set yet.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.1.2.14 `#define SIOCGCANCUSTOMBITTIME _IOWR(RTIOC_TYPE_CAN, 0x04, struct ifreq)`

Get custom bit-time parameters.

Parameters

in,out	<i>arg</i>	Pointer to interface request structure buffer (struct ifreq from linux/if.h). <i>ifr_name</i> must hold a valid CAN interface name, <i>ifr_ifru</i> will be filled with an instance of struct <a href="#">can_bittime</a> .
--------	------------	---

Returns

0 on success, otherwise:

- -EFAULT: It was not possible to access user space memory area at the specified address.
- -ENODEV: No device with specified name exists.
- -EINVAL: No baud rate was set yet.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.1.2.15 `#define SIOCGCANSTATE _IOWR(RTIOC_TYPE_CAN, 0x06, struct ifreq)`

Get current state of CAN controller.

States are divided into main states and additional error indicators. A CAN controller is always in exactly one main state. CAN bus errors are registered by the CAN hardware and collected by the driver. There is one error indicator (bit) per error type. If this IOCTL is triggered the error types which occurred since the last call of this IOCTL are reported and thereafter the error indicators are cleared. See also [CAN controller states](#).

Parameters

in,out	<i>arg</i>	Pointer to interface request structure buffer (struct ifreq from linux/if.h). <i>ifr_name</i> must hold a valid CAN interface name, <i>ifr_ifru</i> will be filled with an instance of <a href="#">can_mode_t</a> .
--------	------------	---

## Returns

0 on success, otherwise:

- -EFAULT: It was not possible to access user space memory area at the specified address.
- -ENODEV: No device with specified name exists.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

### 5.1.2.16 #define SIOCGIFINDEX defined\_by\_kernel\_header\_file

Get CAN interface index by name.

#### Parameters

in,out	<i>arg</i>	Pointer to interface request structure buffer (struct ifreq from linux/if.h). If <i>ifr_name</i> holds a valid CAN interface name <i>ifr_ifindex</i> will be filled with the corresponding interface index.
--------	------------	---

## Returns

0 on success, otherwise:

- -EFAULT: It was not possible to access user space memory area at the specified address.
- -ENODEV: No device with specified name exists.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

## Examples:

[rtcan\\_rtt.c](#), [rtcanconfig.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

### 5.1.2.17 #define SIOCSCANBAUDRATE \_IOW(RTIOC\_TYPE\_CAN, 0x01, struct ifreq)

Set baud rate.

The baudrate must be specified in bits per second. The driver will try to calculate reasonable CAN bit-timing parameters. You can use [SIOCSCANCUSTOMBITTIME](#) to set custom bit-timing.

## Parameters

<i>in</i>	<i>arg</i>	Pointer to interface request structure buffer (struct ifreq from linux/if.h). <i>ifr_name</i> must hold a valid CAN interface name, <i>ifr_ifru</i> must be filled with an instance of <a href="#">can_baudrate_t</a> .
-----------	------------	---

## Returns

0 on success, otherwise:

- -EFAULT: It was not possible to access user space memory area at the specified address.
- -ENODEV: No device with specified name exists.
- -EINVAL: No valid baud rate, see [can\\_baudrate\\_t](#).
- -EDOM : Baud rate not possible.
- -EAGAIN: Request could not be successfully fulfilled. Try again.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

## Note

Setting the baud rate is a configuration task. It should be done deliberately or otherwise CAN messages will likely be lost.

Rescheduling: possible.

## Examples:

[rtcanconfig.c](#).

5.1.2.18 `#define SIOCSCANCTRLMODE _IOW(RTIOC_TYPE_CAN, 0x07, struct ifreq)`

Set special controller modes.

Various special controller modes could be or'ed together (see [CAN\\_CTRLMODE](#) for further information).

## Parameters

<i>in</i>	<i>arg</i>	Pointer to interface request structure buffer (struct ifreq from linux/if.h). <i>ifr_name</i> must hold a valid CAN interface name, <i>ifr_ifru</i> must be filled with an instance of <a href="#">can_ctrlmode_t</a> .
-----------	------------	---

## Returns

0 on success, otherwise:

- -EFAULT: It was not possible to access user space memory area at the specified address.
- -ENODEV: No device with specified name exists.
- -EINVAL: No valid baud rate, see [can\\_baudrate\\_t](#).
- -EAGAIN: Request could not be successfully fulfilled. Try again.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Note

Setting special controller modes is a configuration task. It should be done deliberately or otherwise CAN messages will likely be lost.

Rescheduling: possible.

Examples:

[rtcanconfig.c](#).

5.1.2.19 #define SIOCSCANCUSTOMBITTIME \_IOW(RTIOC\_TYPE\_CAN, 0x03, struct ifreq)

Set custom bit time parameter.

Custom-bit time could be defined in various formats (see struct [can\\_bittime](#)).

Parameters

<i>in</i>	<i>arg</i>	Pointer to interface request structure buffer (struct ifreq from linux/if.h). ifr_name must hold a valid CAN interface name, ifr_ifru must be filled with an instance of struct <a href="#">can_bittime</a> .
-----------	------------	---

Returns

0 on success, otherwise:

- -EFAULT: It was not possible to access user space memory area at the specified address.
- -ENODEV: No device with specified name exists.
- -EINVAL: No valid baud rate, see [can\\_baudrate\\_t](#).
- -EAGAIN: Request could not be successfully fulfilled. Try again.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Note

Setting the bit-time is a configuration task. It should be done deliberately or otherwise CAN messages will likely be lost.

Rescheduling: possible.

Examples:

[rtcanconfig.c](#).

5.1.2.20 #define SIOCSCANMODE\_IOW(RTIOC\_TYPE\_CAN, 0x05, struct ifreq)

Set operation mode of CAN controller.

See [CAN controller modes](#) for available modes.

Parameters

<i>in</i>	<i>arg</i>	Pointer to interface request structure buffer (struct ifreq from linux/if.h). ifr_name must hold a valid CAN interface name, ifr_ifru must be filled with an instance of <a href="#">can_mode_t</a> .
-----------	------------	---

Returns

0 on success, otherwise:

- -EFAULT: It was not possible to access user space memory area at the specified address.
- -ENODEV: No device with specified name exists.
- -EAGAIN: ([CAN\\_MODE\\_START](#), [CAN\\_MODE\\_STOP](#)) Could not successfully set mode, hardware is busy. Try again.
- -EINVAL: ([CAN\\_MODE\\_START](#)) Cannot start controller, set baud rate first.
- -ENETDOWN: ([CAN\\_MODE\\_SLEEP](#)) Cannot go into sleep mode because controller is stopped or bus off.
- -EOPNOTSUPP: unknown mode

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Note

Setting a CAN controller into normal operation after a bus-off can take some time (128 occurrences of 11 consecutive recessive bits). In such a case, although this IOCTL will return immediately with success and [SIOCGCANSTATE](#) will report CAN\_STATE\_ACTIVE, bus-off recovery may still be in progress.

If a controller is bus-off, setting it into stop mode will return no error but the controller remains bus-off.

Rescheduling: possible.

Examples:

[rtcanconfig.c](#).

5.1.2.21 #define SOL\_CAN\_RAW 103

CAN socket levels.

Used for [Socketops](#) for the particular protocols.

Examples:

[rtcan\\_rtt.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

## 5.1.3 Typedef Documentation

### 5.1.3.1 typedef struct **can\_filter** **can\_filter\_t**

Filter for reception of CAN messages.

This filter works as follows: A received CAN ID is AND'ed bitwise with `can_mask` and then compared to `can_id`. This also includes the `CAN_EFF_FLAG` and `CAN_RTR_FLAG` of `CAN_XXX_FLAG`. If this comparison is true, the message will be received by the socket. The logic can be inverted with the `can_id` flag `CAN_INV_FILTER` :

```
* if (can_id & CAN_INV_FILTER) {
*   if ((received_can_id & can_mask) != (can_id & ~CAN_INV_FILTER))
*     accept-message;
* } else {
*   if ((received_can_id & can_mask) == can_id)
*     accept-message;
* }
*
```

Multiple filters can be arranged in a filter list and set with `Sockopts`. If one of these filters matches a CAN ID upon reception of a CAN frame, this frame is accepted.

### 5.1.3.2 typedef struct **can\_frame** **can\_frame\_t**

Raw CAN frame.

Central structure for receiving and sending CAN frames.

Examples:

[rtcanrecv.c](#).

## 5.1.4 Enumeration Type Documentation

### 5.1.4.1 enum **CAN\_BITTIME\_TYPE**

Supported CAN bit-time types.

Enumerator

**CAN\_BITTIME\_STD** Standard bit-time definition according to Bosch.

**CAN\_BITTIME\_BTR** Hardware-specific BTR bit-time definition.

### 5.1.4.2 enum **CAN\_MODE**

Enumerator

**CAN\_MODE\_STOP** Set controller in Stop mode (no reception / transmission possible)

**CAN\_MODE\_START** Set controller into normal operation.

Coming from stopped mode or bus off, the controller begins with no errors in `CAN_STATE_ACTIVE`.

**CAN\_MODE\_SLEEP** Set controller into Sleep mode.

This is only possible if the controller is not stopped or bus-off.

Notice that sleep mode will only be entered when there is no bus activity. If the controller detects bus activity while "sleeping" it will go into operating mode again.

To actively leave sleep mode again trigger `CAN_MODE_START`.



### 5.1.4.3 enum **CAN\_STATE**

Enumerator

**CAN\_STATE\_ERROR\_ACTIVE** CAN controller is error active.

**CAN\_STATE\_ERROR\_WARNING** CAN controller is error active, warning level is reached.

**CAN\_STATE\_ERROR\_PASSIVE** CAN controller is error passive.

**CAN\_STATE\_BUS\_OFF** CAN controller went into Bus Off.

**CAN\_STATE\_SCANNING\_BAUDRATE** CAN controller is scanning to get the baudrate.

**CAN\_STATE\_STOPPED** CAN controller is in stopped mode.

**CAN\_STATE\_SLEEPING** CAN controller is in Sleep mode.

## 5.2 Real-Time Driver Model

### Modules

- [User API](#)
- [Driver Development API](#)
- [Device Profiles](#)

### Typedefs

- typedef uint64\_t [nanosecs\\_abs\\_t](#)  
*RTDM type for representing absolute dates.*
- typedef int64\_t [nanosecs\\_rel\\_t](#)  
*RTDM type for representing relative intervals.*

### API Versioning

- #define [RTDM\\_API\\_VER](#) 8  
*Common user and driver API version.*
- #define [RTDM\\_API\\_MIN\\_COMPAT\\_VER](#) 6  
*Minimum API revision compatible with the current release.*

### RTDM\_TIMEOUT\_XXX

#### Special timeout values

- #define [RTDM\\_TIMEOUT\\_INFINITE](#) 0  
*Block forever.*
- #define [RTDM\\_TIMEOUT\\_NONE](#) (-1)  
*Any negative timeout means non-blocking.*

#### 5.2.1 Detailed Description

The Real-Time Driver Model (RTDM) provides a unified interface to both users and developers of real-time device drivers. Specifically, it addresses the constraints of mixed RT/non-RT systems like Xenomai. RTDM conforms to POSIX semantics (IEEE Std 1003.1) where available and applicable.

#### **API Revision: 8**

#### 5.2.2 Macro Definition Documentation

##### 5.2.2.1 #define RTDM\_TIMEOUT\_INFINITE 0

Block forever.

##### 5.2.2.2 #define RTDM\_TIMEOUT\_NONE (-1)

Any negative timeout means non-blocking.

## 5.2.3 Typedef Documentation

### 5.2.3.1 typedef uint64\_t **nanosecs\_abs\_t**

RTDM type for representing absolute dates.

Its base type is a 64 bit unsigned integer. The unit is 1 nanosecond.

Examples:

[rtcanrecv.c](#).

### 5.2.3.2 typedef int64\_t **nanosecs\_rel\_t**

RTDM type for representing relative intervals.

Its base type is a 64 bit signed integer. The unit is 1 nanosecond. Relative intervals can also encode the special timeouts "infinite" and "non-blocking", see [RTDM\\_TIMEOUT\\_XXX](#).

Examples:

[rtcanrecv.c](#).

## 5.3 User API

### Files

- file [rtm.h](#)  
*Real-Time Driver Model for Xenomai, user API header.*

### Functions

- int [rt\\_dev\\_open](#) (const char \*path, int oflag,...)  
*Open a device.*
- int [rt\\_dev\\_socket](#) (int protocol\_family, int socket\_type, int protocol)  
*Create a socket.*
- int [rt\\_dev\\_close](#) (int fd)  
*Close a device or socket.*
- int [rt\\_dev\\_ioctl](#) (int fd, int request,...)  
*Issue an IOCTL.*
- ssize\_t [rt\\_dev\\_read](#) (int fd, void \*buf, size\_t nbyte)  
*Read from device.*
- ssize\_t [rt\\_dev\\_write](#) (int fd, const void \*buf, size\_t nbyte)  
*Write to device.*
- ssize\_t [rt\\_dev\\_recvmsg](#) (int fd, struct user\_msghdr \*msg, int flags)  
*Receive message from socket.*
- ssize\_t [rt\\_dev\\_recvfrom](#) (int fd, void \*buf, size\_t len, int flags, struct sockaddr \*from, socklen\_t \*fromlen)  
*Receive message from socket.*
- ssize\_t [rt\\_dev\\_recv](#) (int fd, void \*buf, size\_t len, int flags)  
*Receive message from socket.*
- ssize\_t [rt\\_dev\\_sendmsg](#) (int fd, const struct user\_msghdr \*msg, int flags)  
*Transmit message to socket.*
- ssize\_t [rt\\_dev\\_sendto](#) (int fd, const void \*buf, size\_t len, int flags, const struct sockaddr \*to, socklen\_t tolen)  
*Transmit message to socket.*
- ssize\_t [rt\\_dev\\_send](#) (int fd, const void \*buf, size\_t len, int flags)  
*Transmit message to socket.*
- int [rt\\_dev\\_bind](#) (int fd, const struct sockaddr \*my\_addr, socklen\_t addrlen)  
*Bind to local address.*
- int [rt\\_dev\\_connect](#) (int fd, const struct sockaddr \*serv\_addr, socklen\_t addrlen)  
*Connect to remote address.*
- int [rt\\_dev\\_listen](#) (int fd, int backlog)  
*Listen for incoming connection requests.*
- int [rt\\_dev\\_accept](#) (int fd, struct sockaddr \*addr, socklen\_t \*addrlen)  
*Accept a connection requests.*
- int [rt\\_dev\\_shutdown](#) (int fd, int how)  
*Shut down parts of a connection.*
- int [rt\\_dev\\_getsockopt](#) (int fd, int level, int optname, void \*optval, socklen\_t \*optlen)  
*Get socket option.*
- int [rt\\_dev\\_setsockopt](#) (int fd, int level, int optname, const void \*optval, socklen\_t optlen)  
*Set socket option.*
- int [rt\\_dev\\_getsockname](#) (int fd, struct sockaddr \*name, socklen\_t \*namelen)  
*Get local socket address.*
- int [rt\\_dev\\_getpeername](#) (int fd, struct sockaddr \*name, socklen\_t \*namelen)  
*Get socket destination address.*

### 5.3.1 Detailed Description

This is the upper interface of RTDM provided to application programs both in kernel and user space. Note that certain functions may not be implemented by every device. Refer to the [Device Profiles](#) for precise information.

### 5.3.2 Function Documentation

#### 5.3.2.1 `int rt_dev_accept ( int fd, struct sockaddr * addr, socklen_t * addrlen )`

Accept a connection requests.

Parameters

<code>in</code>	<code>fd</code>	File descriptor as returned by <code>rt_dev_socket()</code>
<code>out</code>	<code>addr</code>	Buffer for remote address
<code>in,out</code>	<code>addrlen</code>	Address buffer size

Returns

0 on success, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

`accept()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

#### 5.3.2.2 `int rt_dev_bind ( int fd, const struct sockaddr * my_addr, socklen_t addrlen )`

Bind to local address.

Parameters

<code>in</code>	<code>fd</code>	File descriptor as returned by <code>rt_dev_socket()</code>
<code>in</code>	<code>my_addr</code>	Address buffer
<code>in</code>	<code>addrlen</code>	Address buffer size

Returns

0 on success, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

`bind()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

Examples:

[rtcanrecv.c](#), and [rtcansend.c](#).

---

### 5.3.2.3 int rt\_dev\_close ( int fd )

Close a device or socket.

## Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_open()</code> or <code>rt_dev_socket()</code>
-----------	-----------	---

## Returns

0 on success, otherwise a negative error code.

## Note

If the matching `rt_dev_open()` or `rt_dev_socket()` call took place in non-real-time context, `rt_dev_close()` must be issued within non-real-time as well. Otherwise, the call will fail.

## Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

## See Also

`close()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.3.2.4 `int rt_dev_connect ( int fd, const struct sockaddr * serv_addr, socklen_t addrlen )`

Connect to remote address.

## Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_socket()</code>
<i>in</i>	<i>serv_addr</i>	Address buffer
<i>in</i>	<i>addrlen</i>	Address buffer size

## Returns

0 on success, otherwise negative error code

## Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

## See Also

`connect()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.3.2.5 `int rt_dev_getpeername ( int fd, struct sockaddr * name, socklen_t * namelen )`

Get socket destination address.

## Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_socket()</code>
-----------	-----------	---

out	<i>name</i>	Address buffer
in,out	<i>namelen</i>	Address buffer size

Returns

0 on success, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

getpeername() in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.3.2.6 int rt\_dev\_getsockname ( int fd, struct sockaddr \* name, socklen\_t \* namelen )

Get local socket address.

Parameters

in	<i>fd</i>	File descriptor as returned by rt_dev_socket()
out	<i>name</i>	Address buffer
in,out	<i>namelen</i>	Address buffer size

Returns

0 on success, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

getsockname() in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.3.2.7 int rt\_dev\_getsockopt ( int fd, int level, int optname, void \* optval, socklen\_t \* optlen )

Get socket option.

Parameters

in	<i>fd</i>	File descriptor as returned by rt_dev_socket()
in	<i>level</i>	Addressed stack level
in	<i>optname</i>	Option name ID
out	<i>optval</i>	Value buffer
in,out	<i>optlen</i>	Value buffer size



Returns

0 on success, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

`getsockopt()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.3.2.8 `int rt_dev_ioctl ( int fd, int request, ... )`

Issue an IOCTL.

Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_open()</code> or <code>rt_dev_socket()</code>
<i>in</i>	<i>request</i>	IOCTL code
	<i>...</i>	Optional third argument, depending on IOCTL function ( <code>void *</code> or <code>unsigned long</code> )

Returns

Positiv value on success, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

`ioctl()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.3.2.9 `int rt_dev_listen ( int fd, int backlog )`

Listen for incoming connection requests.

Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_socket()</code>
<i>in</i>	<i>backlog</i>	Maximum queue length

Returns

0 on success, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

`listen()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

---

5.3.2.10 `int rt_dev_open ( const char * path, int oflag, ... )`

Open a device.

## Parameters

<i>in</i>	<i>path</i>	Device name
<i>in</i>	<i>oflag</i>	Open flags
	<i>...</i>	Further parameters will be ignored.

## Returns

Positive file descriptor value on success, otherwise a negative error code.

## Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

## See Also

`open()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

### 5.3.2.11 `ssize_t rt_dev_read ( int fd, void * buf, size_t nbyte )`

Read from device.

## Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_open()</code>
<i>out</i>	<i>buf</i>	Input buffer
<i>in</i>	<i>nbyte</i>	Number of bytes to read

## Returns

Number of bytes read, otherwise negative error code

## Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

## See Also

`read()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

### 5.3.2.12 `ssize_t rt_dev_rcv ( int fd, void * buf, size_t len, int flags )`

Receive message from socket.

## Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_socket()</code>
<i>out</i>	<i>buf</i>	Message buffer
<i>in</i>	<i>len</i>	Message buffer size
<i>in</i>	<i>flags</i>	Message flags

Returns

Number of bytes received, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

recv() in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.3.2.13 ssize\_t rt\_dev\_recvfrom ( int fd, void \* buf, size\_t len, int flags, struct sockaddr \* from, socklen\_t \* fromlen )

Receive message from socket.

Parameters

in	<i>fd</i>	File descriptor as returned by <code>rt_dev_socket()</code>
out	<i>buf</i>	Message buffer
in	<i>len</i>	Message buffer size
in	<i>flags</i>	Message flags
out	<i>from</i>	Buffer for message sender address
in,out	<i>fromlen</i>	Address buffer size

Returns

Number of bytes received, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

recvfrom() in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

References `rt_dev_recvfrom()`.

Referenced by `rt_dev_recvfrom()`.

5.3.2.14 ssize\_t rt\_dev\_recvmsg ( int fd, struct user\_msghdr \* msg, int flags )

Receive message from socket.

Parameters

in	<i>fd</i>	File descriptor as returned by <code>rt_dev_socket()</code>
in,out	<i>msg</i>	Message descriptor
in	<i>flags</i>	Message flags

## Returns

Number of bytes received, otherwise negative error code

## Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

## See Also

`recvmsg()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

## Examples:

[rtcanrecv.c](#).

5.3.2.15 `ssize_t rt_dev_send ( int fd, const void * buf, size_t len, int flags )`

Transmit message to socket.

## Parameters

<code>in</code>	<code>fd</code>	File descriptor as returned by <code>rt_dev_socket()</code>
<code>in</code>	<code>buf</code>	Message buffer
<code>in</code>	<code>len</code>	Message buffer size
<code>in</code>	<code>flags</code>	Message flags

## Returns

Number of bytes sent, otherwise negative error code

## Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

## See Also

`send()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

## Examples:

[rtcansend.c](#).

5.3.2.16 `ssize_t rt_dev_sendmsg ( int fd, const struct user_msghdr * msg, int flags )`

Transmit message to socket.

## Parameters

<code>in</code>	<code>fd</code>	File descriptor as returned by <code>rt_dev_socket()</code>
<code>in</code>	<code>msg</code>	Message descriptor

<i>in</i>	<i>flags</i>	Message flags
-----------	--------------	---------------

Returns

Number of bytes sent, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

`sendmsg()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.3.2.17 `ssize_t rt_dev_sendto ( int fd, const void * buf, size_t len, int flags, const struct sockaddr * to, socklen_t tolen )`

Transmit message to socket.

Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_socket()</code>
<i>in</i>	<i>buf</i>	Message buffer
<i>in</i>	<i>len</i>	Message buffer size
<i>in</i>	<i>flags</i>	Message flags
<i>in</i>	<i>to</i>	Buffer for message destination address
<i>in</i>	<i>tolen</i>	Address buffer size

Returns

Number of bytes sent, otherwise negative error code

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

See Also

`sendto()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

Examples:

[rtcansend.c](#).

5.3.2.18 `int rt_dev_setsockopt ( int fd, int level, int optname, const void * optval, socklen_t optlen )`

Set socket option.

Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_socket()</code>
<i>in</i>	<i>level</i>	Addressed stack level
<i>in</i>	<i>optname</i>	Option name ID
<i>in</i>	<i>optval</i>	Value buffer
<i>in</i>	<i>optlen</i>	Value buffer size

#### Returns

0 on success, otherwise negative error code

#### Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

#### See Also

`setsockopt()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

#### Examples:

[rtcanrecv.c](#), and [rtcansend.c](#).

#### 5.3.2.19 `int rt_dev_shutdown ( int fd, int how )`

Shut down parts of a connection.

#### Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_socket()</code>
<i>in</i>	<i>how</i>	Specifies the part to be shut down ( <code>SHUT_XXX</code> )

#### Returns

0 on success, otherwise negative error code

#### Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

#### See Also

`shutdown()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

#### 5.3.2.20 `int rt_dev_socket ( int protocol_family, int socket_type, int protocol )`

Create a socket.

#### Parameters

<i>in</i>	<i>protocol_family</i>	Protocol family ( <code>PF_XXX</code> )
<i>in</i>	<i>socket_type</i>	Socket type ( <code>SOCK_XXX</code> )
<i>in</i>	<i>protocol</i>	Protocol ID, 0 for default

## Returns

Positive file descriptor value on success, otherwise a negative error code.

## Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

## See Also

`socket()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.3.2.21 `ssize_t rt_dev_write ( int fd, const void * buf, size_t nbyte )`

Write to device.

## Parameters

<i>in</i>	<i>fd</i>	File descriptor as returned by <code>rt_dev_open()</code>
<i>in</i>	<i>buf</i>	Output buffer
<i>in</i>	<i>nbyte</i>	Number of bytes to write

## Returns

Number of bytes written, otherwise negative error code

## Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

## See Also

`write()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>



## 5.4 Real-time IPC protocols

**Profile Revision: 1**

### Files

- file [rtipc.h](#)  
*This file is part of the Xenomai project.*

### Data Structures

- struct [rtipc\\_port\\_label](#)  
*Port label information structure.*
- struct [sockaddr\\_ipc](#)  
*Socket address structure for the RTIPC address family.*

### Typedefs

- typedef int16\_t [rtipc\\_port\\_t](#)  
*Port number type for the RTIPC address family.*

### Supported operations

Standard socket operations supported by the RTIPC protocols.

- int [socket\\_\\_AF\\_RTIPC](#) (int domain=AF\_RTIPC, int type=SOCK\_DGRAM, int protocol)  
*Create an endpoint for communication in the AF\_RTIPC domain.*
- int [close\\_\\_AF\\_RTIPC](#) (int sockfd)  
*Close a RTIPC socket descriptor.*
- int [bind\\_\\_AF\\_RTIPC](#) (int sockfd, const struct [sockaddr\\_ipc](#) \*addr, socklen\_t addrlen)  
*Bind a RTIPC socket to a port.*
- int [connect\\_\\_AF\\_RTIPC](#) (int sockfd, const struct [sockaddr\\_ipc](#) \*addr, socklen\_t addrlen)  
*Initiate a connection on a RTIPC socket.*
- int [setsockopt\\_\\_AF\\_RTIPC](#) (int sockfd, int level, int optname, const void \*optval, socklen\_t optlen)  
*Set options on RTIPC sockets.*
- int [getsockopt\\_\\_AF\\_RTIPC](#) (int sockfd, int level, int optname, void \*optval, socklen\_t \*optlen)  
*Get options on RTIPC sockets.*
- ssize\_t [sendmsg\\_\\_AF\\_RTIPC](#) (int sockfd, const struct user\_msghdr \*msg, int flags)  
*Send a message on a RTIPC socket.*
- ssize\_t [recvmsg\\_\\_AF\\_RTIPC](#) (int sockfd, struct user\_msghdr \*msg, int flags)  
*Receive a message from a RTIPC socket.*
- int [getsockname\\_\\_AF\\_RTIPC](#) (int sockfd, struct [sockaddr\\_ipc](#) \*addr, socklen\_t \*addrlen)  
*Get socket name.*
- int [getpeername\\_\\_AF\\_RTIPC](#) (int sockfd, struct [sockaddr\\_ipc](#) \*addr, socklen\_t \*addrlen)  
*Get socket peer.*

## RTIPC protocol list

protocols for the PF\_RTIPC protocol family

- enum { `IPCPROTO_IPC` = 0, `IPCPROTO_XDDP` = 1, `IPCPROTO_IDDP` = 2, `IPCPROTO_BUFP` = 3 }

## XDDP socket options

Setting and getting XDDP socket options.

- #define `XDDP_LABEL` 1  
*XDDP label assignment.*
- #define `XDDP_POOLSZ` 2  
*XDDP local pool size configuration.*
- #define `XDDP_BUFSZ` 3  
*XDDP streaming buffer size configuration.*
- #define `XDDP_MONITOR` 4  
*XDDP monitoring callback.*

## XDDP events

Specific events occurring on XDDP channels, which can be monitored via the `XDDP_MONITOR` socket option.

- #define `XDDP_EVTIN` 1  
*Monitor writes to the non real-time endpoint.*
- #define `XDDP_EVTOUT` 2  
*Monitor reads from the non real-time endpoint.*
- #define `XDDP_EVTDOWN` 3  
*Monitor close from the non real-time endpoint.*
- #define `XDDP_EVTNOBUF` 4  
*Monitor memory shortage for non real-time datagrams.*

## IDDP socket options

Setting and getting IDDP socket options.

- #define `IDDP_LABEL` 1  
*IDDP label assignment.*
- #define `IDDP_POOLSZ` 2  
*IDDP local pool size configuration.*

## BUFP socket options

Setting and getting BUFP socket options.

- #define `BUFP_LABEL` 1  
*BUFP label assignment.*
- #define `BUFP_BUFSZ` 2  
*BUFP buffer size configuration.*

## Socket level options

Setting and getting supported standard socket level options.

- #define `SO_SNDTIMEO` defined\_by\_kernel\_header\_file  
*IPCPROTO\_IDDP and IPCPROTO\_BUFPP protocols support the standard SO\_SNDTIMEO socket option, from the SOL\_SOCKET level.*
- #define `SO_RCVTIMEO` defined\_by\_kernel\_header\_file  
*All RTIPC protocols support the standard SO\_RCVTIMEO socket option, from the SOL\_SOCKET level.*

### 5.4.1 Detailed Description

#### Profile Revision: 1

Device Characteristics

**Device Flags:** RTDM\_PROTOCOL\_DEVICE  
**Protocol Family:** PF\_RTIPC  
**Socket Type:** SOCK\_DGRAM  
**Device Class:** RTDM\_CLASS\_RTIPC

### 5.4.2 Macro Definition Documentation

#### 5.4.2.1 #define BUFP\_BUFSZ 2

BUFP buffer size configuration.

All messages written to a BUFP socket are buffered in a single per-socket memory area. Configuring the size of such buffer prior to binding the socket to a destination port is mandatory.

It is not allowed to configure a buffer size after the socket was bound. However, multiple configuration calls are allowed prior to the binding; the last value set will be used.

Note

: the buffer memory is obtained from the host allocator by the [bind call](#).

Parameters

in	level	<a href="#">SOL_BUFPP</a>
in	optname	<b>BUFP_BUFSZ</b>
in	optval	Pointer to a variable of type size_t, containing the required size of the buffer to reserve at binding time
in	optlen	sizeof(size_t)

Returns

0 is returned upon success. Otherwise:

- -EFAULT (Invalid data address given)
- -EALREADY (socket already bound)
- -EINVAL (*optlen* is invalid or *\*optval* is zero)

Calling context:

RT/non-RT

Examples:

[bufp-label.c](#), and [bufp-readwrite.c](#).

#### 5.4.2.2 #define BUFP\_LABEL 1

BUFP label assignment.

ASCII label strings can be attached to BUFP ports, in order to connect sockets to them in a more descriptive way than using plain numeric port values.

When available, this label will be registered when binding, in addition to the port number (see [BUFP port binding](#)).

It is not allowed to assign a label after the socket was bound. However, multiple assignment calls are allowed prior to the binding; the last label set will be used.

Parameters

in	<i>level</i>	<a href="#">SOL_BUFP</a>
in	<i>optname</i>	<b>BUFP_LABEL</b>
in	<i>optval</i>	Pointer to struct <a href="#">rtipc_port_label</a>
in	<i>optlen</i>	sizeof(struct <a href="#">rtipc_port_label</a> )

Returns

0 is returned upon success. Otherwise:

- -EFAULT (Invalid data address given)
- -EALREADY (socket already bound)
- -EINVAL (*optlen* is invalid)

Calling context:

RT/non-RT

Examples:

[bufp-label.c](#).

#### 5.4.2.3 #define IDDP\_LABEL 1

IDDP label assignment.

ASCII label strings can be attached to IDDP ports, in order to connect sockets to them in a more descriptive way than using plain numeric port values.

When available, this label will be registered when binding, in addition to the port number (see [IDDP port binding](#)).

It is not allowed to assign a label after the socket was bound. However, multiple assignment calls are allowed prior to the binding; the last label set will be used.

## Parameters

<code>in</code>	<i>level</i>	<a href="#">SOL_IDDP</a>
<code>in</code>	<i>optname</i>	<b>IDDP_LABEL</b>
<code>in</code>	<i>optval</i>	Pointer to struct <a href="#">rtipc_port_label</a>
<code>in</code>	<i>optlen</i>	sizeof(struct <a href="#">rtipc_port_label</a> )

## Returns

0 is returned upon success. Otherwise:

- -EFAULT (Invalid data address given)
- -EALREADY (socket already bound)
- -EINVAL (*optlen* is invalid)

## Calling context:

RT/non-RT

## Examples:

[iddp-label.c](#).

### 5.4.2.4 #define IDDP\_POOLSZ 2

IDDP local pool size configuration.

By default, the memory needed to convey the data is pulled from Xenomai's system pool. Setting a local pool size overrides this default for the socket.

If a non-zero size was configured, a local pool is allocated at binding time. This pool will provide storage for pending datagrams.

It is not allowed to configure a local pool size after the socket was bound. However, multiple configuration calls are allowed prior to the binding; the last value set will be used.

## Note

: the pool memory is obtained from the host allocator by the [bind call](#).

## Parameters

<code>in</code>	<i>level</i>	<a href="#">SOL_IDDP</a>
<code>in</code>	<i>optname</i>	<b>IDDP_POOLSZ</b>
<code>in</code>	<i>optval</i>	Pointer to a variable of type <code>size_t</code> , containing the required size of the local pool to reserve at binding time
<code>in</code>	<i>optlen</i>	sizeof( <code>size_t</code> )

## Returns

0 is returned upon success. Otherwise:

- -EFAULT (Invalid data address given)
- -EALREADY (socket already bound)
- -EINVAL (*optlen* is invalid or *\*optval* is zero)

Calling context:

RT/non-RT

Examples:

[iddp-sendrecv.c](#).

#### 5.4.2.5 #define SO\_RCVTIMEO defined\_by\_kernel\_header\_file

All RTIPC protocols support the standard SO\_RCVTIMEO socket option, from the SOL\_SOCKET level.

See Also

setsockopt(), getsockopt() in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399/>

Examples:

[xddp-label.c](#).

#### 5.4.2.6 #define SO\_SNDTIMEO defined\_by\_kernel\_header\_file

IPPROTO\_IDDP and IPPROTO\_BUF protocols support the standard SO\_SNDTIMEO socket option, from the SOL\_SOCKET level.

See Also

setsockopt(), getsockopt() in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399/>

#### 5.4.2.7 #define XDDP\_BUFSZ 3

XDDP streaming buffer size configuration.

In addition to sending datagrams, real-time threads may stream data in a byte-oriented mode through the port as well. This increases the bandwidth and reduces the overhead, when the overall data to send to the Linux domain is collected by bits, and keeping the message boundaries is not required.

This feature is enabled when a non-zero buffer size is set for the socket. In that case, the real-time data accumulates into the streaming buffer when MSG\_MORE is passed to any of the [send functions](#), until:

- the receiver from the Linux domain wakes up and consumes it,
- a different source port attempts to send data to the same destination port,
- MSG\_MORE is absent from the send flags,
- the buffer is full,

whichever comes first.

Setting *\*optval* to zero disables the streaming buffer, in which case all sendings are conveyed in separate datagrams, regardless of MSG\_MORE.

Note

only a single streaming buffer exists per socket. When this buffer is full, the real-time data stops accumulating and sending operations resume in mere datagram mode. Accumulation may happen again after some or all data in the streaming buffer is consumed from the Linux domain endpoint.

The streaming buffer size may be adjusted multiple times during the socket lifetime; the latest configuration change will take effect when the accumulation resumes after the previous buffer was flushed.

## Parameters

in	<i>level</i>	<a href="#">SOL_XDDP</a>
in	<i>optname</i>	<b>XDDP_BUFSZ</b>
in	<i>optval</i>	Pointer to a variable of type <code>size_t</code> , containing the required size of the streaming buffer
in	<i>optlen</i>	<code>sizeof(size_t)</code>

## Returns

0 is returned upon success. Otherwise:

- -EFAULT (Invalid data address given)
- -ENOMEM (Not enough memory)
- -EINVAL (*optlen* is invalid)

## Calling context:

RT/non-RT

## Examples:

[xddp-stream.c](#).

### 5.4.2.8 #define XDDP\_EVTDOWN 3

**Monitor** close from the non real-time endpoint.

XDDP\_EVTDOWN is sent when the non real-time endpoint is closed. The argument is always 0.

### 5.4.2.9 #define XDDP\_EVTIN 1

**Monitor** writes to the non real-time endpoint.

XDDP\_EVTIN is sent when data is written to the non real-time endpoint the socket is bound to (i.e. via `/dev/rtpN`), which means that some input is pending for the real-time endpoint. The argument is the size of the incoming message.

### 5.4.2.10 #define XDDP\_EVTNOBUF 4

**Monitor** memory shortage for non real-time datagrams.

XDDP\_EVTNOBUF is sent when no memory is available from the pool to hold the message currently sent from the non real-time endpoint. The argument is the size of the failed allocation. Upon return from the callback, the caller will block and retry until enough space is available from the pool; during that process, the callback might be invoked multiple times, each time a new attempt to get the required memory fails.

### 5.4.2.11 #define XDDP\_EVTOUT 2

**Monitor** reads from the non real-time endpoint.

XDDP\_EVTOUT is sent when the non real-time endpoint successfully reads a complete message (i.e. via `/dev/rtpN`). The argument is the size of the outgoing message.

#### 5.4.2.12 #define XDDP\_LABEL 1

XDDP label assignment.

ASCII label strings can be attached to XDDP ports, so that opening the non-RT endpoint can be done by specifying this symbolic device name rather than referring to a raw pseudo-device entry (i.e. /dev/rtpN).

When available, this label will be registered when binding, in addition to the port number (see [XDDP port binding](#)).

It is not allowed to assign a label after the socket was bound. However, multiple assignment calls are allowed prior to the binding; the last label set will be used.

Parameters

in	level	<a href="#">SOL_XDDP</a>
in	optname	<b>XDDP_LABEL</b>
in	optval	Pointer to struct <a href="#">rtipc_port_label</a>
in	optlen	sizeof(struct <a href="#">rtipc_port_label</a> )

Returns

0 is returned upon success. Otherwise:

- -EFAULT (Invalid data address given)
- -EALREADY (socket already bound)
- -EINVAL (*optlen* invalid)

Calling context:

RT/non-RT

Examples:

[xddp-label.c](#).

#### 5.4.2.13 #define XDDP\_MONITOR 4

XDDP monitoring callback.

Other RTDM drivers may install a user-defined callback via the [rtdm\\_setsockopt](#) call from the inter-driver API, in order to collect particular events occurring on the channel.

This notification mechanism is particularly useful to monitor a channel asynchronously while performing other tasks.

The user-provided routine will be passed the RTDM file descriptor of the socket receiving the event, the event code, and an optional argument. Four events are currently defined, see [XDDP\\_EVENTS](#).

The XDDP\_EVTIN and XDDP\_EVTOUT events are fired on behalf of a fully atomic context; therefore, care must be taken to keep their overhead low. In those cases, the Xenomai services that may be called from the callback are restricted to the set allowed to a real-time interrupt handler.

Parameters

in	level	<a href="#">SOL_XDDP</a>
----	-------	--------------------------



<i>in</i>	<i>optname</i>	<b>XDDP_MONITOR</b>
<i>in</i>	<i>optval</i>	Pointer to a pointer to function of type <code>int (*)(int fd, int event, long arg)</code> , containing the address of the user-defined callback. Passing a NULL callback pointer in <i>optval</i> disables monitoring.
<i>in</i>	<i>optlen</i>	<code>sizeof(int (*)(int fd, int event, long arg))</code>

#### Returns

0 is returned upon success. Otherwise:

- -EFAULT (Invalid data address given)
- -EPERM (Operation not allowed from user-space)
- -EINVAL (*optlen* is invalid)

#### Calling context:

RT/non-RT, kernel space only

#### 5.4.2.14 #define XDDP\_POOLSZ 2

XDDP local pool size configuration.

By default, the memory needed to convey the data is pulled from Xenomai's system pool. Setting a local pool size overrides this default for the socket.

If a non-zero size was configured, a local pool is allocated at binding time. This pool will provide storage for pending datagrams.

It is not allowed to configure a local pool size after the socket was bound. However, multiple configuration calls are allowed prior to the binding; the last value set will be used.

#### Note

: the pool memory is obtained from the host allocator by the [bind call](#).

#### Parameters

<i>in</i>	<i>level</i>	<a href="#">SOL_XDDP</a>
<i>in</i>	<i>optname</i>	<b>XDDP_POOLSZ</b>
<i>in</i>	<i>optval</i>	Pointer to a variable of type <code>size_t</code> , containing the required size of the local pool to reserve at binding time
<i>in</i>	<i>optlen</i>	<code>sizeof(size_t)</code>

#### Returns

0 is returned upon success. Otherwise:

- -EFAULT (Invalid data address given)
- -EALREADY (socket already bound)
- -EINVAL (*optlen* invalid or *\*optval* is zero)

#### Calling context:

RT/non-RT

#### Examples:

[xddp-echo.c](#).

## 5.4.3 Enumeration Type Documentation

### 5.4.3.1 anonymous enum

#### Enumerator

**IPCPROTO\_IPC** Default protocol (IDDP)

**IPCPROTO\_XDDP** Cross-domain datagram protocol (RT <-> non-RT). Real-time Xenomai threads and regular Linux threads may want to exchange data in a way that does not require the former to leave the real-time domain (i.e. primary mode). The RTDM-based XDDP protocol is available for this purpose.

On the Linux domain side, pseudo-device files named `/dev/rtp<minor>` give regular POSIX threads access to non real-time communication endpoints, via the standard character-based I/O interface. On the Xenomai domain side, sockets may be bound to XDDP ports, which act as proxies to send and receive data to/from the associated pseudo-device files. Ports and pseudo-device minor numbers are paired, meaning that e.g. socket port 7 will proxy the traffic to/from `/dev/rtp7`.

All data sent through a bound/connected XDDP socket via `sendto(2)` or `write(2)` will be passed to the peer endpoint in the Linux domain, and made available for reading via the standard `read(2)` system call. Conversely, all data sent using `write(2)` through the non real-time endpoint will be conveyed to the real-time socket endpoint, and made available to the `recvfrom(2)` or `read(2)` system calls.

**IPCPROTO\_IDDP** Intra-domain datagram protocol (RT <-> RT). The RTDM-based IDDP protocol enables real-time threads to exchange datagrams within the Xenomai domain, via socket endpoints.

**IPCPROTO\_BUF** Buffer protocol (RT <-> RT, byte-oriented). The RTDM-based BUFP protocol implements a lightweight, byte-oriented, one-way Producer-Consumer data path. All messages written are buffered into a single memory area in strict FIFO order, until read by the consumer.

This protocol always prevents short writes, and only allows short reads when a potential deadlock situation arises (i.e. readers and writers waiting for each other indefinitely).

## 5.4.4 Function Documentation

### 5.4.4.1 `int bind__AF_RTIPC ( int sockfd, const struct sockaddr_ipc * addr, socklen_t addrlen )`

Bind a RTIPC socket to a port.

Bind the socket to a destination port.

#### Parameters

<code>in</code>	<code>addr</code>	The address to bind the socket to (see struct <a href="#">sockaddr_ipc</a> ). The meaning of such address depends on the RTIPC protocol in use for the socket:
-----------------	-------------------	--

- **IPCPROTO\_XDDP**

This action creates an endpoint for channelling traffic between the Xenomai and Linux domains.

`sipc_family` must be `AF_RTIPC`, `sipc_port` is either `-1`, or a valid free port number between `0` and `CONFIG_XENO_OPT_PIPE_NRDEV-1`.

If `sipc_port` is `-1`, a free port will be assigned automatically.

Upon success, the pseudo-device `/dev/rtpN` will be reserved for this communication channel, where `N` is the assigned port number. The non real-time side shall open this device to exchange data over the bound socket.

If a label was assigned (see [XDDP\\_LABEL](#)) prior to binding the socket to a port, a registry link referring to the created pseudo-device will be automatically set up as `/proc/xenomai/registry/rtipc/xddp/label`, where *label* is the label string passed to `setsockopt()` for the [XDDP\\_LABEL](#) option.

- IPCPROTO\_IDDP

This action creates an endpoint for exchanging datagrams within the Xenomai domain.

*sipc\_family* must be `AF_RTIPC`, *sipc\_port* is either `-1`, or a valid free port number between 0 and `CONFIG_XENO_OPT_IDDP_NRPORT-1`.

If *sipc\_port* is `-1`, a free port will be assigned automatically. The real-time peer shall connect to the same port for exchanging data over the bound socket.

If a label was assigned (see [IDDP\\_LABEL](#)) prior to binding the socket to a port, a registry link referring to the assigned port number will be automatically set up as `/proc/xenomai/registry/rtipc/iddp/label`, where *label* is the label string passed to `setsockopt()` for the [IDDP\\_LABEL](#) option.

- IPCPROTO\_BUFPP

This action creates an endpoint for a one-way byte stream within the Xenomai domain.

*sipc\_family* must be `AF_RTIPC`, *sipc\_port* is either `-1`, or a valid free port number between 0 and `CONFIG_XENO_OPT_BUFPP_NRPORT-1`.

If *sipc\_port* is `-1`, an available port will be assigned automatically. The real-time peer shall connect to the same port for exchanging data over the bound socket.

If a label was assigned (see [BUFP\\_LABEL](#)) prior to binding the socket to a port, a registry link referring to the assigned port number will be automatically set up as `/proc/xenomai/registry/rtipc/bufp/label`, where *label* is the label string passed to `setsockopt()` for the [BUFP\\_LABEL](#) option.

#### Returns

In addition to the standard error codes for `bind(2)`, the following specific error code may be returned:

- `-EFAULT` (Invalid data address given)
- `-ENOMEM` (Not enough memory)
- `-EINVAL` (Invalid parameter)
- `-EADDRINUSE` (Socket already bound to a port, or no port available)

#### Calling context:

non-RT

#### 5.4.4.2 `int close__AF_RTIPC ( int sockfd )`

Close a RTIPC socket descriptor.

Blocking calls to any of the [sendmsg](#) or [recvmsg](#) functions will be unblocked when the socket is closed and return with an error.

#### Returns

In addition to the standard error codes for `close(2)`, the following specific error code may be returned: none

#### Calling context:

non-RT

---

5.4.4.3 int connect\_\_AF\_RTIPC ( int sockfd, const struct **sockaddr\_ipc** \* addr, socklen\_t addrlen )

Initiate a connection on a RTIPC socket.

## Parameters

<code>in</code>	<code>addr</code>	The address to connect the socket to (see struct <a href="#">sockaddr_ipc</a> ).
-----------------	-------------------	--

- If `sipc_port` is a valid port for the protocol, it is used verbatim and the connection succeeds immediately, regardless of whether the destination is bound at the time of the call.
- If `sipc_port` is -1 and a label was assigned to the socket, `connect()` blocks for the requested amount of time (see [SO\\_RCVTIMEO](#)) until a socket is bound to the same label via `bind(2)` (see [XDDP\\_LABEL](#), [IDDP\\_LABEL](#), [BUFP\\_LABEL](#)), in which case a connection is established between both endpoints.
- If `sipc_port` is -1 and no label was assigned to the socket, the default destination address is cleared, meaning that any subsequent write to the socket will return `-EDESTADDRREQ`, until a valid destination address is set via `connect(2)` or `bind(2)`.

## Returns

In addition to the standard error codes for `connect(2)`, the following specific error code may be returned: none.

## Calling context:

RT/non-RT

5.4.4.4 `int getpeername__AF_RTIPC ( int sockfd, struct sockaddr_ipc * addr, socklen_t * addrlen )`

### Get socket peer.

The name of the remote endpoint for the socket is copied back (see struct [sockaddr\\_ipc](#)). This is the default destination address for messages sent on the socket. It can be set either explicitly via `connect(2)`, or implicitly via `bind(2)` if no `connect(2)` was called prior to binding the socket to a port, in which case both the local and remote names are equal.

## Returns

In addition to the standard error codes for `getpeername(2)`, the following specific error code may be returned: none.

## Calling context:

RT/non-RT

5.4.4.5 `int getsockname__AF_RTIPC ( int sockfd, struct sockaddr_ipc * addr, socklen_t * addrlen )`

### Get socket name.

The name of the local endpoint for the socket is copied back (see struct [sockaddr\\_ipc](#)).

## Returns

In addition to the standard error codes for `getsockname(2)`, the following specific error code may be returned: none.

## Calling context:

RT/non-RT

5.4.4.6 int getsockopt\_\_AF\_RTIPC ( int sockfd, int level, int optname, void \* optval, socklen\_t \* optlen )

Get options on RTIPC sockets.

These functions allow to get various socket options. Supported Levels and Options:

- Level [SOL\\_SOCKET](#)
- Level [SOL\\_XDDP](#)
- Level [SOL\\_IDDP](#)
- Level [SOL\\_BUF](#)

Returns

In addition to the standard error codes for `getsockopt(2)`, the following specific error code may be returned: follow the option links above.

Calling context:

RT/non-RT

5.4.4.7 ssize\_t recvmsg\_\_AF\_RTIPC ( int sockfd, struct user\_msghdr \* msg, int flags )

Receive a message from a RTIPC socket.

Parameters

<i>in</i>	<i>flags</i>	Operation flags:
-----------	--------------	------------------

- `MSG_DONTWAIT` Non-blocking I/O operation. The caller will not be blocked whenever no message is immediately available for receipt at the time of the call, but will rather return with `-EWOULDBLOCK`.

Note

[IPCPROTO\\_BUF](#) does not allow for short reads and always returns the requested amount of bytes, except in one situation: whenever some writer is waiting for sending data upon a buffer full condition, while the caller would have to wait for receiving a complete message. This is usually the sign of a pathological use of the `BUF` socket, like defining an incorrect buffer size via [BUF\\_BUFSZ](#). In that case, a short read is allowed to prevent a deadlock.

Returns

In addition to the standard error codes for `recvmsg(2)`, the following specific error code may be returned: none.

Calling context:

RT

5.4.4.8 ssize\_t sendmsg\_\_AF\_RTIPC ( int sockfd, const struct user\_msghdr \* msg, int flags )

Send a message on a RTIPC socket.

## Parameters

<code>in</code>	<code>flags</code>	Operation flags:
-----------------	--------------------	------------------

- `MSG_OOB` Send out-of-band message. For all RTIPC protocols except `IPCPROTO_BUFPP`, sending out-of-band data actually means pushing them to the head of the receiving queue, so that the reader will always receive them before normal messages. `IPCPROTO_BUFPP` does not support out-of-band sending.
- `MSG_DONTWAIT` Non-blocking I/O operation. The caller will not be blocked whenever the message cannot be sent immediately at the time of the call (e.g. memory shortage), but will rather return with `-EWOULDBLOCK`. Unlike other RTIPC protocols, `IPCPROTO_XDDP` accepts but never considers `MSG_DONTWAIT` since writing to a real-time XDDP endpoint is inherently a non-blocking operation.
- `MSG_MORE` Accumulate data before sending. This flag is accepted by the `IPCPROTO_XDDP` protocol only, and tells the send service to accumulate the outgoing data into an internal streaming buffer, instead of issuing a datagram immediately for it. See `XDDP_BUFSZ` for more.

## Note

No RTIPC protocol allows for short writes, and only complete messages are sent to the peer.

## Returns

In addition to the standard error codes for `sendmsg(2)`, the following specific error code may be returned: none.

## Calling context:

RT

5.4.4.9 `int setsockopt__AF_RTIPC ( int sockfd, int level, int optname, const void * optval, socklen_t optlen )`

Set options on RTIPC sockets.

These functions allow to set various socket options. Supported Levels and Options:

- Level `SOL_SOCKET`
- Level `SOL_XDDP`
- Level `SOL_IDDP`
- Level `SOL_BUFPP`

## Returns

In addition to the standard error codes for `setsockopt(2)`, the following specific error code may be returned: follow the option links above.

## Calling context:

non-RT

5.4.4.10 `int socket__AF_RTIPC ( int domain = AF_RTIPC, int type = SOCK_DGRAM, int protocol )`

Create an endpoint for communication in the `AF_RTIPC` domain.

## Parameters

<code>in</code>	<i>protocol</i>	Any of <a href="#">IPCPROTO_XDDP</a> , <a href="#">IPCPROTO_IDDP</a> , or <a href="#">IPCPROTO_BUF</a> . <a href="#">IPCPROTO_IPC</a> is also valid, and refers to the default RTIPC protocol, namely <a href="#">IPCPROTO_IDDP</a> .
-----------------	-----------------	---

## Returns

In addition to the standard error codes for `socket(2)`, the following specific error code may be returned:

- `-ENOPROTOOPT` (Protocol is known, but not compiled in the RTIPC driver). See [RTIPC protocols](#) for available protocols.

## Calling context:

non-RT



## 5.5 Serial Devices

### Files

- file [rtserial.h](#)  
*Real-Time Driver Model for Xenomai, serial device profile header.*

### Data Structures

- struct [rtser\\_config](#)  
*Serial device configuration.*
- struct [rtser\\_status](#)  
*Serial device status.*
- struct [rtser\\_event](#)  
*Additional information about serial device events.*

### Macros

- #define [RTSER\\_RTIOC\\_BREAK\\_CTL\\_IOR](#)(RTIOC\_TYPE\_SERIAL, 0x06, int)  
*Set or clear break on UART output line.*

### RTSER\_DEF\_BAUD

#### Default baud rate

- #define **RTSER\_DEF\_BAUD** 9600

### RTSER\_xxx\_PARITY

#### Number of parity bits

- #define **RTSER\_NO\_PARITY** 0x00
- #define **RTSER\_ODD\_PARITY** 0x01
- #define **RTSER\_EVEN\_PARITY** 0x03
- #define **RTSER\_DEF\_PARITY** RTSER\_NO\_PARITY

### RTSER\_xxx\_BITS

#### Number of data bits

- #define **RTSER\_5\_BITS** 0x00
- #define **RTSER\_6\_BITS** 0x01
- #define **RTSER\_7\_BITS** 0x02
- #define **RTSER\_8\_BITS** 0x03
- #define **RTSER\_DEF\_BITS** RTSER\_8\_BITS

## RTSER\_XXX\_STOPB

Number of stop bits

- `#define RTSER_1_STOPB 0x00`  
*valid only in combination with 5 data bits*
- `#define RTSER_1_5_STOPB 0x01`  
*valid only in combination with 5 data bits*
- `#define RTSER_2_STOPB 0x01`  
*valid only in combination with 5 data bits*
- `#define RTSER_DEF_STOPB RTSER_1_STOPB`  
*valid only in combination with 5 data bits*

## RTSER\_XXX\_HAND

Handshake mechanisms

- `#define RTSER_NO_HAND 0x00`
- `#define RTSER_RTSCTS_HAND 0x01`
- `#define RTSER_DEF_HAND RTSER_NO_HAND`

## RTSER\_RS485\_XXX

RS485 mode with automatic RTS handling

- `#define RTSER_RS485_DISABLE 0x00`
- `#define RTSER_RS485_ENABLE 0x01`
- `#define RTSER_DEF_RS485 RTSER_RS485_DISABLE`

## RTSER\_FIFO\_XXX

Reception FIFO interrupt threshold

- `#define RTSER_FIFO_DEPTH_1 0x00`
- `#define RTSER_FIFO_DEPTH_4 0x40`
- `#define RTSER_FIFO_DEPTH_8 0x80`
- `#define RTSER_FIFO_DEPTH_14 0xC0`
- `#define RTSER_DEF_FIFO_DEPTH RTSER_FIFO_DEPTH_1`

## RTSER\_TIMEOUT\_XXX

Special timeout values, see also [RTDM\\_TIMEOUT\\_XXX](#)

- `#define RTSER_TIMEOUT_INFINITE RTDM_TIMEOUT_INFINITE`
- `#define RTSER_TIMEOUT_NONE RTDM_TIMEOUT_NONE`
- `#define RTSER_DEF_TIMEOUT RTDM_TIMEOUT_INFINITE`

## RTSER\_XXX\_TIMESTAMP\_HISTORY

Timestamp history control

- #define **RTSER\_RX\_TIMESTAMP\_HISTORY** 0x01
- #define **RTSER\_DEF\_TIMESTAMP\_HISTORY** 0x00

## RTSER\_EVENT\_XXX

Events bits

- #define **RTSER\_EVENT\_RXPEND** 0x01
- #define **RTSER\_EVENT\_ERRPEND** 0x02
- #define **RTSER\_EVENT\_MODEMHI** 0x04
- #define **RTSER\_EVENT\_MODEMLO** 0x08
- #define **RTSER\_EVENT\_TXEMPTY** 0x10
- #define **RTSER\_DEF\_EVENT\_MASK** 0x00

## RTSER\_SET\_XXX

Configuration mask bits

- #define **RTSER\_SET\_BAUD** 0x0001
- #define **RTSER\_SET\_PARITY** 0x0002
- #define **RTSER\_SET\_DATA\_BITS** 0x0004
- #define **RTSER\_SET\_STOP\_BITS** 0x0008
- #define **RTSER\_SET\_HANDSHAKE** 0x0010
- #define **RTSER\_SET\_FIFO\_DEPTH** 0x0020
- #define **RTSER\_SET\_TIMEOUT\_RX** 0x0100
- #define **RTSER\_SET\_TIMEOUT\_TX** 0x0200
- #define **RTSER\_SET\_TIMEOUT\_EVENT** 0x0400
- #define **RTSER\_SET\_TIMESTAMP\_HISTORY** 0x0800
- #define **RTSER\_SET\_EVENT\_MASK** 0x1000
- #define **RTSER\_SET\_RS485** 0x2000

## RTSER\_LSR\_XXX

Line status bits

- #define **RTSER\_LSR\_DATA** 0x01
- #define **RTSER\_LSR\_OVERRUN\_ERR** 0x02
- #define **RTSER\_LSR\_PARITY\_ERR** 0x04
- #define **RTSER\_LSR\_FRAMING\_ERR** 0x08
- #define **RTSER\_LSR\_BREAK\_IND** 0x10
- #define **RTSER\_LSR\_THR\_EMPTY** 0x20
- #define **RTSER\_LSR\_TRANSM\_EMPTY** 0x40
- #define **RTSER\_LSR\_FIFO\_ERR** 0x80
- #define **RTSER\_SOFT\_OVERRUN\_ERR** 0x0100

## RTSER\_MSR\_XXX

### Modem status bits

- #define **RTSER\_MSR\_DCTS** 0x01
- #define **RTSER\_MSR\_DDSR** 0x02
- #define **RTSER\_MSR\_TERI** 0x04
- #define **RTSER\_MSR\_DDCD** 0x08
- #define **RTSER\_MSR\_CTS** 0x10
- #define **RTSER\_MSR\_DSR** 0x20
- #define **RTSER\_MSR\_RI** 0x40
- #define **RTSER\_MSR\_DCD** 0x80

## RTSER\_MCR\_XXX

### Modem control bits

- #define **RTSER\_MCR\_DTR** 0x01
- #define **RTSER\_MCR\_RTS** 0x02
- #define **RTSER\_MCR\_OUT1** 0x04
- #define **RTSER\_MCR\_OUT2** 0x08
- #define **RTSER\_MCR\_LOOP** 0x10

## RTSER\_BREAK\_XXX

### Break control

- typedef struct [rtser\\_config](#) [rtser\\_config\\_t](#)  
*Serial device configuration.*
- typedef struct [rtser\\_status](#) [rtser\\_status\\_t](#)  
*Serial device status.*
- typedef struct [rtser\\_event](#) [rtser\\_event\\_t](#)  
*Additional information about serial device events.*
- #define **RTSER\_BREAK\_CLR** 0x00  
*Serial device configuration.*
- #define **RTSER\_BREAK\_SET** 0x01  
*Serial device configuration.*
- #define **RTIOC\_TYPE\_SERIAL** RTDM\_CLASS\_SERIAL  
*Serial device configuration.*

## Sub-Classes of RTDM\_CLASS\_SERIAL

- #define **RTDM\_SUBCLASS\_16550A** 0

## IOCTLs

### Serial device IOCTLs

- #define **RTSER\_RTIOC\_GET\_CONFIG** \_IOR(RTIOC\_TYPE\_SERIAL, 0x00, struct [rtser\\_config](#))  
*Get serial device configuration.*

- #define `RTSER_RTIOC_SET_CONFIG` `_IOW(RTIOC_TYPE_SERIAL, 0x01, struct rtser_config)`  
*Set serial device configuration.*
- #define `RTSER_RTIOC_GET_STATUS` `_IOR(RTIOC_TYPE_SERIAL, 0x02, struct rtser_status)`  
*Get serial device status.*
- #define `RTSER_RTIOC_GET_CONTROL` `_IOR(RTIOC_TYPE_SERIAL, 0x03, int)`  
*Get serial device's modem control register.*
- #define `RTSER_RTIOC_SET_CONTROL` `_IOW(RTIOC_TYPE_SERIAL, 0x04, int)`  
*Set serial device's modem control register.*
- #define `RTSER_RTIOC_WAIT_EVENT` `_IOR(RTIOC_TYPE_SERIAL, 0x05, struct rtser_event)`  
*Wait on serial device events according to previously set mask.*

## 5.5.1 Detailed Description

This is the common interface a RTDM-compliant serial device has to provide. Feel free to comment on this profile via the Xenomai mailing list ([Xenomai@xenomai.org](mailto:Xenomai@xenomai.org)) or directly to the author ([jan.kiszka@web.de](mailto:jan.kiszka@web.de)).

**Profile Revision:** 3

### Device Characteristics

**Device Flags:** `RTDM_NAMED_DEVICE`, `RTDM_EXCLUSIVE`  
**Device Name:** `"rtser<N>"`, `N >= 0`  
**Device Class:** `RTDM_CLASS_SERIAL`

### Supported Operations

#### Open

Environments: non-RT (RT optional, deprecated)

Specific return values: none

#### Close

Environments: non-RT (RT optional, deprecated)

Specific return values: none

#### IOCTL

Mandatory Environments: see [below](#)

Specific return values: see [below](#)

#### Read

Environments: RT (non-RT optional)

Specific return values:

- `-ETIMEDOUT`
- `-EINTR` (interrupted explicitly or by signal)
- `-EAGAIN` (no data available in non-blocking mode)
- `-EBADF` (device has been closed while reading)
- `-EIO` (hardware error or broken bit stream)

#### Write

Environments: RT (non-RT optional)

Specific return values:

- `-ETIMEDOUT`
- `-EINTR` (interrupted explicitly or by signal)
- `-EAGAIN` (no data written in non-blocking mode)
- `-EBADF` (device has been closed while writing)

---

## 5.5.2 Macro Definition Documentation

5.5.2.1 `#define RTSER_RTIOC_BREAK_CTL_IOR(RTIOC_TYPE_SERIAL, 0x06, int)`

Set or clear break on UART output line.

## Parameters

<i>in</i>	<i>arg</i>	RTSER_BREAK_SET or RTSER_BREAK_CLR (int)
-----------	------------	--

## Returns

0 on success, otherwise negative error code

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

## Note

A set break condition may also be cleared on UART line reconfiguration.

Rescheduling: never.

```
5.5.2.2 #define RTSER_RTIOC_GET_CONFIG _IOR(RTIOC_TYPE_SERIAL, 0x00, struct  
      rtser_config)
```

Get serial device configuration.

## Parameters

<i>out</i>	<i>arg</i>	Pointer to configuration buffer (struct <a href="#">rtser_config</a> )
------------	------------	--

## Returns

0 on success, otherwise negative error code

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

```
5.5.2.3 #define RTSER_RTIOC_GET_CONTROL _IOR(RTIOC_TYPE_SERIAL, 0x03, int)
```

Get serial device's modem control register.

## Parameters

out	<i>arg</i>	Pointer to variable receiving the content (int, see <a href="#">RTSER_MCR_xxx</a> )
-----	------------	---

## Returns

0 on success, otherwise negative error code

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.5.2.4 `#define RTSER_RTIOC_GET_STATUS _IOR(RTIOC_TYPE_SERIAL, 0x02, struct rtser_status)`

Get serial device status.

## Parameters

out	<i>arg</i>	Pointer to status buffer (struct <a href="#">rtser_status</a> )
-----	------------	---

## Returns

0 on success, otherwise negative error code

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

## Note

The error states `RTSER_LSR_OVERRUN_ERR`, `RTSER_LSR_PARITY_ERR`, `RTSER_LSR_FRAMING_ERR`, and `RTSER_SOFT_OVERRUN_ERR` that may have occurred during previous read accesses to the device will be saved for being reported via this IOCTL. Upon return from `RTSER_RTIOC_GET_STATUS`, the saved state will be cleared.

Rescheduling: never.

5.5.2.5 `#define RTSER_RTIOC_SET_CONFIG _IOW(RTIOC_TYPE_SERIAL, 0x01, struct rtser_config)`

Set serial device configuration.



## Parameters

<i>in</i>	<i>arg</i>	Pointer to configuration buffer (struct <a href="#">rtser_config</a> )
-----------	------------	--

## Returns

0 on success, otherwise:

- -EPERM is returned if the caller's context is invalid, see note below.
- -ENOMEM is returned if a new history buffer for timestamps cannot be allocated.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

## Note

If [rtser\\_config](#) contains a valid `timestamp_history` and the addressed device has been opened in non-real-time context, this IOCTL must be issued in non-real-time context as well. Otherwise, this command will fail.

Rescheduling: never.

Examples:

[cross-link.c](#).

5.5.2.6 `#define RTSER_RTIOC_SET_CONTROL_IOW(RTIOC_TYPE_SERIAL, 0x04, int)`

Set serial device's modem control register.

## Parameters

<i>in</i>	<i>arg</i>	New control register content (int, see <a href="#">RTSER_MCR_xxx</a> )
-----------	------------	--

## Returns

0 on success, otherwise negative error code

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.5.2.7 `#define RTSER_RTIOC_WAIT_EVENT_IOR(RTIOC_TYPE_SERIAL, 0x05, struct rtser_event)`

Wait on serial device events according to previously set mask.

## Parameters

out	<i>arg</i>	Pointer to event information buffer (struct <a href="#">rtser_event</a> )
-----	------------	---

## Returns

0 on success, otherwise:

- -EBUSY is returned if another task is already waiting on events of this device.
- -EBADF is returned if the file descriptor is invalid or the device has just been closed.

## Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: possible.

## Examples:

[cross-link.c](#).

## 5.6 Testing Devices

### Files

- file [rttesting.h](#)  
*Real-Time Driver Model for Xenomai, testing device profile header.*

### Sub-Classes of RTDM\_CLASS\_TESTING

- #define [RTDM\\_SUBCLASS\\_TIMERBENCH](#) 0  
*subclass name: "timerbench"*
- #define [RTDM\\_SUBCLASS\\_IRQBENCH](#) 1  
*subclass name: "irqbench"*
- #define [RTDM\\_SUBCLASS\\_SWITCHTEST](#) 2  
*subclass name: "switchtest"*
- #define [RTDM\\_SUBCLASS\\_RTDMTEST](#) 3  
*subclass name: "rtdm"*

### IOCTLs

#### Testing device IOCTLs

- #define [RTTST\\_RTIOC\\_INTERM\\_BENCH\\_RES](#) \_IOW(RTIOC\_TYPE\_TESTING, 0x00, struct rttst\_interm\_bench\_res)
- #define [RTTST\\_RTIOC\\_TMBENCH\\_START](#) \_IOW(RTIOC\_TYPE\_TESTING, 0x10, struct rttst\_tmbench\_config)
- #define [RTTST\\_RTIOC\\_TMBENCH\\_STOP](#) \_IOW(RTIOC\_TYPE\_TESTING, 0x11, struct rttst\_overall\_bench\_res)
- #define [RTTST\\_RTIOC\\_IRQBENCH\\_START](#) \_IOW(RTIOC\_TYPE\_TESTING, 0x20, struct rttst\_irqbench\_config)
- #define [RTTST\\_RTIOC\\_IRQBENCH\\_STOP](#) \_IO(RTIOC\_TYPE\_TESTING, 0x21)
- #define [RTTST\\_RTIOC\\_IRQBENCH\\_GET\\_STATS](#) \_IOR(RTIOC\_TYPE\_TESTING, 0x22, struct rttst\_irqbench\_stats)
- #define [RTTST\\_RTIOC\\_IRQBENCH\\_WAIT\\_IRQ](#) \_IO(RTIOC\_TYPE\_TESTING, 0x23)
- #define [RTTST\\_RTIOC\\_IRQBENCH\\_REPLY\\_IRQ](#) \_IO(RTIOC\_TYPE\_TESTING, 0x24)
- #define [RTTST\\_RTIOC\\_SWTEST\\_SET\\_TASKS\\_COUNT](#) \_IOW(RTIOC\_TYPE\_TESTING, 0x30, unsigned long)
- #define [RTTST\\_RTIOC\\_SWTEST\\_SET\\_CPU](#) \_IOW(RTIOC\_TYPE\_TESTING, 0x31, unsigned long)
- #define [RTTST\\_RTIOC\\_SWTEST\\_REGISTER\\_UTASK](#) \_IOW(RTIOC\_TYPE\_TESTING, 0x32, struct rttst\_swtest\_task)
- #define [RTTST\\_RTIOC\\_SWTEST\\_CREATE\\_KTASK](#) \_IOWR(RTIOC\_TYPE\_TESTING, 0x33, struct rttst\_swtest\_task)
- #define [RTTST\\_RTIOC\\_SWTEST\\_PEND](#) \_IOR(RTIOC\_TYPE\_TESTING, 0x34, struct rttst\_swtest\_task)
- #define [RTTST\\_RTIOC\\_SWTEST\\_SWITCH\\_TO](#) \_IOR(RTIOC\_TYPE\_TESTING, 0x35, struct rttst\_swtest\_dir)
- #define [RTTST\\_RTIOC\\_SWTEST\\_GET\\_SWITCHES\\_COUNT](#) \_IOR(RTIOC\_TYPE\_TESTING, 0x36, unsigned long)
- #define [RTTST\\_RTIOC\\_SWTEST\\_GET\\_LAST\\_ERROR](#) \_IOR(RTIOC\_TYPE\_TESTING, 0x37, struct rttst\_swtest\_error)
- #define [RTTST\\_RTIOC\\_SWTEST\\_SET\\_PAUSE](#) \_IOW(RTIOC\_TYPE\_TESTING, 0x38, unsigned long)
- #define [RTTST\\_RTIOC\\_RTDM\\_DEFER\\_CLOSE](#) \_IOW(RTIOC\_TYPE\_TESTING, 0x40, unsigned long)

## 5.6.1 Detailed Description

This group of devices is intended to provide in-kernel testing results. Feel free to comment on this profile via the Xenomai mailing list ([xenomai@xenomai.org](mailto:xenomai@xenomai.org)) or directly to the author ([jan.kiszka@web.de](mailto:jan.kiszka@web.de)).

**Profile Revision:** 2

### Device Characteristics

**Device Flags:** RTDM\_NAMED\_DEVICE

**Device Name:** "rtttest[-<subclass>]<N>", N >= 0, optional subclass name to simplify device discovery

**Device Class:** RTDM\_CLASS\_TESTING

### Supported Operations

#### **Open**

Environments: non-RT (RT optional, deprecated)

Specific return values: none

#### **Close**

Environments: non-RT (RT optional, deprecated)

Specific return values: none

#### **IOCTL**

Mandatory Environments: see [TSTIOCTLs](#) below

Specific return values: see [TSTIOCTLs](#) below

## 5.7 Inter-Driver API

### Functions

- struct `rtdm_dev_context` \* `rtdm_context_get` (int fd)  
*Retrieve and lock a device context.*
- int `rtdm_select_bind` (int fd, `rtdm_selector_t` \*selector, enum `rtdm_selecttype` type, unsigned fd\_index)  
*Bind a selector to specified event types of a given file descriptor.*
- void `rtdm_context_lock` (struct `rtdm_dev_context` \*context)  
*Increment context reference counter.*
- void `rtdm_context_unlock` (struct `rtdm_dev_context` \*context)  
*Decrement context reference counter.*
- void `rtdm_context_put` (struct `rtdm_dev_context` \*context)  
*Release a device context obtained via `rtdm_context_get()`*
- int `rtdm_open` (const char \*path, int oflag,...)  
*Open a device.*
- int `rtdm_socket` (int protocol\_family, int socket\_type, int protocol)  
*Create a socket.*
- int `rtdm_close` (int fd)  
*Close a device or socket.*
- int `rtdm_ioctl` (int fd, int request,...)  
*Issue an IOCTL.*
- `ssize_t` `rtdm_read` (int fd, void \*buf, `size_t` nbyte)  
*Read from device.*
- `ssize_t` `rtdm_write` (int fd, const void \*buf, `size_t` nbyte)  
*Write to device.*
- `ssize_t` `rtdm_recvmsg` (int fd, struct `user_msghdr` \*msg, int flags)  
*Receive message from socket.*
- `ssize_t` `rtdm_recvfrom` (int fd, void \*buf, `size_t` len, int flags, struct `sockaddr` \*from, `socklen_t` \*fromlen)  
*Receive message from socket.*
- `ssize_t` `rtdm_recv` (int fd, void \*buf, `size_t` len, int flags)  
*Receive message from socket.*
- `ssize_t` `rtdm_sendmsg` (int fd, const struct `user_msghdr` \*msg, int flags)  
*Transmit message to socket.*
- `ssize_t` `rtdm_sendto` (int fd, const void \*buf, `size_t` len, int flags, const struct `sockaddr` \*to, `socklen_t` tolen)  
*Transmit message to socket.*
- `ssize_t` `rtdm_send` (int fd, const void \*buf, `size_t` len, int flags)  
*Transmit message to socket.*
- int `rtdm_bind` (int fd, const struct `sockaddr` \*my\_addr, `socklen_t` addrlen)  
*Bind to local address.*
- int `rtdm_connect` (int fd, const struct `sockaddr` \*serv\_addr, `socklen_t` addrlen)  
*Connect to remote address.*
- int `rtdm_listen` (int fd, int backlog)  
*Listen for incoming connection requests.*
- int `rtdm_accept` (int fd, struct `sockaddr` \*addr, `socklen_t` \*addrlen)  
*Accept a connection requests.*
- int `rtdm_shutdown` (int fd, int how)  
*Shut down parts of a connection.*

- int [rtm\\_getsockopt](#) (int fd, int level, int optname, void \*optval, socklen\_t \*optlen)  
*Get socket option.*
- int [rtm\\_setsockopt](#) (int fd, int level, int optname, const void \*optval, socklen\_t optlen)  
*Set socket option.*
- int [rtm\\_getsockname](#) (int fd, struct sockaddr \*name, socklen\_t \*namelen)  
*Get local socket address.*
- int [rtm\\_getpeername](#) (int fd, struct sockaddr \*name, socklen\_t \*namelen)  
*Get socket destination address.*

## 5.7.1 Detailed Description

## 5.7.2 Function Documentation

### 5.7.2.1 int rtdm\_accept ( int fd, struct sockaddr \* addr, socklen\_t \* addrlen )

Accept a connection requests.

Refer to [rt\\_dev\\_accept\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

### 5.7.2.2 int rtdm\_bind ( int fd, const struct sockaddr \* my\_addr, socklen\_t addrlen )

Bind to local address.

Refer to [rt\\_dev\\_bind\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

### 5.7.2.3 int rtdm\_close ( int fd )

Close a device or socket.

Refer to [rt\\_dev\\_close\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

### 5.7.2.4 int rtdm\_connect ( int fd, const struct sockaddr \* serv\_addr, socklen\_t addrlen )

Connect to remote address.

Refer to [rt\\_dev\\_connect\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

---

5.7.2.5 struct **rtm\_dev\_context**\* rtdm\_context\_get ( int fd )

Retrieve and lock a device context.

## Parameters

<code>in</code>	<code>fd</code>	File descriptor
-----------------	-----------------	-----------------

## Returns

Pointer to associated device context, or NULL on error

## Note

The device context has to be unlocked using [rtm\\_context\\_put\(\)](#) when it is no longer referenced.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

References `rtm_dev_context::close_lock_count`, and `rtm_dev_context::fd`.

Referenced by `rtm_select_bind()`.

5.7.2.6 `void rtm_context_lock ( struct rtm_dev_context * context )`

Increment context reference counter.

## Parameters

<code>in</code>	<code>context</code>	Device context
-----------------	----------------------	----------------

## Note

[rtm\\_context\\_get\(\)](#) automatically increments the lock counter. You only need to call this function in special scenarios, e.g. when keeping additional references to the context structure that have different lifetimes. Only use [rtm\\_context\\_lock\(\)](#) on contexts that are currently locked via an earlier [rtm\\_context\\_get\(\)/rtm\\_context\\_lock\(\)](#) or while running a device operation handler.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.7.2.7 `void rtm_context_put ( struct rtm_dev_context * context )`

Release a device context obtained via [rtm\\_context\\_get\(\)](#)



## Parameters

<code>in</code>	<code>context</code>	Device context
-----------------	----------------------	----------------

## Note

Every successful call to [rtm\\_context\\_get\(\)](#) must be matched by a [rtm\\_context\\_put\(\)](#) invocation.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.7.2.8 void [rtm\\_context\\_unlock](#) ( struct **rtm\_dev\_context** \* context )

Decrement context reference counter.

## Parameters

<code>in</code>	<code>context</code>	Device context
-----------------	----------------------	----------------

## Note

Every call to [rtm\\_context\\_locked\(\)](#) must be matched by a [rtm\\_context\\_unlock\(\)](#) invocation.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

Referenced by [rtm\\_select\\_bind\(\)](#).

5.7.2.9 int [rtm\\_getpeername](#) ( int fd, struct sockaddr \* name, socklen\_t \* namelen )

Get socket destination address.

Refer to [rt\\_dev\\_getpeername\(\)](#) for parameters and return values

## Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.10 `int rtdm_getsockname ( int fd, struct sockaddr * name, socklen_t * namelen )`

Get local socket address.

Refer to [rt\\_dev\\_getsockname\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.11 `int rtdm_getsockopt ( int fd, int level, int optname, void * optval, socklen_t * optlen )`

Get socket option.

Refer to [rt\\_dev\\_getsockopt\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.12 `int rtdm_ioctl ( int fd, int request, ... )`

Issue an IOCTL.

Refer to [rt\\_dev\\_ioctl\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.13 `int rtdm_listen ( int fd, int backlog )`

Listen for incoming connection requests.

Refer to [rt\\_dev\\_listen\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.14 `int rtdm_open ( const char * path, int oflag, ... )`

Open a device.

Refer to [rt\\_dev\\_open\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.15 `ssize_t rtdm_read ( int fd, void * buf, size_t nbyte )`

Read from device.

Refer to `rt_dev_read()` for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.16 `ssize_t rtdm_rcv ( int fd, void * buf, size_t len, int flags )`

Receive message from socket.

Refer to `rt_dev_rcv()` for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.17 `ssize_t rtdm_rcvfrom ( int fd, void * buf, size_t len, int flags, struct sockaddr * from, socklen_t * fromlen )`

Receive message from socket.

Refer to `rt_dev_rcvfrom()` for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.18 `ssize_t rtdm_rcvmsg ( int fd, struct user_msghdr * msg, int flags )`

Receive message from socket.

Refer to `rt_dev_rcvmsg()` for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.19 `int rtdm_select_bind ( int fd, rtdm_selector_t * selector, enum rtdm_selecttype type, unsigned fd_index )`

Bind a selector to specified event types of a given file descriptor.

This function is invoked by higher RTOS layers implementing select-like services. It shall not be called directly by RTDM drivers.

Parameters

<code>in</code>	<code>fd</code>	File descriptor to bind to
<code>in,out</code>	<code>selector</code>	Selector object that shall be bound to the given event
<code>in</code>	<code>type</code>	Event type the caller is interested in
<code>in</code>	<code>fd_index</code>	Index in the file descriptor set of the caller

Returns

0 on success, otherwise:

- -EBADF is returned if the file descriptor *fd* cannot be resolved.
- -EINVAL is returned if *type* or *fd\_index* are invalid.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

References `rtdm_dev_context::ops`, `rtdm_context_get()`, `rtdm_context_unlock()`, and `rtdm_operations::select_bind`.

5.7.2.20 `ssize_t rtdm_send ( int fd, const void * buf, size_t len, int flags )`

Transmit message to socket.

Refer to [rt\\_dev\\_send\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.21 `ssize_t rtdm_sendmsg ( int fd, const struct user_msghdr * msg, int flags )`

Transmit message to socket.

Refer to [rt\\_dev\\_sendmsg\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.22 `ssize_t rtdm_sendto ( int fd, const void * buf, size_t len, int flags, const struct sockaddr * to, socklen_t tolen )`

Transmit message to socket.

Refer to [rt\\_dev\\_sendto\(\)](#) for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.23 `int rtdm_setsockopt ( int fd, int level, int optname, const void * optval, socklen_t optlen )`

Set socket option.

Refer to `rt_dev_setsockopt()` for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.24 `int rtdm_shutdown ( int fd, int how )`

Shut down parts of a connection.

Refer to `rt_dev_shutdown()` for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.25 `int rtdm_socket ( int protocol_family, int socket_type, int protocol )`

Create a socket.

Refer to `rt_dev_socket()` for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

5.7.2.26 `ssize_t rtdm_write ( int fd, const void * buf, size_t nbyte )`

Write to device.

Refer to `rt_dev_write()` for parameters and return values

Environments:

Depends on driver implementation, see [Device Profiles](#).

Rescheduling: possible.

## 5.8 Device Registration Services

### Modules

- [Synchronisation Services](#)

### Data Structures

- struct [rtdm\\_operations](#)  
*Device operations.*
- struct [rtdm\\_dev\\_context](#)  
*Device context.*
- struct [rtdm\\_device](#)  
*RTDM device.*

### Functions

- static void \* [rtdm\\_context\\_to\\_private](#) (struct [rtdm\\_dev\\_context](#) \*context)  
*Locate the driver private area associated to a device context structure.*
- static struct [rtdm\\_dev\\_context](#) \* [rtdm\\_private\\_to\\_context](#) (void \*dev\_private)  
*Locate a device context structure from its driver private area.*
- int [rtdm\\_dev\\_register](#) (struct [rtdm\\_device](#) \*device)  
*Register a RTDM device.*
- int [rtdm\\_dev\\_unregister](#) (struct [rtdm\\_device](#) \*device, unsigned int poll\_delay)  
*Unregisters a RTDM device.*

### Device Flags

Static flags describing a RTDM device

- #define [RTDM\\_EXCLUSIVE](#) 0x0001  
*If set, only a single instance of the device can be requested by an application.*
- #define [RTDM\\_NAMED\\_DEVICE](#) 0x0010  
*If set, the device is addressed via a clear-text name.*
- #define [RTDM\\_PROTOCOL\\_DEVICE](#) 0x0020  
*If set, the device is addressed via a combination of protocol ID and socket type.*
- #define [RTDM\\_DEVICE\\_TYPE\\_MASK](#) 0x00F0  
*Mask selecting the device type.*

### Context Flags

Dynamic flags describing the state of an open RTDM device (bit numbers)

- #define [RTDM\\_CREATED\\_IN\\_NRT](#) 0  
*Set by RTDM if the device instance was created in non-real-time context.*
- #define [RTDM\\_CLOSING](#) 1  
*Set by RTDM when the device is being closed.*
- #define [RTDM\\_USER\\_CONTEXT\\_FLAG](#) 8 /\* first user-definable flag \*/  
*Lowest bit number the driver developer can use freely.*

## Driver Versioning

Current revisions of RTDM structures, encoding of driver versions. See [API Versioning](#) for the interface revision.

- #define `RTDM_DEVICE_STRUCT_VER` 5  
*Version of struct `rtdm_device`.*
- #define `RTDM_CONTEXT_STRUCT_VER` 3  
*Version of struct `rtdm_dev_context`.*
- #define `RTDM_SECURE_DEVICE` 0x80000000  
*Flag indicating a secure variant of RTDM (not supported here)*
- #define `RTDM_DRIVER_VER`(major, minor, patch) (((major & 0xFF) << 16) | ((minor & 0xFF) << 8) | (patch & 0xFF))  
*Version code constructor for driver revisions.*
- #define `RTDM_DRIVER_MAJOR_VER`(ver) (((ver) >> 16) & 0xFF)  
*Get major version number from driver revision code.*
- #define `RTDM_DRIVER_MINOR_VER`(ver) (((ver) >> 8) & 0xFF)  
*Get minor version number from driver revision code.*
- #define `RTDM_DRIVER_PATCH_VER`(ver) ((ver) & 0xFF)  
*Get patch version number from driver revision code.*

## Operation Handler Prototypes

- typedef int(\* `rtdm_open_handler_t` )(struct `rtdm_dev_context` \*context, `rtdm_user_info_t` \*user\_info, int oflag)  
*Named device open handler.*
- typedef int(\* `rtdm_socket_handler_t` )(struct `rtdm_dev_context` \*context, `rtdm_user_info_t` \*user\_info, int protocol)  
*Socket creation handler for protocol devices.*
- typedef int(\* `rtdm_close_handler_t` )(struct `rtdm_dev_context` \*context, `rtdm_user_info_t` \*user\_info)  
*Close handler.*
- typedef int(\* `rtdm_ioctl_handler_t` )(struct `rtdm_dev_context` \*context, `rtdm_user_info_t` \*user\_info, unsigned int request, void \_\_user \*arg)  
*IOCTL handler.*
- typedef int(\* `rtdm_select_bind_handler_t` )(struct `rtdm_dev_context` \*context, `rtdm_selector_t` \*selector, enum `rtdm_selecttype` type, unsigned fd\_index)  
*Select binding handler.*
- typedef ssize\_t(\* `rtdm_read_handler_t` )(struct `rtdm_dev_context` \*context, `rtdm_user_info_t` \*user\_info, void \*buf, size\_t nbyte)  
*Read handler.*
- typedef ssize\_t(\* `rtdm_write_handler_t` )(struct `rtdm_dev_context` \*context, `rtdm_user_info_t` \*user\_info, const void \*buf, size\_t nbyte)  
*Write handler.*
- typedef ssize\_t(\* `rtdm_recvmsg_handler_t` )(struct `rtdm_dev_context` \*context, `rtdm_user_info_t` \*user\_info, struct `user_msghdr` \*msg, int flags)  
*Receive message handler.*
- typedef ssize\_t(\* `rtdm_sendmsg_handler_t` )(struct `rtdm_dev_context` \*context, `rtdm_user_info_t` \*user\_info, const struct `user_msghdr` \*msg, int flags)  
*Transmit message handler.*

## 5.8.1 Detailed Description

### 5.8.2 Macro Definition Documentation

#### 5.8.2.1 `#define RTDM_CLOSING 1`

Set by RTDM when the device is being closed.

#### 5.8.2.2 `#define RTDM_CREATED_IN_NRT 0`

Set by RTDM if the device instance was created in non-real-time context.

#### 5.8.2.3 `#define RTDM_DEVICE_TYPE_MASK 0x00F0`

Mask selecting the device type.

Referenced by `rtm_dev_register()`, and `rtm_dev_unregister()`.

#### 5.8.2.4 `#define RTDM_EXCLUSIVE 0x0001`

If set, only a single instance of the device can be requested by an application.

Referenced by `rtm_dev_register()`.

#### 5.8.2.5 `#define RTDM_NAMED_DEVICE 0x0010`

If set, the device is addressed via a clear-text name.

Referenced by `rtm_dev_register()`, and `rtm_dev_unregister()`.

#### 5.8.2.6 `#define RTDM_PROTOCOL_DEVICE 0x0020`

If set, the device is addressed via a combination of protocol ID and socket type.

Referenced by `rtm_dev_register()`.

## 5.8.3 Typedef Documentation

#### 5.8.3.1 `typedef int(* rtdm_close_handler_t)(struct rtdm_dev_context *context, rtdm_user_info_t *user_info)`

Close handler.

Parameters

<code>in</code>	<code>context</code>	Context structure associated with opened device instance
<code>in</code>	<code>user_info</code>	Opaque pointer to information about user mode caller, NULL if kernel mode or deferred user mode call

Returns

0 on success. On failure return either `-ENOSYS`, to request that this handler be called again from the opposite realtime/non-realtime context, `-EAGAIN` to request a recall after a grace period, or a valid negative error code according to IEEE Std 1003.1.



## Note

Drivers must be prepared for that case that the close handler is invoked more than once per open context (even if the handler already completed an earlier run successfully). The driver has to avoid releasing resources twice as well as returning false errors on successive close invocations.

## See Also

`close()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.8.3.2 `typedef int(* rtdm_ioctl_handler_t)(struct rtdm_dev_context *context, rtdm_user_info_t *user_info, unsigned int request, void __user *arg)`

## IOCTL handler.

### Parameters

<code>in</code>	<code>context</code>	Context structure associated with opened device instance
<code>in</code>	<code>user_info</code>	Opaque pointer to information about user mode caller, NULL if kernel mode call
<code>in</code>	<code>request</code>	Request number as passed by the user
<code>in,out</code>	<code>arg</code>	Request argument as passed by the user

## Returns

A positive value or 0 on success. On failure return either `-ENOSYS`, to request that the function be called again from the opposite realtime/non-realtime context, or another negative error code.

## See Also

`ioctl()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.8.3.3 `typedef int(* rtdm_open_handler_t)(struct rtdm_dev_context *context, rtdm_user_info_t *user_info, int oflag)`

## Named device open handler.

### Parameters

<code>in</code>	<code>context</code>	Context structure associated with opened device instance
<code>in</code>	<code>user_info</code>	Opaque pointer to information about user mode caller, NULL if kernel mode call
<code>in</code>	<code>oflag</code>	Open flags as passed by the user

## Returns

0 on success. On failure return either `-ENOSYS`, to request that this handler be called again from the opposite realtime/non-realtime context, or another negative error code.

## See Also

`open()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.8.3.4 `typedef ssize_t(* rtdm_read_handler_t)(struct rtdm_dev_context *context, rtdm_user_info_t *user_info, void *buf, size_t nbytes)`

## Read handler.

## Parameters

in	<i>context</i>	Context structure associated with opened device instance
in	<i>user_info</i>	Opaque pointer to information about user mode caller, NULL if kernel mode call
out	<i>buf</i>	Input buffer as passed by the user
in	<i>nbyte</i>	Number of bytes the user requests to read

## Returns

On success, the number of bytes read. On failure return either -ENOSYS, to request that this handler be called again from the opposite realtime/non-realtime context, or another negative error code.

## See Also

read() in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.8.3.5 typedef ssize\_t(\* rtdm\_recvmsg\_handler\_t)(struct **rtdm\_dev\_context** \*context, rtdm\_user\_info\_t \*user\_info, struct user\_msghdr \*msg, int flags)

Receive message handler.

## Parameters

in	<i>context</i>	Context structure associated with opened device instance
in	<i>user_info</i>	Opaque pointer to information about user mode caller, NULL if kernel mode call
in,out	<i>msg</i>	Message descriptor as passed by the user, automatically mirrored to safe kernel memory in case of user mode call
in	<i>flags</i>	Message flags as passed by the user

## Returns

On success, the number of bytes received. On failure return either -ENOSYS, to request that this handler be called again from the opposite realtime/non-realtime context, or another negative error code.

## See Also

recvmsg() in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.8.3.6 typedef int(\* rtdm\_select\_bind\_handler\_t)(struct **rtdm\_dev\_context** \*context, rtdm\_selector\_t \*selector, enum **rtdm\_selecttype** type, unsigned fd\_index)

Select binding handler.

## Parameters

in	<i>context</i>	Context structure associated with opened device instance
in,out	<i>selector</i>	Object that shall be bound to the given event

<i>in</i>	<i>type</i>	Event type the selector is interested in
<i>in</i>	<i>fd_index</i>	Opaque value, to be passed to <code>rtdm_event_select_bind</code> or <code>rtdm_semaphore_select_bind</code> unmodified

#### Returns

0 on success. On failure return either `-ENOSYS`, to request that this handler be called again from the opposite realtime/non-realtime context, or another negative error code.

5.8.3.7 `typedef ssize_t(* rtdm_sendmsg_handler_t)(struct rtdm_dev_context *context, rtdm_user_info_t *user_info, const struct user_msghdr *msg, int flags)`

Transmit message handler.

#### Parameters

<i>in</i>	<i>context</i>	Context structure associated with opened device instance
<i>in</i>	<i>user_info</i>	Opaque pointer to information about user mode caller, NULL if kernel mode call
<i>in</i>	<i>msg</i>	Message descriptor as passed by the user, automatically mirrored to safe kernel memory in case of user mode call
<i>in</i>	<i>flags</i>	Message flags as passed by the user

#### Returns

On success, the number of bytes transmitted. On failure return either `-ENOSYS`, to request that this handler be called again from the opposite realtime/non-realtime context, or another negative error code.

#### See Also

`sendmsg()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

5.8.3.8 `typedef int(* rtdm_socket_handler_t)(struct rtdm_dev_context *context, rtdm_user_info_t *user_info, int protocol)`

Socket creation handler for protocol devices.

#### Parameters

<i>in</i>	<i>context</i>	Context structure associated with opened device instance
<i>in</i>	<i>user_info</i>	Opaque pointer to information about user mode caller, NULL if kernel mode call
<i>in</i>	<i>protocol</i>	Protocol number as passed by the user

#### Returns

0 on success. On failure return either `-ENOSYS`, to request that this handler be called again from the opposite realtime/non-realtime context, or another negative error code.

#### See Also

`socket()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

---

5.8.3.9 `typedef ssize_t(* rtdm_write_handler_t)(struct rtdm_dev_context *context, rtdm_user_info_t *user_info, const void *buf, size_t nbyte)`

Write handler.

## Parameters

<code>in</code>	<code>context</code>	Context structure associated with opened device instance
<code>in</code>	<code>user_info</code>	Opaque pointer to information about user mode caller, NULL if kernel mode call
<code>in</code>	<code>buf</code>	Output buffer as passed by the user
<code>in</code>	<code>nbyte</code>	Number of bytes the user requests to write

## Returns

On success, the number of bytes written. On failure return either `-ENOSYS`, to request that this handler be called again from the opposite realtime/non-realtime context, or another negative error code.

## See Also

`write()` in IEEE Std 1003.1, <http://www.opengroup.org/onlinepubs/009695399>

## 5.8.4 Function Documentation

5.8.4.1 `static void* rtdm_context_to_private ( struct rtdm_dev_context * context ) [inline], [static]`

Locate the driver private area associated to a device context structure.

### Parameters

<code>in</code>	<code>context</code>	Context structure associated with opened device instance
-----------------	----------------------	--

### Returns

The address of the private driver area associated to `context`.

References `rtdm_dev_context::dev_private`.

5.8.4.2 `int rtdm_dev_register ( struct rtdm_device * device )`

Register a RTDM device.

### Parameters

<code>in</code>	<code>device</code>	Pointer to structure describing the new device.
-----------------	---------------------	---

### Returns

0 is returned upon success. Otherwise:

- `-EINVAL` is returned if the device structure contains invalid entries. Check kernel log in this case.
- `-ENOMEM` is returned if the context for an exclusive device cannot be allocated.
- `-EEXIST` is returned if the specified device name or protocol ID is already in use.
- `-EAGAIN` is returned if some `/proc` entry cannot be created.

### Environments:

This service can be called from:

- Kernel module initialization/cleanup code

Rescheduling: never.

References `rtm_operations::close_nrt`, `rtm_operations::close_rt`, `rtm_device::context_size`, `rtm_device::device_class`, `rtm_device::device_flags`, `rtm_device::device_name`, `rtm_device::device_sub_class`, `rtm_device::driver_version`, `rtm_device::open_rt`, `rtm_device::ops`, `rtm_device::proc_name`, `rtm_device::profile_version`, `rtm_device::protocol_family`, `rtm_device::reserved`, `RTDM_DEVICE_STRUCT_VER`, `RTDM_DEVICE_TYPE_MASK`, `RTDM_EXCLUSIVE`, `RTDM_NAMED_DEVICE`, `RTDM_PROTOCOL_DEVICE`, `rtm_operations::select_bind`, `rtm_device::socket_rt`, `rtm_device::socket_type`, and `rtm_device::struct_version`.

5.8.4.3 `int rtdm_dev_unregister ( struct rtdm_device * device, unsigned int poll_delay )`

Unregisters a RTDM device.

Parameters

<code>in</code>	<code>device</code>	Pointer to structure describing the device to be unregistered.
<code>in</code>	<code>poll_delay</code>	Polling delay in milliseconds to check repeatedly for open instances of <code>device</code> , or 0 for non-blocking mode.

Returns

0 is returned upon success. Otherwise:

- `-ENODEV` is returned if the device was not registered.
- `-EAGAIN` is returned if the device is busy with open instances and 0 has been passed for `poll_delay`.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code

Rescheduling: never.

References `rtm_device::device_flags`, `rtm_device::device_name`, `rtm_device::protocol_family`, `rtm_device::reserved`, `RTDM_DEVICE_TYPE_MASK`, `RTDM_NAMED_DEVICE`, and `rtm_device::socket_type`.

5.8.4.4 `static struct rtdm_dev_context* rtdm_private_to_context ( void * dev_private ) [static]`

Locate a device context structure from its driver private area.

Parameters

<code>in</code>	<code>dev_private</code>	Address of a private context area
-----------------	--------------------------	-----------------------------------

Returns

The address of the device context structure defining `dev_private`.

## 5.9 Driver Development API

### Modules

- [Inter-Driver API](#)
- [Device Registration Services](#)
- [Clock Services](#)
- [Task Services](#)
- [Timer Services](#)
- [Synchronisation Services](#)
- [Interrupt Management Services](#)
- [Non-Real-Time Signalling Services](#)
- [Utility Services](#)

### Files

- file [rtdm\\_driver.h](#)  
*Real-Time Driver Model for Xenomai, driver API header.*

#### 5.9.1 Detailed Description

This is the lower interface of RTDM provided to device drivers, currently limited to kernel-space. Real-time drivers should only use functions of this interface in order to remain portable.

## 5.10 Clock Services

### Functions

- [nanosecs\\_abs\\_t rtdm\\_clock\\_read](#) (void)  
*Get system time.*
- [nanosecs\\_abs\\_t rtdm\\_clock\\_read\\_monotonic](#) (void)  
*Get monotonic time.*

#### 5.10.1 Detailed Description

#### 5.10.2 Function Documentation

##### 5.10.2.1 [nanosecs\\_abs\\_t rtdm\\_clock\\_read](#) ( void )

Get system time.

#### Returns

The system time in nanoseconds is returned

#### Note

The resolution of this service depends on the system timer. In particular, if the system timer is running in periodic mode, the return value will be limited to multiples of the timer tick period.

The system timer may have to be started to obtain valid results. Whether this happens automatically (as on Xenomai) or is controlled by the application depends on the RTDM host environment.

#### Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

Referenced by [rtdm\\_ratelimit\(\)](#).

##### 5.10.2.2 [nanosecs\\_abs\\_t rtdm\\_clock\\_read\\_monotonic](#) ( void )

Get monotonic time.

#### Returns

The monotonic time in nanoseconds is returned



## Note

The resolution of this service depends on the system timer. In particular, if the system timer is running in periodic mode, the return value will be limited to multiples of the timer tick period.

The system timer may have to be started to obtain valid results. Whether this happens automatically (as on Xenomai) or is controlled by the application depends on the RTDM host environment.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

## 5.11 Task Services

### Typedefs

- typedef void([\\* rtdm\\_task\\_proc\\_t](#))(void \*arg)  
*Real-time task procedure.*

### Functions

- int [rtdm\\_task\\_init](#) (rtdm\_task\_t \*task, const char \*name, [rtdm\\_task\\_proc\\_t](#) task\_proc, void \*arg, int priority, [nanosecs\\_rel\\_t](#) period)  
*Initialise and start a real-time task.*
- void [rtdm\\_task\\_destroy](#) (rtdm\_task\_t \*task)  
*Destroy a real-time task.*
- void [rtdm\\_task\\_set\\_priority](#) (rtdm\_task\_t \*task, int priority)  
*Adjust real-time task priority.*
- int [rtdm\\_task\\_set\\_period](#) (rtdm\_task\_t \*task, [nanosecs\\_rel\\_t](#) period)  
*Adjust real-time task period.*
- int [rtdm\\_task\\_wait\\_period](#) (void)  
*Wait on next real-time task period.*
- int [rtdm\\_task\\_unblock](#) (rtdm\_task\_t \*task)  
*Activate a blocked real-time task.*
- rtdm\_task\_t \* [rtdm\\_task\\_current](#) (void)  
*Get current real-time task.*
- int [rtdm\\_task\\_sleep](#) ([nanosecs\\_rel\\_t](#) delay)  
*Sleep a specified amount of time.*
- int [rtdm\\_task\\_sleep\\_until](#) ([nanosecs\\_abs\\_t](#) wakeup\_time)  
*Sleep until a specified absolute time.*
- int [rtdm\\_task\\_sleep\\_abs](#) ([nanosecs\\_abs\\_t](#) wakeup\_time, enum [rtdm\\_timer\\_mode](#) mode)  
*Sleep until a specified absolute time.*
- void [rtdm\\_task\\_join\\_nrt](#) (rtdm\_task\_t \*task, unsigned int poll\_delay)  
*Wait on a real-time task to terminate.*
- void [rtdm\\_task\\_busy\\_sleep](#) ([nanosecs\\_rel\\_t](#) delay)  
*Busy-wait a specified amount of time.*

### Task Priority Range

Maximum and minimum task priorities

- #define **RTDM\_TASK\_LOWEST\_PRIORITY** XNSCHED\_LOW\_PRIO
- #define **RTDM\_TASK\_HIGHEST\_PRIORITY** XNSCHED\_HIGH\_PRIO

### Task Priority Modification

Raise or lower task priorities by one level

- #define **RTDM\_TASK\_RAISE\_PRIORITY** (+1)
- #define **RTDM\_TASK\_LOWER\_PRIORITY** (-1)

### 5.11.1 Detailed Description

### 5.11.2 Typedef Documentation

#### 5.11.2.1 typedef void(\* rtdm\_task\_proc\_t)(void \*arg)

Real-time task procedure.

Parameters

<i>in, out</i>	<i>arg</i>	argument as passed to <a href="#">rtdm_task_init()</a>
----------------	------------	--

### 5.11.3 Function Documentation

#### 5.11.3.1 void rtdm\_task\_busy\_sleep ( **nanosecs\_rel\_t** delay )

Busy-wait a specified amount of time.

Parameters

<i>in</i>	<i>delay</i>	Delay in nanoseconds. Note that a zero delay does <b>not</b> have the meaning of <code>RTDM_TIMEOUT_INFINITE</code> here.
-----------	--------------	---

Note

The caller must not be migratable to different CPUs while executing this service. Otherwise, the actual delay will be undefined.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine (should be avoided or kept short)
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never (except due to external interruptions).

#### 5.11.3.2 rtdm\_task\_t\* rtdm\_task\_current ( void )

Get current real-time task.

Returns

Pointer to task handle

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

---

5.11.3.3 void rtdm\_task\_destroy ( rtdm\_task\_t \* task )

Destroy a real-time task.

## Parameters

<i>in,out</i>	<i>task</i>	Task handle as returned by <a href="#">rtdm_task_init()</a>
---------------	-------------	---

## Note

Passing the same task handle to RTDM services after the completion of this function is not allowed.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

```
5.11.3.4 int rtdm_task_init ( rtdm_task_t * task, const char * name, rtdm_task_proc_t task_proc,
                             void * arg, int priority, nanosecs_rel_t period )
```

Initialise and start a real-time task.

After initialising a task, the task handle remains valid and can be passed to RTDM services until either [rtdm\\_task\\_destroy\(\)](#) or [rtdm\\_task\\_join\\_nrt\(\)](#) was invoked.

## Parameters

<i>in,out</i>	<i>task</i>	Task handle
<i>in</i>	<i>name</i>	Optional task name
<i>in</i>	<i>task_proc</i>	Procedure to be executed by the task
<i>in</i>	<i>arg</i>	Custom argument passed to <code>task_proc()</code> on entry
<i>in</i>	<i>priority</i>	Priority of the task, see also <a href="#">Task Priority Range</a>
<i>in</i>	<i>period</i>	Period in nanoseconds of a cyclic task, 0 for non-cyclic mode

## Returns

0 on success, otherwise negative error code

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

```
5.11.3.5 void rtdm_task_join_nrt ( rtdm_task_t * task, unsigned int poll_delay )
```

Wait on a real-time task to terminate.

## Parameters

<i>in, out</i>	<i>task</i>	Task handle as returned by <a href="#">rtdm_task_init()</a>
<i>in</i>	<i>poll_delay</i>	Delay in milliseconds between periodic tests for the state of the real-time task. This parameter is ignored if the termination is internally realised without polling.

## Note

Passing the same task handle to RTDM services after the completion of this function is not allowed. This service does not trigger the termination of the targeted task. The user has to take of this, otherwise [rtdm\\_task\\_join\\_nrt\(\)](#) will never return.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- User-space task (non-RT)

Rescheduling: possible.

5.11.3.6 `int rtdm_task_set_period ( rtdm_task_t * task, nanosecs_rel_t period )`

Adjust real-time task period.

## Parameters

<i>in, out</i>	<i>task</i>	Task handle as returned by <a href="#">rtdm_task_init()</a>
<i>in</i>	<i>period</i>	New period in nanoseconds of a cyclic task, 0 for non-cyclic mode

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.11.3.7 `void rtdm_task_set_priority ( rtdm_task_t * task, int priority )`

Adjust real-time task priority.

## Parameters

<i>in, out</i>	<i>task</i>	Task handle as returned by <a href="#">rtdm_task_init()</a>
<i>in</i>	<i>priority</i>	New priority of the task, see also <a href="#">Task Priority Range</a>

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine

- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.11.3.8 `int rtdm_task_sleep ( nanosecs_rel_t delay )`

Sleep a specified amount of time.

Parameters

<code>in</code>	<code>delay</code>	Delay in nanoseconds, see <a href="#">RTDM_TIMEOUT_XXX</a> for special values.
-----------------	--------------------	--

Returns

0 on success, otherwise:

- -EINTR is returned if calling task has been unblock by a signal or explicitly via [rtdm\\_task\\_unblock\(\)](#).
- -EPERM *may* be returned if an illegal invocation environment is detected.

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: always.

5.11.3.9 `int rtdm_task_sleep_abs ( nanosecs_abs_t wakeup_time, enum rtdm_timer_mode mode )`

Sleep until a specified absolute time.

Parameters

<code>in</code>	<code>wakeup_time</code>	Absolute timeout in nanoseconds
<code>in</code>	<code>mode</code>	Selects the timer mode, see <a href="#">RTDM_TIMERMODE_XXX</a> for details

Returns

0 on success, otherwise:

- -EINTR is returned if calling task has been unblock by a signal or explicitly via [rtdm\\_task\\_unblock\(\)](#).
- -EPERM *may* be returned if an illegal invocation environment is detected.
- -EINVAL is returned if an invalid parameter was passed.

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: always, unless the specified time already passed.

5.11.3.10 `int rtdm_task_sleep_until ( nanosecs_abs_t wakeup_time )`

Sleep until a specified absolute time.

**Deprecated** Use `rtdm_task_sleep_abs` instead!

Parameters

<code>in</code>	<code>wakeup_time</code>	Absolute timeout in nanoseconds
-----------------	--------------------------	---------------------------------

Returns

0 on success, otherwise:

- -EINTR is returned if calling task has been unblock by a signal or explicitly via `rtdm_task_unblock()`.
- -EPERM *may* be returned if an illegal invocation environment is detected.

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: always, unless the specified time already passed.

5.11.3.11 `int rtdm_task_unblock ( rtdm_task_t * task )`

Activate a blocked real-time task.

Returns

Non-zero is returned if the task was actually unblocked from a pending wait state, 0 otherwise.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.11.3.12 `int rtdm_task_wait_period ( void )`

Wait on next real-time task period.



---

## Returns

0 on success, otherwise:

- -EINVAL is returned if calling task is not in periodic mode.
- -ETIMEDOUT is returned if a timer overrun occurred, which indicates that a previous release point has been missed by the calling task.

## Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: always, unless a timer overrun occurred.

## 5.12 Timer Services

### Typedefs

- typedef void(\* [rtdm\\_timer\\_handler\\_t](#))([rtdm\\_timer\\_t](#) \*timer)  
*Timer handler.*

### Functions

- int [rtdm\\_timer\\_init](#) ([rtdm\\_timer\\_t](#) \*timer, [rtdm\\_timer\\_handler\\_t](#) handler, const char \*name)  
*Initialise a timer.*
- void [rtdm\\_timer\\_destroy](#) ([rtdm\\_timer\\_t](#) \*timer)  
*Destroy a timer.*
- int [rtdm\\_timer\\_start](#) ([rtdm\\_timer\\_t](#) \*timer, [nanosecs\\_abs\\_t](#) expiry, [nanosecs\\_rel\\_t](#) interval, enum [rtdm\\_timer\\_mode](#) mode)  
*Start a timer.*
- void [rtdm\\_timer\\_stop](#) ([rtdm\\_timer\\_t](#) \*timer)  
*Stop a timer.*
- int [rtdm\\_timer\\_start\\_in\\_handler](#) ([rtdm\\_timer\\_t](#) \*timer, [nanosecs\\_abs\\_t](#) expiry, [nanosecs\\_rel\\_t](#) interval, enum [rtdm\\_timer\\_mode](#) mode)  
*Start a timer from inside a timer handler.*
- void [rtdm\\_timer\\_stop\\_in\\_handler](#) ([rtdm\\_timer\\_t](#) \*timer)  
*Stop a timer from inside a timer handler.*

### RTDM\_TIMERMODE\_XXX

#### Timer operation modes

- enum [rtdm\\_timer\\_mode](#) { [RTDM\\_TIMERMODE\\_RELATIVE](#) = XN\_RELATIVE, [RTDM\\_TIMERMODE\\_ABSOLUTE](#) = XN\_ABSOLUTE, [RTDM\\_TIMERMODE\\_REALTIME](#) = XN\_REALTIME }

#### 5.12.1 Detailed Description

#### 5.12.2 Typedef Documentation

##### 5.12.2.1 typedef void(\* [rtdm\\_timer\\_handler\\_t](#))([rtdm\\_timer\\_t](#) \*timer)

Timer handler.

Parameters

<code>in</code>	<code>timer</code>	Timer handle as returned by <a href="#">rtdm_timer_init()</a>
-----------------	--------------------	---

#### 5.12.3 Enumeration Type Documentation

##### 5.12.3.1 enum [rtdm\\_timer\\_mode](#)

Enumerator

**[RTDM\\_TIMERMODE\\_RELATIVE](#)** Monotonic timer with relative timeout.

**[RTDM\\_TIMERMODE\\_ABSOLUTE](#)** Monotonic timer with absolute timeout.

**[RTDM\\_TIMERMODE\\_REALTIME](#)** Adjustable timer with absolute timeout.

## 5.12.4 Function Documentation

5.12.4.1 `void rtdm_timer_destroy ( rtdm_timer_t * timer )`

Destroy a timer.

Parameters

<code>in,out</code>	<code>timer</code>	Timer handle as returned by <a href="#">rtdm_timer_init()</a>
---------------------	--------------------	---

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.12.4.2 `int rtdm_timer_init ( rtdm_timer_t * timer, rtdm_timer_handler_t handler, const char * name )`

Initialise a timer.

Parameters

<code>in,out</code>	<code>timer</code>	Timer handle
<code>in</code>	<code>handler</code>	Handler to be called on timer expiry
<code>in</code>	<code>name</code>	Optional timer name

Returns

0 on success, otherwise negative error code

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.12.4.3 `int rtdm_timer_start ( rtdm_timer_t * timer, nanosecs_abs_t expiry, nanosecs_rel_t interval, enum rtdm_timer_mode mode )`

Start a timer.

Parameters

<i>in, out</i>	<i>timer</i>	Timer handle as returned by <a href="#">rtm_timer_init()</a>
<i>in</i>	<i>expiry</i>	Firing time of the timer, <i>mode</i> defines if relative or absolute
<i>in</i>	<i>interval</i>	Relative reload value, > 0 if the timer shall work in periodic mode with the specific interval, 0 for one-shot timers
<i>in</i>	<i>mode</i>	Defines the operation mode, see <a href="#">RTDM_TIMERMODE_xxx</a> for possible values

Returns

0 on success, otherwise:

- -ETIMEDOUT is returned if *expiry* describes an absolute date in the past.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

```
5.12.4.4 int rtdm_timer_start_in_handler ( rtdm_timer_t * timer, nanosecs_abs_t expiry,
nanosecs_rel_t interval, enum rtdm_timer_mode mode )
```

Start a timer from inside a timer handler.

Parameters

<i>in, out</i>	<i>timer</i>	Timer handle as returned by <a href="#">rtm_timer_init()</a>
<i>in</i>	<i>expiry</i>	Firing time of the timer, <i>mode</i> defines if relative or absolute
<i>in</i>	<i>interval</i>	Relative reload value, > 0 if the timer shall work in periodic mode with the specific interval, 0 for one-shot timers
<i>in</i>	<i>mode</i>	Defines the operation mode, see <a href="#">RTDM_TIMERMODE_xxx</a> for possible values

Returns

0 on success, otherwise:

- -ETIMEDOUT is returned if *expiry* describes an absolute date in the past.

Environments:

This service can be called from:

- Timer handler

Rescheduling: never.

```
5.12.4.5 void rtdm_timer_stop ( rtdm_timer_t * timer )
```

Stop a timer.

Parameters

<code>in,out</code>	<code>timer</code>	Timer handle as returned by <a href="#">rtm_timer_init()</a>
---------------------	--------------------	--

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.12.4.6 `void rtdm_timer_stop_in_handler ( rtdm_timer_t * timer )`

Stop a timer from inside a timer handler.

Parameters

<code>in,out</code>	<code>timer</code>	Timer handle as returned by <a href="#">rtm_timer_init()</a>
---------------------	--------------------	--

Environments:

This service can be called from:

- Timer handler

Rescheduling: never.

## 5.13 Synchronisation Services

### Functions

- int `rtdm_select_bind` (int fd, `rtdm_selector_t` \*selector, enum `rtdm_selecttype` type, unsigned fd\_index)  
*Bind a selector to specified event types of a given file descriptor.*

### RTDM\_SELECTTYPE\_XXX

Event types select can bind to

- enum `rtdm_selecttype` { `RTDM_SELECTTYPE_READ` = XNSELECT\_READ, `RTDM_SELECTTYPE_WRITE` = XNSELECT\_WRITE, `RTDM_SELECTTYPE_EXCEPT` = XNSELECT\_EXCEPT }

### Global Lock across Scheduler Invocation

- #define `RTDM_EXECUTE_ATOMICALLY`(code\_block)  
*Execute code block atomically.*

### Spinlock with Preemption Deactivation

- typedef `rthal_spinlock_t` `rtdm_lock_t`  
*Lock variable.*
- typedef unsigned long `rtdm_lockctx_t`  
*Variable to save the context while holding a lock.*
- #define `RTDM_LOCK_UNLOCKED` `RTHAL_SPIN_LOCK_UNLOCKED`  
*Static lock initialisation.*
- #define `rtdm_lock_init`(lock) `rthal_spin_lock_init`(lock)  
*Dynamic lock initialisation.*
- #define `rtdm_lock_get`(lock) `rthal_spin_lock`(lock)  
*Acquire lock from non-preemptible contexts.*
- #define `rtdm_lock_put`(lock)  
*Release lock without preemption restoration.*
- #define `rtdm_lock_get_irqsave`(lock, context)  
*Acquire lock and disable preemption.*
- #define `rtdm_lock_put_irqrestore`(lock, context)  
*Release lock and restore preemption state.*
- #define `rtdm_lock_irqsave`(context) `rthal_local_irq_save`(context)  
*Disable preemption locally.*
- #define `rtdm_lock_irqrestore`(context) `rthal_local_irq_restore`(context)  
*Restore preemption state.*

### Timeout Sequence Management

- void `rtdm_toseq_init` (`rtdm_toseq_t` \*timeout\_seq, `nanosecs_rel_t` timeout)  
*Initialise a timeout sequence.*
- `EXPORT_SYMBOL_GPL` (`rtdm_toseq_init`)  
*Initialise a timeout sequence.*

## Event Services

- void `rtm_event_init` (`rtm_event_t *event`, unsigned long pending)  
*Initialise an event.*
- `EXPORT_SYMBOL_GPL` (`rtm_event_init`)  
*Initialise an event.*
- void `rtm_event_destroy` (`rtm_event_t *event`)  
*Destroy an event.*
- void `rtm_event_pulse` (`rtm_event_t *event`)  
*Signal an event occurrence to currently listening waiters.*
- void `rtm_event_signal` (`rtm_event_t *event`)  
*Signal an event occurrence.*
- `EXPORT_SYMBOL_GPL` (`rtm_event_signal`)  
*Initialise an event.*
- int `rtm_event_wait` (`rtm_event_t *event`)  
*Wait on event occurrence.*
- `EXPORT_SYMBOL_GPL` (`rtm_event_wait`)  
*Initialise an event.*
- int `rtm_event_timedwait` (`rtm_event_t *event`, `nanosecs_rel_t` timeout, `rtm_toseq_t *timeout_seq`)  
*Wait on event occurrence with timeout.*
- `EXPORT_SYMBOL_GPL` (`rtm_event_timedwait`)  
*Initialise an event.*
- void `rtm_event_clear` (`rtm_event_t *event`)  
*Clear event state.*
- `EXPORT_SYMBOL_GPL` (`rtm_event_clear`)  
*Initialise an event.*
- int `rtm_event_select_bind` (`rtm_event_t *event`, `rtm_selector_t *selector`, enum `rtm_selecttype` type, unsigned `fd_index`)  
*Bind a selector to an event.*
- `EXPORT_SYMBOL_GPL` (`rtm_event_select_bind`)  
*Initialise an event.*

## Semaphore Services

- void `rtm_sem_init` (`rtm_sem_t *sem`, unsigned long value)  
*Initialise a semaphore.*
- `EXPORT_SYMBOL_GPL` (`rtm_sem_init`)  
*Initialise a semaphore.*
- void `rtm_sem_destroy` (`rtm_sem_t *sem`)  
*Destroy a semaphore.*
- int `rtm_sem_down` (`rtm_sem_t *sem`)  
*Decrement a semaphore.*
- `EXPORT_SYMBOL_GPL` (`rtm_sem_down`)  
*Initialise a semaphore.*
- int `rtm_sem_timeddown` (`rtm_sem_t *sem`, `nanosecs_rel_t` timeout, `rtm_toseq_t *timeout_seq`)  
*Decrement a semaphore with timeout.*
- `EXPORT_SYMBOL_GPL` (`rtm_sem_timeddown`)  
*Initialise a semaphore.*
- void `rtm_sem_up` (`rtm_sem_t *sem`)  
*Increment a semaphore.*

- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_sem\\_up\)](#)  
*Initialise a semaphore.*
- `int rtdm_sem_select_bind (rtdm_sem_t *sem, rtdm_selector_t *selector, enum rtdm_selecttype type, unsigned fd_index)`  
*Bind a selector to a semaphore.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_sem\\_select\\_bind\)](#)  
*Initialise a semaphore.*

## Mutex Services

- `void rtdm_mutex_init (rtdm_mutex_t *mutex)`  
*Initialise a mutex.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_mutex\\_init\)](#)  
*Initialise a mutex.*
- `void rtdm_mutex_destroy (rtdm_mutex_t *mutex)`  
*Destroy a mutex.*
- `void rtdm_mutex_unlock (rtdm_mutex_t *mutex)`  
*Release a mutex.*
- `int rtdm_mutex_lock (rtdm_mutex_t *mutex)`  
*Request a mutex.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_mutex\\_lock\)](#)  
*Initialise a mutex.*
- `int rtdm_mutex_timedlock (rtdm_mutex_t *mutex, nanosecs_rel_t timeout, rtdm_toseq_t *timeout_seq)`  
*Request a mutex with timeout.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_mutex\\_timedlock\)](#)  
*Initialise a mutex.*

### 5.13.1 Detailed Description

### 5.13.2 Macro Definition Documentation

#### 5.13.2.1 #define RTDM\_EXECUTE\_ATOMICALY( code\_block )

##### Value:

```
{
    <ENTER_ATOMIC_SECTION>
    code_block;
    <LEAVE_ATOMIC_SECTION>
}
```

Execute code block atomically.

Generally, it is illegal to suspend the current task by calling [rtdm\\_task\\_sleep\(\)](#), [rtdm\\_event\\_wait\(\)](#), etc. while holding a spinlock. In contrast, this macro allows to combine several operations including a potentially rescheduling call to an atomic code block with respect to other [RTDM\\_EXECUTE\\_ATOMICALY\(\)](#) blocks. The macro is a light-weight alternative for protecting code blocks via mutexes, and it can even be used to synchronise real-time and non-real-time contexts.



## Parameters

<i>code_block</i>	Commands to be executed atomically
-------------------	------------------------------------

## Note

It is not allowed to leave the code block explicitly by using `break`, `return`, `goto`, etc. This would leave the global lock held during the code block execution in an inconsistent state. Moreover, do not embed complex operations into the code block. Consider that they will be executed under preemption lock with interrupts switched-off. Also note that invocation of rescheduling calls may break the atomicity until the task gains the CPU again.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible, depends on functions called within *code\_block*.

5.13.2.2 `#define rtdm_lock_get( lock ) rthal_spin_lock(lock)`

Acquire lock from non-preemptible contexts.

## Parameters

<i>lock</i>	Address of lock variable
-------------	--------------------------

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.2.3 `#define rtdm_lock_get_irqsave( lock, context )`

## Value:

```
do {
    rthal_spin_lock_irqsave(lock, context); \
    __xnpod_lock_sched(); \
} while (0)
```

Acquire lock and disable preemption.

Parameters

<i>lock</i>	Address of lock variable
<i>context</i>	name of local variable to store the context in

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

Referenced by `rtdm_ratelimit()`.

5.13.2.4 `#define rtdm_lock_init( lock ) rthal_spin_lock_init(lock)`

Dynamic lock initialisation.

Parameters

<i>lock</i>	Address of lock variable
-------------	--------------------------

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.2.5 `#define rtdm_lock_irqrestore( context ) rthal_local_irq_restore(context)`

Restore preemption state.

Parameters

<i>context</i>	name of local variable which stored the context
----------------	---

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.13.2.6 `#define rtdm_lock_irqsave( context ) rthal_local_irq_save(context)`

Disable preemption locally.

Parameters

<i>context</i>	name of local variable to store the context in
----------------	--

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.2.7 `#define rtdm_lock_put( lock )`

**Value:**

```
do {
    rthal_spin_unlock(lock); \
    __xnpod_unlock_sched(); \
} while (0)
```

Release lock without preemption restoration.

Parameters

<i>lock</i>	Address of lock variable
-------------	--------------------------

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.2.8 `#define rtdm_lock_put_irqrestore( lock, context )`

**Value:**

```
do {
    rthal_spin_unlock(lock); \
    __xnpod_unlock_sched(); \
    rthal_local_irq_restore(context); \
} while (0)
```

Release lock and restore preemption state.

## Parameters

<i>lock</i>	Address of lock variable
<i>context</i>	name of local variable which stored the context

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

Referenced by `rtdm_ratelimit()`.

## 5.13.3 Enumeration Type Documentation

### 5.13.3.1 enum `rtdm_selecttype`

Enumerator

**`RTDM_SELECTTYPE_READ`** Select input data availability events.

**`RTDM_SELECTTYPE_WRITE`** Select output buffer availability events.

**`RTDM_SELECTTYPE_EXCEPT`** Select exceptional events.

## 5.13.4 Function Documentation

### 5.13.4.1 `EXPORT_SYMBOL_GPL ( rtdm_toseq_init )`

Initialise a timeout sequence.

This service initialises a timeout sequence handle according to the given timeout value. Timeout sequences allow to maintain a continuous *timeout* across multiple calls of blocking synchronisation services. A typical application scenario is given below.

Parameters

<code>in,out</code>	<i>timeout_seq</i>	Timeout sequence handle
<code>in</code>	<i>timeout</i>	Relative timeout in nanoseconds, see <a href="#">RTDM_TIMEOUT_XXX</a> for special values

Application Scenario:

```
int device_service_routine(...)
{
    rtdm_toseq_t timeout_seq;
    ...

    rtdm_toseq_init(&timeout_seq, timeout);
    ...
    while (received < requested) {
        ret = rtdm_event_timedwait(&data_available, timeout, &timeout_seq);
        if (ret < 0) // including -ETIMEDOUT
            break;

        // receive some data
        ...
    }
    ...
}
*
```

Using a timeout sequence in such a scenario avoids that the user-provided relative `timeout` is restarted on every call to `rtm_event_timedwait()`, potentially causing an overall delay that is larger than specified by `timeout`. Moreover, all functions supporting timeout sequences also interpret special timeout values (infinite and non-blocking), disburdening the driver developer from handling them separately.

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: never.

#### 5.13.4.2 EXPORT\_SYMBOL\_GPL ( `rtm_event_init` )

Initialise an event.

Parameters

<code>in,out</code>	<code>event</code>	Event handle
<code>in</code>	<code>pending</code>	Non-zero if event shall be initialised as set, 0 otherwise

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.3 EXPORT\_SYMBOL\_GPL ( `rtm_event_signal` )

Initialise an event.

Parameters

<code>in,out</code>	<code>event</code>	Event handle
<code>in</code>	<code>pending</code>	Non-zero if event shall be initialised as set, 0 otherwise

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.4 EXPORT\_SYMBOL\_GPL ( `rtm_event_wait` )

Initialise an event.

Parameters

<i>in,out</i>	<i>event</i>	Event handle
<i>in</i>	<i>pending</i>	Non-zero if event shall be initialised as set, 0 otherwise

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.5 EXPORT\_SYMBOL\_GPL ( `rtdm_event_timedwait` )

Initialise an event.

Parameters

<i>in,out</i>	<i>event</i>	Event handle
<i>in</i>	<i>pending</i>	Non-zero if event shall be initialised as set, 0 otherwise

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.6 EXPORT\_SYMBOL\_GPL ( `rtdm_event_clear` )

Initialise an event.

Parameters

<i>in,out</i>	<i>event</i>	Event handle
<i>in</i>	<i>pending</i>	Non-zero if event shall be initialised as set, 0 otherwise

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.7 EXPORT\_SYMBOL\_GPL ( `rtdm_event_select_bind` )

Initialise an event.

Parameters

<i>in, out</i>	<i>event</i>	Event handle
<i>in</i>	<i>pending</i>	Non-zero if event shall be initialised as set, 0 otherwise

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.8 EXPORT\_SYMBOL\_GPL ( **rtdm\_sem\_init** )

Initialise a semaphore.

Parameters

<i>in, out</i>	<i>sem</i>	Semaphore handle
<i>in</i>	<i>value</i>	Initial value of the semaphore

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.9 EXPORT\_SYMBOL\_GPL ( **rtdm\_sem\_down** )

Initialise a semaphore.

Parameters

<i>in, out</i>	<i>sem</i>	Semaphore handle
<i>in</i>	<i>value</i>	Initial value of the semaphore

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.10 EXPORT\_SYMBOL\_GPL ( **rtdm\_sem\_timeddown** )

Initialise a semaphore.

Parameters

<i>in,out</i>	<i>sem</i>	Semaphore handle
<i>in</i>	<i>value</i>	Initial value of the semaphore

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.11 EXPORT\_SYMBOL\_GPL ( *rtdm\_sem\_up* )

Initialise a semaphore.

Parameters

<i>in,out</i>	<i>sem</i>	Semaphore handle
<i>in</i>	<i>value</i>	Initial value of the semaphore

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.12 EXPORT\_SYMBOL\_GPL ( *rtdm\_sem\_select\_bind* )

Initialise a semaphore.

Parameters

<i>in,out</i>	<i>sem</i>	Semaphore handle
<i>in</i>	<i>value</i>	Initial value of the semaphore

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.13 EXPORT\_SYMBOL\_GPL ( *rtdm\_mutex\_init* )

Initialise a mutex.

This function initialises a basic mutex with priority inversion protection. "Basic", as it does not allow a mutex owner to recursively lock the same mutex again.



Parameters

<code>in,out</code>	<code>mutex</code>	Mutex handle
---------------------	--------------------	--------------

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.14 EXPORT\_SYMBOL\_GPL ( `rtdm_mutex_lock` )

Initialise a mutex.

This function initialises a basic mutex with priority inversion protection. "Basic", as it does not allow a mutex owner to recursively lock the same mutex again.

Parameters

<code>in,out</code>	<code>mutex</code>	Mutex handle
---------------------	--------------------	--------------

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.15 EXPORT\_SYMBOL\_GPL ( `rtdm_mutex_timedlock` )

Initialise a mutex.

This function initialises a basic mutex with priority inversion protection. "Basic", as it does not allow a mutex owner to recursively lock the same mutex again.

Parameters

<code>in,out</code>	<code>mutex</code>	Mutex handle
---------------------	--------------------	--------------

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

#### 5.13.4.16 `void rtdm_event_clear ( rtdm_event_t * event )`

Clear event state.

Parameters

<i>in, out</i>	<i>event</i>	Event handle as returned by <a href="#">rtdm_event_init()</a>
----------------	--------------	---

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.4.17 `void rtdm_event_destroy ( rtdm_event_t * event )`

Destroy an event.

Parameters

<i>in, out</i>	<i>event</i>	Event handle as returned by <a href="#">rtdm_event_init()</a>
----------------	--------------	---

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.13.4.18 `void rtdm_event_init ( rtdm_event_t * event, unsigned long pending )`

Initialise an event.

Parameters

<i>in, out</i>	<i>event</i>	Event handle
<i>in</i>	<i>pending</i>	Non-zero if event shall be initialised as set, 0 otherwise

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.4.19 `void rtdm_event_pulse ( rtdm_event_t * event )`

Signal an event occurrence to currently listening waiters.

This function wakes up all current waiters of the given event, but it does not change the event state. Subsequently callers of [rtdm\\_event\\_wait\(\)](#) or [rtdm\\_event\\_timedwait\(\)](#) will therefore be blocked first.

Parameters

<i>in,out</i>	<i>event</i>	Event handle as returned by <a href="#">rtdm_event_init()</a>
---------------	--------------	---

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.13.4.20 `int rtdm_event_select_bind ( rtdm_event_t * event, rtdm_selector_t * selector, enum rtdm_selecttype type, unsigned fd_index )`

Bind a selector to an event.

This functions binds the given selector to an event so that the former is notified when the event state changes. Typically the select binding handler will invoke this service.

Parameters

<i>in,out</i>	<i>event</i>	Event handle as returned by <a href="#">rtdm_event_init()</a>
<i>in,out</i>	<i>selector</i>	Selector as passed to the select binding handler
<i>in</i>	<i>type</i>	Type of the bound event as passed to the select binding handler
<i>in</i>	<i>fd_index</i>	File descriptor index as passed to the select binding handler

Returns

0 on success, otherwise:

- -ENOMEM is returned if there is insufficient memory to establish the dynamic binding.
- -EINVAL is returned if *type* or *fd\_index* are invalid.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.4.21 `void rtdm_event_signal ( rtdm_event_t * event )`

Signal an event occurrence.

This function sets the given event and wakes up all current waiters. If no waiter is presently registered, the next call to [rtdm\\_event\\_wait\(\)](#) or [rtdm\\_event\\_timedwait\(\)](#) will return immediately.

## Parameters

<i>in, out</i>	<i>event</i>	Event handle as returned by <a href="#">rtdm_event_init()</a>
----------------	--------------	---

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.13.4.22 `int rtdm_event_timedwait ( rtdm_event_t * event, nanosecs_rel_t timeout, rtdm_toseq_t * timeout_seq )`

Wait on event occurrence with timeout.

This function waits or tests for the occurrence of the given event, taking the provided timeout into account. On successful return, the event is reset.

## Parameters

<i>in, out</i>	<i>event</i>	Event handle as returned by <a href="#">rtdm_event_init()</a>
<i>in</i>	<i>timeout</i>	Relative timeout in nanoseconds, see <a href="#">RTDM_TIMEOUT_XXX</a> for special values
<i>in, out</i>	<i>timeout_seq</i>	Handle of a timeout sequence as returned by <a href="#">rtdm_toseq_init()</a> or NULL

## Returns

0 on success, otherwise:

- -ETIMEDOUT is returned if the if the request has not been satisfied within the specified amount of time.
- -EINTR is returned if calling task has been unblock by a signal or explicitly via [rtdm\\_task\\_unblock\(\)](#).
- -EIDRM is returned if *event* has been destroyed.
- -EPERM *may* be returned if an illegal invocation environment is detected.
- -EWOULDBLOCK is returned if a negative *timeout* (i.e., non-blocking operation) has been specified.

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: possible.

Referenced by [rtdm\\_event\\_wait\(\)](#).

---

5.13.4.23 `int rtdm_event_wait ( rtdm_event_t * event )`

Wait on event occurrence.

This is the light-weight version of [rtdm\\_event\\_timedwait\(\)](#), implying an infinite timeout.

## Parameters

<code>in,out</code>	<code>event</code>	Event handle as returned by <a href="#">rtdm_event_init()</a>
---------------------	--------------------	---

## Returns

0 on success, otherwise:

- -EINTR is returned if calling task has been unblock by a signal or explicitly via [rtdm\\_task\\_unblock\(\)](#).
- -EIDRM is returned if `event` has been destroyed.
- -EPERM *may* be returned if an illegal invocation environment is detected.

## Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: possible.

References [rtdm\\_event\\_timedwait\(\)](#).

5.13.4.24 `void rtdm_mutex_destroy ( rtdm_mutex_t * mutex )`

Destroy a mutex.

## Parameters

<code>in,out</code>	<code>mutex</code>	Mutex handle as returned by <a href="#">rtdm_mutex_init()</a>
---------------------	--------------------	---

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.13.4.25 `void rtdm_mutex_init ( rtdm_mutex_t * mutex )`

Initialise a mutex.

This function initialises a basic mutex with priority inversion protection. "Basic", as it does not allow a mutex owner to recursively lock the same mutex again.

## Parameters

<code>in,out</code>	<code>mutex</code>	Mutex handle
---------------------	--------------------	--------------

Environments:

This service can be called from:

- Kernel module initialization/cleanup code

- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.4.26 `int rtdm_mutex_lock ( rtdm_mutex_t * mutex )`

Request a mutex.

This is the light-weight version of [rtdm\\_mutex\\_timedlock\(\)](#), implying an infinite timeout.

Parameters

<code>in, out</code>	<code>mutex</code>	Mutex handle as returned by <a href="#">rtdm_mutex_init()</a>
----------------------	--------------------	---

Returns

0 on success, otherwise:

- -EIDRM is returned if `mutex` has been destroyed.
- -EPERM *may* be returned if an illegal invocation environment is detected.

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: possible.

References [rtdm\\_mutex\\_timedlock\(\)](#).

5.13.4.27 `int rtdm_mutex_timedlock ( rtdm_mutex_t * mutex, nanosecs_rel_t timeout, rtdm_toseq_t * timeout_seq )`

Request a mutex with timeout.

This function tries to acquire the given mutex. If it is not available, the caller is blocked unless non-blocking operation was selected.

Parameters

<code>in, out</code>	<code>mutex</code>	Mutex handle as returned by <a href="#">rtdm_mutex_init()</a>
<code>in</code>	<code>timeout</code>	Relative timeout in nanoseconds, see <a href="#">RTDM_TIMEOUT_xxx</a> for special values
<code>in, out</code>	<code>timeout_seq</code>	Handle of a timeout sequence as returned by <a href="#">rtdm_toseq_init()</a> or NULL

Returns

0 on success, otherwise:

- -ETIMEDOUT is returned if the if the request has not been satisfied within the specified amount of time.

- -EWOULDBLOCK is returned if *timeout* is negative and the semaphore value is currently not positive.
- -EIDRM is returned if *mutex* has been destroyed.
- -EPERM *may* be returned if an illegal invocation environment is detected.

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: possible.

Referenced by `rtm_mutex_lock()`.

5.13.4.28 `void rtdm_mutex_unlock ( rtdm_mutex_t * mutex )`

Release a mutex.

This function releases the given mutex, waking up a potential waiter which was blocked upon `rtm_mutex_lock()` or `rtm_mutex_timedlock()`.

Parameters

<code>in, out</code>	<code>mutex</code>	Mutex handle as returned by <code>rtm_mutex_init()</code>
----------------------	--------------------	---

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: possible.

5.13.4.29 `int rtdm_select_bind ( int fd, rtdm_selector_t * selector, enum rtdm_selecttype type, unsigned fd_index )`

Bind a selector to specified event types of a given file descriptor.

This function is invoked by higher RTOS layers implementing select-like services. It shall not be called directly by RTDM drivers.

Parameters

<code>in</code>	<code>fd</code>	File descriptor to bind to
<code>in, out</code>	<code>selector</code>	Selector object that shall be bound to the given event
<code>in</code>	<code>type</code>	Event type the caller is interested in
<code>in</code>	<code>fd_index</code>	Index in the file descriptor set of the caller

Returns

0 on success, otherwise:

- -EBADF is returned if the file descriptor *fd* cannot be resolved.
- -EINVAL is returned if *type* or *fd\_index* are invalid.



Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

References `rtm_dev_context::ops`, `rtm_context_get()`, `rtm_context_unlock()`, and `rtm_operations::select_bind`.

5.13.4.30 `void rtdm_sem_destroy ( rtdm_sem_t * sem )`

Destroy a semaphore.

Parameters

<code>in,out</code>	<code>sem</code>	Semaphore handle as returned by <a href="#">rtm_sem_init()</a>
---------------------	------------------	--

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.13.4.31 `int rtdm_sem_down ( rtdm_sem_t * sem )`

Decrement a semaphore.

This is the light-weight version of [rtm\\_sem\\_timeddown\(\)](#), implying an infinite timeout.

Parameters

<code>in,out</code>	<code>sem</code>	Semaphore handle as returned by <a href="#">rtm_sem_init()</a>
---------------------	------------------	--

Returns

0 on success, otherwise:

- `-EINTR` is returned if calling task has been unblock by a signal or explicitly via [rtm\\_task\\_unblock\(\)](#).
- `-EIDRM` is returned if `sem` has been destroyed.
- `-EPERM` may be returned if an illegal invocation environment is detected.

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: possible.

References [rtm\\_sem\\_timeddown\(\)](#).

5.13.4.32 void rtdm\_sem\_init ( rtdm\_sem\_t \* sem, unsigned long value )

Initialise a semaphore.

Parameters

in,out	<i>sem</i>	Semaphore handle
in	<i>value</i>	Initial value of the semaphore

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.4.33 int rtdm\_sem\_select\_bind ( rtdm\_sem\_t \* sem, rtdm\_selector\_t \* selector, enum **rtdm\_selecttype** type, unsigned fd\_index )

Bind a selector to a semaphore.

This functions binds the given selector to the semaphore so that the former is notified when the semaphore state changes. Typically the select binding handler will invoke this service.

Parameters

in,out	<i>sem</i>	Semaphore handle as returned by <a href="#">rtdm_sem_init()</a>
in,out	<i>selector</i>	Selector as passed to the select binding handler
in	<i>type</i>	Type of the bound event as passed to the select binding handler
in	<i>fd_index</i>	File descriptor index as passed to the select binding handler

Returns

0 on success, otherwise:

- -ENOMEM is returned if there is insufficient memory to establish the dynamic binding.
- -EINVAL is returned if *type* or *fd\_index* are invalid.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.13.4.34 int rtdm\_sem\_timeddown ( rtdm\_sem\_t \* sem, **nanosecs\_rel\_t** timeout, rtdm\_toseq\_t \* timeout\_seq )

Decrement a semaphore with timeout.

This function tries to decrement the given semaphore's value if it is positive on entry. If not, the caller is blocked unless non-blocking operation was selected.

## Parameters

in,out	<i>sem</i>	Semaphore handle as returned by <a href="#">rtm_sem_init()</a>
in	<i>timeout</i>	Relative timeout in nanoseconds, see <a href="#">RTDM_TIMEOUT_xxx</a> for special values
in,out	<i>timeout_seq</i>	Handle of a timeout sequence as returned by <a href="#">rtm_toseq_init()</a> or NULL

## Returns

0 on success, otherwise:

- -ETIMEDOUT is returned if the if the request has not been satisfied within the specified amount of time.
- -EWOULDBLOCK is returned if *timeout* is negative and the semaphore value is currently not positive.
- -EINTR is returned if calling task has been unblock by a signal or explicitly via [rtm\\_task\\_unblock\(\)](#).
- -EIDRM is returned if *sem* has been destroyed.
- -EPERM *may* be returned if an illegal invocation environment is detected.

## Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: possible.

Referenced by [rtm\\_sem\\_down\(\)](#).

5.13.4.35 void [rtm\\_sem\\_up](#) ( [rtm\\_sem\\_t](#) \* *sem* )

Increment a semaphore.

This function increments the given semaphore's value, waking up a potential waiter which was blocked upon [rtm\\_sem\\_down\(\)](#).

## Parameters

in,out	<i>sem</i>	Semaphore handle as returned by <a href="#">rtm_sem_init()</a>
--------	------------	--

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.13.4.36 void rtdm\_toseq\_init ( rtdm\_toseq\_t \* timeout\_seq, nanosecs\_rel\_t timeout )

Initialise a timeout sequence.

This service initialises a timeout sequence handle according to the given timeout value. Timeout sequences allow to maintain a continuous *timeout* across multiple calls of blocking synchronisation services. A typical application scenario is given below.

Parameters

in,out	<i>timeout_seq</i>	Timeout sequence handle
in	<i>timeout</i>	Relative timeout in nanoseconds, see <a href="#">RTDM_TIMEOUT_xxx</a> for special values

Application Scenario:

```
int device_service_routine(...)
{
    rtdm_toseq_t timeout_seq;
    ...

    rtdm_toseq_init(&timeout_seq, timeout);
    ...
    while (received < requested) {
        ret = rtdm_event_timedwait(&data_available, timeout, &timeout_seq);
        if (ret < 0) // including -ETIMEDOUT
            break;

        // receive some data
        ...
    }
    ...
}
*
```

Using a timeout sequence in such a scenario avoids that the user-provided relative *timeout* is restarted on every call to [rtdm\\_event\\_timedwait\(\)](#), potentially causing an overall delay that is larger than specified by *timeout*. Moreover, all functions supporting timeout sequences also interpret special timeout values (infinite and non-blocking), disburdening the driver developer from handling them separately.

Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT)

Rescheduling: never.

## 5.14 Interrupt Management Services

### Macros

- `#define rtdm_irq_get_arg(irq_handle, type) ((type *)irq_handle->cookie)`  
*Retrieve IRQ handler argument.*

### Typedefs

- `typedef int(* rtdm_irq_handler_t)(rtdm_irq_t *irq_handle)`  
*Interrupt handler.*

### Functions

- `int rtdm_irq_request(rtdm_irq_t *irq_handle, unsigned int irq_no, rtdm_irq_handler_t handler, unsigned long flags, const char *device_name, void *arg)`  
*Register an interrupt handler.*
- `int rtdm_irq_free(rtdm_irq_t *irq_handle)`  
*Release an interrupt handler.*
- `int rtdm_irq_enable(rtdm_irq_t *irq_handle)`  
*Enable interrupt line.*
- `int rtdm_irq_disable(rtdm_irq_t *irq_handle)`  
*Disable interrupt line.*

### RTDM\_IRQTYPE\_XXX

#### Interrupt registrations flags

- `#define RTDM_IRQTYPE_SHARED XN_ISR_SHARED`  
*Enable IRQ-sharing with other real-time drivers.*
- `#define RTDM_IRQTYPE_EDGE XN_ISR_EDGE`  
*Mark IRQ as edge-triggered, relevant for correct handling of shared edge-triggered IRQs.*

### RTDM\_IRQ\_XXX

#### Return flags of interrupt handlers

- `#define RTDM_IRQ_NONE XN_ISR_NONE`  
*Unhandled interrupt.*
- `#define RTDM_IRQ_HANDLED XN_ISR_HANDLED`  
*Denote handled interrupt.*

### 5.14.1 Detailed Description

### 5.14.2 Macro Definition Documentation

#### 5.14.2.1 `#define rtdm_irq_get_arg( irq_handle, type ) ((type *)irq_handle->cookie)`

Retrieve IRQ handler argument.

## Parameters

<i>irq_handle</i>	IRQ handle
<i>type</i>	Type of the pointer to return

## Returns

The argument pointer registered on [rtm\\_irq\\_request\(\)](#) is returned, type-casted to the specified *type*.

## Environments:

This service can be called from:

- Interrupt service routine

Rescheduling: never.

## 5.14.3 Typedef Documentation

5.14.3.1 `typedef int(* rtdm_irq_handler_t)(rtdm_irq_t *irq_handle)`

Interrupt handler.

## Parameters

<i>in</i>	<i>irq_handle</i>	IRQ handle as returned by <a href="#">rtm_irq_request()</a>
-----------	-------------------	---

## Returns

0 or a combination of [RTDM\\_IRQ\\_XXX](#) flags

## 5.14.4 Function Documentation

5.14.4.1 `int rtdm_irq_disable ( rtdm_irq_t * irq_handle )`

Disable interrupt line.

## Parameters

<i>in, out</i>	<i>irq_handle</i>	IRQ handle as returned by <a href="#">rtm_irq_request()</a>
----------------	-------------------	---

## Returns

0 on success, otherwise negative error code

## Note

This service is for exceptional use only. Drivers should always prefer interrupt masking at device level (via corresponding control registers etc.) over masking at line level. Keep in mind that the latter is incompatible with IRQ line sharing and can also be more costly as interrupt controller access requires broader synchronization. Also, certain IRQ types may not allow the invocation over RT and interrupt contexts. The caller is responsible for excluding such conflicts.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.14.4.2 `int rtdm_irq_enable ( rtdm_irq_t * irq_handle )`

Enable interrupt line.

Parameters

<code>in,out</code>	<code>irq_handle</code>	IRQ handle as returned by <a href="#">rtdm_irq_request()</a>
---------------------	-------------------------	--

Returns

0 on success, otherwise negative error code

Note

This service is for exceptional use only. Drivers should always prefer interrupt masking at device level (via corresponding control registers etc.) over masking at line level. Keep in mind that the latter is incompatible with IRQ line sharing and can also be more costly as interrupt controller access requires broader synchronization. Also, certain IRQ types may not allow the invocation over RT and interrupt contexts. The caller is responsible for excluding such conflicts.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

5.14.4.3 `int rtdm_irq_free ( rtdm_irq_t * irq_handle )`

Release an interrupt handler.

Parameters

<code>in,out</code>	<code>irq_handle</code>	IRQ handle as returned by <a href="#">rtdm_irq_request()</a>
---------------------	-------------------------	--

Returns

0 on success, otherwise negative error code

## Note

The caller is responsible for shutting down the IRQ source at device level before invoking this service. In turn, `rtm_irq_free` ensures that any pending event on the given IRQ line is fully processed on return from this service.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- User-space task (non-RT)

Rescheduling: never.

```
5.14.4.4 int rtdm_irq_request ( rtdm_irq_t * irq_handle, unsigned int irq_no, rtdm_irq_handler_t
    handler, unsigned long flags, const char * device_name, void * arg )
```

Register an interrupt handler.

This function registers the provided handler with an IRQ line and enables the line.

## Parameters

in,out	<i>irq_handle</i>	IRQ handle
in	<i>irq_no</i>	Line number of the addressed IRQ
in	<i>handler</i>	Interrupt handler
in	<i>flags</i>	Registration flags, see <a href="#">RTDM_IRQTYPE_XXX</a> for details
in	<i>device_name</i>	Device name to show up in real-time IRQ lists
in	<i>arg</i>	Pointer to be passed to the interrupt handler on invocation

## Returns

0 on success, otherwise:

- -EINVAL is returned if an invalid parameter was passed.
- -EBUSY is returned if the specified IRQ line is already in use.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- User-space task (non-RT)

Rescheduling: never.



## 5.15 Non-Real-Time Signalling Services

### Typedefs

- typedef void(\* [rtdm\\_nrtsig\\_handler\\_t](#))(rtdm\_nrtsig\_t nrt\_sig, void \*arg)  
*Non-real-time signal handler.*

### Functions

- int [rtdm\\_nrtsig\\_init](#) (rtdm\_nrtsig\_t \*nrt\_sig, [rtdm\\_nrtsig\\_handler\\_t](#) handler, void \*arg)  
*Register a non-real-time signal handler.*
- void [rtdm\\_nrtsig\\_destroy](#) (rtdm\_nrtsig\_t \*nrt\_sig)  
*Release a non-realtime signal handler.*
- void [rtdm\\_nrtsig\\_pend](#) (rtdm\_nrtsig\_t \*nrt\_sig)  
*Trigger non-real-time signal.*

#### 5.15.1 Detailed Description

These services provide a mechanism to request the execution of a specified handler in non-real-time context. The triggering can safely be performed in real-time context without suffering from unknown delays. The handler execution will be deferred until the next time the real-time subsystem releases the CPU to the non-real-time part.

#### 5.15.2 Typedef Documentation

##### 5.15.2.1 typedef void(\* [rtdm\\_nrtsig\\_handler\\_t](#))(rtdm\_nrtsig\_t nrt\_sig, void \*arg)

Non-real-time signal handler.

Parameters

in	<i>nrt_sig</i>	Signal handle as returned by <a href="#">rtdm_nrtsig_init()</a>
in	<i>arg</i>	Argument as passed to <a href="#">rtdm_nrtsig_init()</a>

Note

The signal handler will run in soft-IRQ context of the non-real-time subsystem. Note the implications of this context, e.g. no invocation of blocking operations.

#### 5.15.3 Function Documentation

##### 5.15.3.1 void [rtdm\\_nrtsig\\_destroy](#) ( rtdm\_nrtsig\_t \* nrt\_sig )

Release a non-realtime signal handler.

Parameters

in,out	<i>nrt_sig</i>	Signal handle
--------	----------------	---------------

Environments:

This service can be called from:

- Kernel module initialization/cleanup code

- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.15.3.2 `int rtdm_nrtsig_init ( rtdm_nrtsig_t * nrt_sig, rtdm_nrtsig_handler_t handler, void * arg )`

Register a non-real-time signal handler.

Parameters

<code>in,out</code>	<code>nrt_sig</code>	Signal handle
<code>in</code>	<code>handler</code>	Non-real-time signal handler
<code>in</code>	<code>arg</code>	Custom argument passed to <code>handler()</code> on each invocation

Returns

0 on success, otherwise:

- -EAGAIN is returned if no free signal slot is available.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.15.3.3 `void rtdm_nrtsig_pend ( rtdm_nrtsig_t * nrt_sig )`

Trigger non-real-time signal.

Parameters

<code>in,out</code>	<code>nrt_sig</code>	Signal handle
---------------------	----------------------	---------------

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never in real-time context, possible in non-real-time environments.

## 5.16 Utility Services

### Functions

- int [rtdm\\_mmap\\_to\\_user](#) (rtdm\_user\_info\_t \*user\_info, void \*src\_addr, size\_t len, int prot, void \*\*pptr, struct vm\_operations\_struct \*vm\_ops, void \*vm\_private\_data)  
*Map a kernel memory range into the address space of the user.*
- int [rtdm\\_iomap\\_to\\_user](#) (rtdm\_user\_info\_t \*user\_info, phys\_addr\_t src\_addr, size\_t len, int prot, void \*\*pptr, struct vm\_operations\_struct \*vm\_ops, void \*vm\_private\_data)  
*Map an I/O memory range into the address space of the user.*
- int [rtdm\\_munmap](#) (rtdm\_user\_info\_t \*user\_info, void \*ptr, size\_t len)  
*Unmap a user memory range.*
- int [rtdm\\_ratelimit](#) (struct rtdm\_ratelimit\_state \*rs, const char \*func)  
*Enforces a rate limit.*
- void [rtdm\\_printk\\_ratelimited](#) (const char \*format,...)  
*Real-time safe rate-limited message printing on kernel console.*
- void [rtdm\\_printk](#) (const char \*format,...)  
*Real-time safe message printing on kernel console.*
- void \* [rtdm\\_malloc](#) (size\_t size)  
*Allocate memory block in real-time context.*
- void [rtdm\\_free](#) (void \*ptr)  
*Release real-time memory block.*
- int [rtdm\\_read\\_user\\_ok](#) (rtdm\_user\_info\_t \*user\_info, const void \_\_user \*ptr, size\_t size)  
*Check if read access to user-space memory block is safe.*
- int [rtdm\\_rw\\_user\\_ok](#) (rtdm\_user\_info\_t \*user\_info, const void \_\_user \*ptr, size\_t size)  
*Check if read/write access to user-space memory block is safe.*
- int [rtdm\\_copy\\_from\\_user](#) (rtdm\_user\_info\_t \*user\_info, void \*dst, const void \_\_user \*src, size\_t size)  
*Copy user-space memory block to specified buffer.*
- int [rtdm\\_safe\\_copy\\_from\\_user](#) (rtdm\_user\_info\_t \*user\_info, void \*dst, const void \_\_user \*src, size\_t size)  
*Check if read access to user-space memory block and copy it to specified buffer.*
- int [rtdm\\_copy\\_to\\_user](#) (rtdm\_user\_info\_t \*user\_info, void \_\_user \*dst, const void \*src, size\_t size)  
*Copy specified buffer to user-space memory block.*
- int [rtdm\\_safe\\_copy\\_to\\_user](#) (rtdm\_user\_info\_t \*user\_info, void \_\_user \*dst, const void \*src, size\_t size)  
*Check if read/write access to user-space memory block is safe and copy specified buffer to it.*
- int [rtdm\\_strncpy\\_from\\_user](#) (rtdm\_user\_info\_t \*user\_info, char \*dst, const char \_\_user \*src, size\_t count)  
*Copy user-space string to specified buffer.*
- int [rtdm\\_in\\_rt\\_context](#) (void)  
*Test if running in a real-time task.*
- int [rtdm\\_rt\\_capable](#) (rtdm\_user\_info\_t \*user\_info)  
*Test if the caller is capable of running in real-time context.*

#### 5.16.1 Detailed Description

#### 5.16.2 Function Documentation

- 5.16.2.1 int [rtdm\\_copy\\_from\\_user](#) ( rtdm\_user\_info\_t \* user\_info, void \* dst, const void \_\_user \* src, size\_t size )

Copy user-space memory block to specified buffer.

## Parameters

<code>in</code>	<code>user_info</code>	User information pointer as passed to the invoked device operation handler
<code>in</code>	<code>dst</code>	Destination buffer address
<code>in</code>	<code>src</code>	Address of the user-space memory block
<code>in</code>	<code>size</code>	Size of the memory block

## Returns

0 on success, otherwise:

- -EFAULT is returned if an invalid memory area was accessed.

## Note

Before invoking this service, verify via [rtdm\\_read\\_user\\_ok\(\)](#) that the provided user-space address can securely be accessed.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

```
5.16.2.2 int rtdm_copy_to_user ( rtdm_user_info_t * user_info, void __user * dst, const void * src,
                               size_t size )
```

Copy specified buffer to user-space memory block.

## Parameters

<code>in</code>	<code>user_info</code>	User information pointer as passed to the invoked device operation handler
<code>in</code>	<code>dst</code>	Address of the user-space memory block
<code>in</code>	<code>src</code>	Source buffer address
<code>in</code>	<code>size</code>	Size of the memory block

## Returns

0 on success, otherwise:

- -EFAULT is returned if an invalid memory area was accessed.

## Note

Before invoking this service, verify via [rtdm\\_rw\\_user\\_ok\(\)](#) that the provided user-space address can securely be accessed.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.16.2.3 void rtdm\_free ( void \* ptr )

Release real-time memory block.

Parameters

<i>in</i>	<i>ptr</i>	Pointer to memory block as returned by <a href="#">rtdm_malloc()</a>
-----------	------------	--

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine (consider the overhead!)
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.16.2.4 int rtdm\_in\_rt\_context ( void )

Test if running in a real-time task.

Returns

Non-zero is returned if the caller resides in real-time context, 0 otherwise.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.16.2.5 int rtdm\_iomap\_to\_user ( rtdm\_user\_info\_t \* user\_info, phys\_addr\_t src\_addr, size\_t len, int prot, void \*\* pptr, struct vm\_operations\_struct \* vm\_ops, void \* vm\_private\_data )

Map an I/O memory range into the address space of the user.

## Parameters

in	<i>user_info</i>	User information pointer as passed to the invoked device operation handler
in	<i>src_addr</i>	physical I/O address to be mapped
in	<i>len</i>	Length of the memory range
in	<i>prot</i>	Protection flags for the user's memory range, typically either PROT_READ or PROT_READ PROT_WRITE
in,out	<i>pptr</i>	Address of a pointer containing the desired user address or NULL on entry and the finally assigned address on return
in	<i>vm_ops</i>	vm_operations to be executed on the vma_area of the user memory range or NULL
in	<i>vm_private_data</i>	Private data to be stored in the vma_area, primarily useful for vm_operation handlers

## Returns

0 on success, otherwise (most common values):

- -EINVAL is returned if an invalid start address, size, or destination address was passed.
- -ENOMEM is returned if there is insufficient free memory or the limit of memory mapping for the user process was reached.
- -EAGAIN is returned if too much memory has been already locked by the user process.
- -EPERM may be returned if an illegal invocation environment is detected.

## Note

RTDM supports two models for unmapping the user memory range again. One is explicit unmaping via [rtdm\\_munmap\(\)](#), either performed when the user requests it via an IOCTL etc. or when the related device is closed. The other is automatic unmapping, triggered by the user invoking standard munmap() or by the termination of the related process. To track release of the mapping and therefore relinquishment of the referenced physical memory, the caller of [rtdm\\_iomap\\_to\\_user\(\)](#) can pass a `vm_operations_struct` on invocation, defining a close handler for the `vma_area`. See Linux documentaion (e.g. Linux Device Drivers book) on virtual memory management for details.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- User-space task (non-RT)

Rescheduling: possible.

5.16.2.6 void\* rtdm\_malloc ( size\_t size )

Allocate memory block in real-time context.

## Parameters

<i>in</i>	<i>size</i>	Requested size of the memory block
-----------	-------------	------------------------------------

#### Returns

The pointer to the allocated block is returned on success, NULL otherwise.

#### Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine (consider the overhead!)
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

```
5.16.2.7 int rtdm_mmap_to_user ( rtdm_user_info_t * user_info, void * src_addr, size_t len, int prot,
void ** pptr, struct vm_operations_struct * vm_ops, void * vm_private_data )
```

Map a kernel memory range into the address space of the user.

#### Parameters

<i>in</i>	<i>user_info</i>	User information pointer as passed to the invoked device operation handler
<i>in</i>	<i>src_addr</i>	Kernel virtual address to be mapped
<i>in</i>	<i>len</i>	Length of the memory range
<i>in</i>	<i>prot</i>	Protection flags for the user's memory range, typically either PROT_READ or PROT_READ PROT_WRITE
<i>in,out</i>	<i>pptr</i>	Address of a pointer containing the desired user address or NULL on entry and the finally assigned address on return
<i>in</i>	<i>vm_ops</i>	vm_operations to be executed on the vma_area of the user memory range or NULL
<i>in</i>	<i>vm_private_data</i>	Private data to be stored in the vma_area, primarily useful for vm_operation handlers

#### Returns

0 on success, otherwise (most common values):

- -EINVAL is returned if an invalid start address, size, or destination address was passed.
- -ENOMEM is returned if there is insufficient free memory or the limit of memory mapping for the user process was reached.
- -EAGAIN is returned if too much memory has been already locked by the user process.
- -EPERM may be returned if an illegal invocation environment is detected.

## Note

This service only works on memory regions allocated via `kmalloc()` or `vmalloc()`. To map physical I/O memory to user-space use `rtdm_iomap_to_user()` instead.

RTDM supports two models for unmapping the user memory range again. One is explicit unmapping via `rtdm_munmap()`, either performed when the user requests it via an IOCTL etc. or when the related device is closed. The other is automatic unmapping, triggered by the user invoking standard `munmap()` or by the termination of the related process. To track release of the mapping and therefore relinquishment of the referenced physical memory, the caller of `rtdm_mmap_to_user()` can pass a `vm_operations_struct` on invocation, defining a close handler for the `vm_area`. See Linux documentaion (e.g. Linux Device Drivers book) on virtual memory management for details.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- User-space task (non-RT)

Rescheduling: possible.

5.16.2.8 `int rtdm_munmap ( rtdm_user_info_t * user_info, void * ptr, size_t len )`

Unmap a user memory range.

## Parameters

<code>in</code>	<code>user_info</code>	User information pointer as passed to <code>rtdm_mmap_to_user()</code> when requesting to map the memory range
<code>in</code>	<code>ptr</code>	User address or the memory range
<code>in</code>	<code>len</code>	Length of the memory range

## Returns

0 on success, otherwise:

- `-EINVAL` is returned if an invalid address or size was passed.
- `-EPERM` *may* be returned if an illegal invocation environment is detected.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- User-space task (non-RT)

Rescheduling: possible.

5.16.2.9 `void rtdm_printk ( const char * format, ... )`

Real-time safe message printing on kernel console.



## Parameters

<i>in</i>	<i>format</i>	Format string (conforming standard <code>printf()</code> )
	...	Arguments referred by <i>format</i>

## Returns

On success, this service returns the number of characters printed. Otherwise, a negative error code is returned.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine (consider the overhead!)
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never in real-time context, possible in non-real-time environments.

5.16.2.10 `void rtdm_printk_ratelimited ( const char * format, ... )`

Real-time safe rate-limited message printing on kernel console.

## Parameters

<i>in</i>	<i>format</i>	Format string (conforming standard <code>printf()</code> )
	...	Arguments referred by <i>format</i>

## Returns

On success, this service returns the number of characters printed. Otherwise, a negative error code is returned.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Interrupt service routine (consider the overhead!)
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never in real-time context, possible in non-real-time environments.

5.16.2.11 `int rtdm_ratelimit ( struct rtdm_ratelimit_state * rs, const char * func )`

Enforces a rate limit.

This function enforces a rate limit: not more than `->burst` callbacks in every `->interval`.

## Parameters

<code>in, out</code>	<code>rtdm_ratelimit- _state</code>	data
<code>in</code>	<code>name</code>	of calling function

## Returns

0 means callback will be suppressed and 1 means go ahead and do it

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: possible.

References `rtdm_clock_read()`, `rtdm_lock_get_irqsave`, and `rtdm_lock_put_irqrestore`.

```
5.16.2.12 int rtdm_read_user_ok ( rtdm_user_info_t * user_info, const void __user * ptr, size_t size )
```

Check if read access to user-space memory block is safe.

## Parameters

<code>in</code>	<code>user_info</code>	User information pointer as passed to the invoked device operation handler
<code>in</code>	<code>ptr</code>	Address of the user-provided memory block
<code>in</code>	<code>size</code>	Size of the memory block

## Returns

Non-zero is return when it is safe to read from the specified memory block, 0 otherwise.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

```
5.16.2.13 int rtdm_rt_capable ( rtdm_user_info_t * user_info )
```

Test if the caller is capable of running in real-time context.

## Parameters

<i>in</i>	<i>user_info</i>	User information pointer as passed to the invoked device operation handler
-----------	------------------	--

## Returns

Non-zero is returned if the caller is able to execute in real-time context (independent of its current execution mode), 0 otherwise.

## Note

This function can be used by drivers that provide different implementations for the same service depending on the execution mode of the caller. If a caller requests such a service in non-real-time context but is capable of running in real-time as well, it might be appropriate for the driver to reject the request via `-ENOSYS` so that RTDM can switch the caller and restart the request in real-time context.

## Environments:

This service can be called from:

- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

```
5.16.2.14 int rtdm_rw_user_ok ( rtdm_user_info_t * user_info, const void __user * ptr, size_t size )
```

Check if read/write access to user-space memory block is safe.

## Parameters

<i>in</i>	<i>user_info</i>	User information pointer as passed to the invoked device operation handler
<i>in</i>	<i>ptr</i>	Address of the user-provided memory block
<i>in</i>	<i>size</i>	Size of the memory block

## Returns

Non-zero is return when it is safe to read from or write to the specified memory block, 0 otherwise.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

```
5.16.2.15 int rtdm_safe_copy_from_user ( rtdm_user_info_t * user_info, void * dst, const void __user * src, size_t size )
```

Check if read access to user-space memory block and copy it to specified buffer.

## Parameters

<code>in</code>	<code>user_info</code>	User information pointer as passed to the invoked device operation handler
<code>in</code>	<code>dst</code>	Destination buffer address
<code>in</code>	<code>src</code>	Address of the user-space memory block
<code>in</code>	<code>size</code>	Size of the memory block

## Returns

0 on success, otherwise:

- -EFAULT is returned if an invalid memory area was accessed.

## Note

This service is a combination of `rtm_read_user_ok` and `rtm_copy_from_user`.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

```
5.16.2.16 int rtdm_safe_copy_to_user ( rtdm_user_info_t * user_info, void __user * dst, const void *
      src, size_t size )
```

Check if read/write access to user-space memory block is safe and copy specified buffer to it.

## Parameters

<code>in</code>	<code>user_info</code>	User information pointer as passed to the invoked device operation handler
<code>in</code>	<code>dst</code>	Address of the user-space memory block
<code>in</code>	<code>src</code>	Source buffer address
<code>in</code>	<code>size</code>	Size of the memory block

## Returns

0 on success, otherwise:

- -EFAULT is returned if an invalid memory area was accessed.

## Note

This service is a combination of `rtm_rw_user_ok` and `rtm_copy_to_user`.

## Environments:

This service can be called from:

- Kernel module initialization/cleanup code

- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

5.16.2.17 `int rtdm_strncpy_from_user ( rtdm_user_info_t * user_info, char * dst, const char __user * src, size_t count )`

Copy user-space string to specified buffer.

Parameters

<code>in</code>	<code>user_info</code>	User information pointer as passed to the invoked device operation handler
<code>in</code>	<code>dst</code>	Destination buffer address
<code>in</code>	<code>src</code>	Address of the user-space string
<code>in</code>	<code>count</code>	Maximum number of bytes to copy, including the trailing '0'

Returns

Length of the string on success (not including the trailing '0'), otherwise:

- -EFAULT is returned if an invalid memory area was accessed.

Note

This services already includes a check of the source address, calling [rtdm\\_read\\_user\\_ok\(\)](#) for `src` explicitly is not required.

Environments:

This service can be called from:

- Kernel module initialization/cleanup code
- Kernel-based task
- User-space task (RT, non-RT)

Rescheduling: never.

## 5.17 Device Profiles

### Modules

- [CAN Devices](#)
- [Real-time IPC protocols](#)

**Profile Revision: 1**

- [Serial Devices](#)
- [Testing Devices](#)

### Data Structures

- struct [rtdm\\_device\\_info](#)  
*Device information.*

### Typedefs

- typedef struct [rtdm\\_device\\_info](#) [rtdm\\_device\\_info\\_t](#)  
*Device information.*

### RTDM\_CLASS\_XXX

#### Device classes

- #define **RTDM\_CLASS\_PARPORT** 1
- #define **RTDM\_CLASS\_SERIAL** 2
- #define **RTDM\_CLASS\_CAN** 3
- #define **RTDM\_CLASS\_NETWORK** 4
- #define **RTDM\_CLASS\_RTMAC** 5
- #define **RTDM\_CLASS\_TESTING** 6
- #define **RTDM\_CLASS\_RTIPC** 7
- #define **RTDM\_CLASS\_EXPERIMENTAL** 224
- #define **RTDM\_CLASS\_MAX** 255

### Device Naming

Maximum length of device names (excluding the final null character)

- #define **RTDM\_MAX\_DEVNAME\_LEN** 31

### RTDM\_PURGE\_XXX\_BUFFER

Flags selecting buffers to be purged

- #define **RTDM\_PURGE\_RX\_BUFFER** 0x0001
- #define **RTDM\_PURGE\_TX\_BUFFER** 0x0002

## Common IOCTLs

The following IOCTLs are common to all device profiles.

- #define [RTIOC\\_DEVICE\\_INFO](#) \_IOR(RTIOC\_TYPE\_COMMON, 0x00, struct rtdm\_device\_info)  
*Retrieve information about a device or socket.*
- #define [RTIOC\\_PURGE](#) \_IOW(RTIOC\_TYPE\_COMMON, 0x10, int)  
*Purge internal device or socket buffers.*

### 5.17.1 Detailed Description

Device profiles define which operation handlers a driver of a certain class has to implement, which name or protocol it has to register, which IOCTLs it has to provide, and further details. Sub-classes can be defined in order to extend a device profile with more hardware-specific functions.

### 5.17.2 Macro Definition Documentation

5.17.2.1 #define [RTIOC\\_DEVICE\\_INFO](#) \_IOR(RTIOC\_TYPE\_COMMON, 0x00, struct rtdm\_device\_info)

Retrieve information about a device or socket.

Parameters

out	arg	Pointer to information buffer (struct <a href="#">rtdm_device_info</a> )
-----	-----	--

5.17.2.2 #define [RTIOC\\_PURGE](#) \_IOW(RTIOC\_TYPE\_COMMON, 0x10, int)

Purge internal device or socket buffers.

Parameters

in	arg	Purge mask, see <a href="#">RTDM_PURGE_xxx_BUFFER</a>
----	-----	---





# Chapter 6

## Data Structure Documentation

### 6.1 `can_bittime` Struct Reference

Custom CAN bit-time definition.

```
#include <rtcan.h>
```

#### Data Fields

- [`can\_bittime\_type\_t`](#) type  
*Type of bit-time definition.*
- struct [`can\_bittime\_std`](#) `std`  
*Standard bit-time.*
- struct [`can\_bittime\_btr`](#) `btr`  
*Hardware-specific BTR bit-time.*

#### 6.1.1 Detailed Description

Custom CAN bit-time definition.

Examples:

[rtcanconfig.c](#).

The documentation for this struct was generated from the following file:

- `include/rtdm/rtcan.h`

### 6.2 `can_bittime_btr` Struct Reference

Hardware-specific BTR bit-times.

```
#include <rtcan.h>
```

#### Data Fields

- `uint8_t` `btr0`

---

*Bus timing register 0.*

- [uint8\\_t btr1](#)

*Bus timing register 1.*

## 6.2.1 Detailed Description

Hardware-specific BTR bit-times.

The documentation for this struct was generated from the following file:

- [include/rtdm/rtcan.h](#)

## 6.3 can\_bittime\_std Struct Reference

Standard bit-time parameters according to Bosch.

```
#include <rtcan.h>
```

### Data Fields

- [uint32\\_t brp](#)  
*Baud rate prescaler.*
- [uint8\\_t prop\\_seg](#)  
*from 1 to 8*
- [uint8\\_t phase\\_seg1](#)  
*from 1 to 8*
- [uint8\\_t phase\\_seg2](#)  
*from 1 to 8*
- [uint8\\_t sjw:7](#)  
*from 1 to 4*
- [uint8\\_t sam:1](#)  
*1 - enable triple sampling*

## 6.3.1 Detailed Description

Standard bit-time parameters according to Bosch.

The documentation for this struct was generated from the following file:

- [include/rtdm/rtcan.h](#)

## 6.4 can\_filter Struct Reference

Filter for reception of CAN messages.

```
#include <rtcan.h>
```

## Data Fields

- `uint32_t can_id`  
*CAN ID which must match with incoming IDs after passing the mask.*
- `uint32_t can_mask`  
*Mask which is applied to incoming IDs.*

### 6.4.1 Detailed Description

Filter for reception of CAN messages.

This filter works as follows: A received CAN ID is AND'ed bitwise with `can_mask` and then compared to `can_id`. This also includes the `CAN_EFF_FLAG` and `CAN_RTR_FLAG` of `CAN_XXX_FLAG`. If this comparison is true, the message will be received by the socket. The logic can be inverted with the `can_id` flag `CAN_INV_FILTER` :

```
* if (can_id & CAN_INV_FILTER) {
*   if ((received_can_id & can_mask) != (can_id & ~CAN_INV_FILTER))
*     accept-message;
* } else {
*   if ((received_can_id & can_mask) == can_id)
*     accept-message;
* }
*
```

Multiple filters can be arranged in a filter list and set with `Sockopts`. If one of these filters matches a CAN ID upon reception of a CAN frame, this frame is accepted.

Examples:

`rtcan_rtt.c`, and `rtcanrecv.c`.

### 6.4.2 Field Documentation

#### 6.4.2.1 `uint32_t can_filter::can_id`

CAN ID which must match with incoming IDs after passing the mask.

The filter logic can be inverted with the flag `CAN_INV_FILTER`.

Examples:

`rtcanrecv.c`.

#### 6.4.2.2 `uint32_t can_filter::can_mask`

Mask which is applied to incoming IDs.

See [CAN ID masks](#) if exactly one CAN ID should come through.

The documentation for this struct was generated from the following file:

- `include/rtdm/rtcan.h`

## 6.5 `can_frame` Struct Reference

Raw CAN frame.

```
#include <rtcan.h>
```

## Public Member Functions

- `uint8_t data[8] \_\_attribute\_\_ ((aligned(8)))`  
*Payload data bytes.*

## Data Fields

- `can_id_t can_id`  
*CAN ID of the frame.*
- `uint8_t can_dlc`  
*Size of the payload in bytes.*

### 6.5.1 Detailed Description

Raw CAN frame.

Central structure for receiving and sending CAN frames.

Examples:

[rtcan\\_rtt.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

### 6.5.2 Field Documentation

#### 6.5.2.1 `can_id_t can_frame::can_id`

CAN ID of the frame.

See [CAN ID flags](#) for special bits.

Examples:

[rtcan\\_rtt.c](#).

The documentation for this struct was generated from the following file:

- `include/rtdm/rtcan.h`

## 6.6 `rtdm_dev_context` Struct Reference

Device context.

```
#include <rtdm_driver.h>
```

## Data Fields

- unsigned long `context_flags`  
*Context flags, see [Context Flags](#) for details.*
- int `fd`  
*Associated file descriptor.*
- atomic\_t `close_lock_count`  
*Lock counter of context, held while structure is referenced by an operation handler.*

- struct `rtm_operations` \* `ops`  
*Set of active device operation handlers.*
- struct `rtm_device` \* `device`  
*Reference to owning device.*
- struct `rtm_devctx_reserved` `reserved`  
*Data stored by RTDM inside a device context (internal use only)*
- char `dev_private` [0]  
*Begin of driver defined context data structure.*

## 6.6.1 Detailed Description

Device context.

A device context structure is associated with every open device instance. RTDM takes care of its creation and destruction and passes it to the operation handlers when being invoked.

Drivers can attach arbitrary data immediately after the official structure. The size of this data is provided via `rtm_device.context_size` during device registration.

The documentation for this struct was generated from the following file:

- `include/rtm/rtm_driver.h`

## 6.7 rtm\_device Struct Reference

RTDM device.

```
#include <rtm_driver.h>
```

### Data Fields

- int `struct_version`  
*Revision number of this structure, see [Driver Versioning](#) defines.*
- int `device_flags`  
*Device flags, see [Device Flags](#) for details.*
- size\_t `context_size`  
*Size of driver defined appendix to struct `rtm_dev_context`.*
- char `device_name` [RTDM\_MAX\_DEVNAME\_LEN+1]  
*Named device identification (orthogonal to Linux device name space)*
- int `protocol_family`  
*Protocol device identification: protocol family (PF\_XXX)*
- int `socket_type`  
*Protocol device identification: socket type (SOCK\_XXX)*
- `rtm_open_handler_t` `open_rt`  
*Named device instance creation for real-time contexts, optional (but deprecated) if `open_nrt` is non-NULL, ignored for protocol devices.*
- `rtm_open_handler_t` `open_nrt`  
*Named device instance creation for non-real-time contexts, optional if `open_rt` is non-NULL, ignored for protocol devices.*
- `rtm_socket_handler_t` `socket_rt`  
*Protocol socket creation for real-time contexts, optional (but deprecated) if `socket_nrt` is non-NULL, ignored for named devices.*

- `rtm_socket_handler_t socket_nrt`  
*Protocol socket creation for non-real-time contexts, optional if `socket_rt` is non-NULL, ignored for named devices.*
- struct `rtm_operations ops`  
*Default operations on newly opened device instance.*
- int `device_class`  
*Device class ID, see `RTDM_CLASS_xxx`.*
- int `device_sub_class`  
*Device sub-class, see `RTDM_SUBCLASS_xxx` definition in the [Device Profiles](#).*
- int `profile_version`  
*Supported device profile version.*
- const char \* `driver_name`  
*Informational driver name (reported via `/proc`)*
- int `driver_version`  
*Driver version, see [Driver Versioning](#) defines.*
- const char \* `peripheral_name`  
*Informational peripheral name the device is attached to (reported via `/proc`)*
- const char \* `provider_name`  
*Informational driver provider name (reported via `/proc`)*
- const char \* `proc_name`  
*Name of `/proc` entry for the device, must not be NULL.*
- int `device_id`  
*Driver definable device ID.*
- void \* `device_data`  
*Driver definable device data.*
- struct `rtm_dev_reserved reserved`  
*Data stored by RTDM inside a registered device (internal use only)*

## 6.7.1 Detailed Description

RTDM device.

This structure specifies a RTDM device. As some fields, especially the reserved area, will be modified by RTDM during runtime, the structure must not reside in write-protected memory.

## 6.7.2 Field Documentation

### 6.7.2.1 `rtm_open_handler_t rtm_device::open_rt`

Named device instance creation for real-time contexts, optional (but deprecated) if `open_nrt` is non-NULL, ignored for protocol devices.

**Deprecated** Only use non-real-time open handler in new drivers.

Referenced by `rtm_dev_register()`.

### 6.7.2.2 `rtdm_socket_handler_t` `rtdm_device::socket_rt`

Protocol socket creation for real-time contexts, optional (but deprecated) if `socket_nrt` is non-NULL, ignored for named devices.

**Deprecated** Only use non-real-time socket creation handler in new drivers.

Referenced by `rtdm_dev_register()`.

The documentation for this struct was generated from the following file:

- [include/rtdm/rtdm\\_driver.h](#)

## 6.8 `rtdm_device_info` Struct Reference

Device information.

```
#include <rtdm.h>
```

### Data Fields

- int [device\\_flags](#)  
*Device flags, see [Device Flags](#) for details.*
- int [device\\_class](#)  
*Device class ID, see [RTDM\\_CLASS\\_xxx](#).*
- int [device\\_sub\\_class](#)  
*Device sub-class, either [RTDM\\_SUBCLASS\\_GENERIC](#) or a [RTDM\\_SUBCLASS\\_xxx](#) definition of the related [Device Profile](#).*
- int [profile\\_version](#)  
*Supported device profile version.*

### 6.8.1 Detailed Description

Device information.

The documentation for this struct was generated from the following file:

- [include/rtdm/rtdm.h](#)

## 6.9 `rtdm_operations` Struct Reference

Device operations.

```
#include <rtdm_driver.h>
```

### Data Fields

#### Common Operations

- [rtdm\\_close\\_handler\\_t close\\_rt](#)  
*Close handler for real-time contexts (optional, deprecated)*
- [rtdm\\_close\\_handler\\_t close\\_nrt](#)

- [rtm\\_close\\_handler\\_t rtm\\_close](#)  
*Close handler for non-real-time contexts (required)*
- [rtm\\_ioctl\\_handler\\_t rtm\\_ioctl\\_rt](#)  
*IOCTL from real-time context (optional)*
- [rtm\\_ioctl\\_handler\\_t rtm\\_ioctl\\_nrt](#)  
*IOCTL from non-real-time context (optional)*
- [rtm\\_select\\_bind\\_handler\\_t rtm\\_select\\_bind](#)  
*Select binding handler for any context (optional)*

### Stream-Oriented Device Operations

- [rtm\\_read\\_handler\\_t rtm\\_read\\_rt](#)  
*Read handler for real-time context (optional)*
- [rtm\\_read\\_handler\\_t rtm\\_read\\_nrt](#)  
*Read handler for non-real-time context (optional)*
- [rtm\\_write\\_handler\\_t rtm\\_write\\_rt](#)  
*Write handler for real-time context (optional)*
- [rtm\\_write\\_handler\\_t rtm\\_write\\_nrt](#)  
*Write handler for non-real-time context (optional)*

### Message-Oriented Device Operations

- [rtm\\_recvmsg\\_handler\\_t rtm\\_recvmsg\\_rt](#)  
*Receive message handler for real-time context (optional)*
- [rtm\\_recvmsg\\_handler\\_t rtm\\_recvmsg\\_nrt](#)  
*Receive message handler for non-real-time context (optional)*
- [rtm\\_sendmsg\\_handler\\_t rtm\\_sendmsg\\_rt](#)  
*Transmit message handler for real-time context (optional)*
- [rtm\\_sendmsg\\_handler\\_t rtm\\_sendmsg\\_nrt](#)  
*Transmit message handler for non-real-time context (optional)*

## 6.9.1 Detailed Description

Device operations.

## 6.9.2 Field Documentation

### 6.9.2.1 [rtm\\_close\\_handler\\_t rtm\\_operations::close\\_rt](#)

Close handler for real-time contexts (optional, deprecated)

**Deprecated** Only use non-real-time close handler in new drivers.

Referenced by [rtm\\_dev\\_register\(\)](#).

The documentation for this struct was generated from the following file:

- [include/rtdm/rtdm\\_driver.h](#)

## 6.10 [rtipc\\_port\\_label](#) Struct Reference

Port label information structure.

```
#include <rtipc.h>
```



## Data Fields

- char [label](#) [XNOBJECT\_NAME\_LEN]  
*Port label string, null-terminated.*

### 6.10.1 Detailed Description

Port label information structure.

Examples:

[bufp-label.c](#), [iddp-label.c](#), and [xddp-label.c](#).

### 6.10.2 Field Documentation

#### 6.10.2.1 char `rtipc_port_label::label`[XNOBJECT\_NAME\_LEN]

Port label string, null-terminated.

The documentation for this struct was generated from the following file:

- [include/rtdm/rtipc.h](#)

## 6.11 `rtser_config` Struct Reference

Serial device configuration.

```
#include <rtserial.h>
```

## Data Fields

- int [config\\_mask](#)  
*mask specifying valid fields, see [RTSER\\_SET\\_xxx](#)*
- int [baud\\_rate](#)  
*baud rate, default [RTSER\\_DEF\\_BAUD](#)*
- int [parity](#)  
*number of parity bits, see [RTSER\\_xxx\\_PARITY](#)*
- int [data\\_bits](#)  
*number of data bits, see [RTSER\\_xxx\\_BITS](#)*
- int [stop\\_bits](#)  
*number of stop bits, see [RTSER\\_xxx\\_STOPB](#)*
- int [handshake](#)  
*handshake mechanisms, see [RTSER\\_xxx\\_HAND](#)*
- int [fifo\\_depth](#)  
*reception FIFO interrupt threshold, see [RTSER\\_FIFO\\_xxx](#)*
- [nanosecs\\_rel\\_t rx\\_timeout](#)  
*reception timeout, see [RTSER\\_TIMEOUT\\_xxx](#) for special values*
- [nanosecs\\_rel\\_t tx\\_timeout](#)  
*transmission timeout, see [RTSER\\_TIMEOUT\\_xxx](#) for special values*
- [nanosecs\\_rel\\_t event\\_timeout](#)  
*event timeout, see [RTSER\\_TIMEOUT\\_xxx](#) for special values*

- int [timestamp\\_history](#)  
*enable timestamp history, see [RTSER\\_XXX\\_TIMESTAMP\\_HISTORY](#)*
- int [event\\_mask](#)  
*event mask to be used with [RTSER\\_RTIOC\\_WAIT\\_EVENT](#), see [RTSER\\_EVENT\\_XXX](#)*
- int [rs485](#)  
*enable RS485 mode, see [RTSER\\_RS485\\_XXX](#)*

### 6.11.1 Detailed Description

Serial device configuration.

Examples:

[cross-link.c](#).

The documentation for this struct was generated from the following file:

- [include/rtdm/rtserial.h](#)

## 6.12 rtser\_event Struct Reference

Additional information about serial device events.

```
#include <rtserial.h>
```

### Data Fields

- int [events](#)  
*signalled events, see [RTSER\\_EVENT\\_XXX](#)*
- int [rx\\_pending](#)  
*number of pending input characters*
- [nanosecs\\_abs\\_t last\\_timestamp](#)  
*last interrupt timestamp*
- [nanosecs\\_abs\\_t rxpend\\_timestamp](#)  
*reception timestamp of oldest character in input queue*

### 6.12.1 Detailed Description

Additional information about serial device events.

Examples:

[cross-link.c](#).

The documentation for this struct was generated from the following file:

- [include/rtdm/rtserial.h](#)

## 6.13 rtser\_status Struct Reference

Serial device status.

```
#include <rtserial.h>
```

## Data Fields

- int [line\\_status](#)  
*line status register, see [RTSER\\_LSR\\_xxx](#)*
- int [modem\\_status](#)  
*modem status register, see [RTSER\\_MSR\\_xxx](#)*

### 6.13.1 Detailed Description

Serial device status.

The documentation for this struct was generated from the following file:

- [include/rtdm/rtserial.h](#)

## 6.14 sockaddr\_can Struct Reference

Socket address structure for the CAN address family.

```
#include <rtcan.h>
```

## Data Fields

- sa\_family\_t [can\\_family](#)  
*CAN address family, must be `AF_CAN`.*
- int [can\\_ifindex](#)  
*Interface index of CAN controller.*

### 6.14.1 Detailed Description

Socket address structure for the CAN address family.

Examples:

[rtcan\\_rtt.c](#), [rtcanrecv.c](#), and [rtcansend.c](#).

### 6.14.2 Field Documentation

#### 6.14.2.1 int sockaddr\_can::can\_ifindex

Interface index of CAN controller.

See [SIOCGIFINDEX](#).

The documentation for this struct was generated from the following file:

- [include/rtdm/rtcan.h](#)

## 6.15 sockaddr\_ipc Struct Reference

Socket address structure for the RTIPC address family.

```
#include <rtipc.h>
```

## Data Fields

- [sa\\_family\\_t](#) [sipc\\_family](#)  
*RTIPC address family, must be AF\_RTIPC.*
- [rtipc\\_port\\_t](#) [sipc\\_port](#)  
*Port number.*

### 6.15.1 Detailed Description

Socket address structure for the RTIPC address family.

Examples:

[bufp-label.c](#), [bufp-readwrite.c](#), [iddp-label.c](#), [iddp-sendrecv.c](#), [xddp-echo.c](#), [xddp-label.c](#), and [xddp-stream.c](#).

### 6.15.2 Field Documentation

#### 6.15.2.1 [rtipc\\_port\\_t](#) `sockaddr_ipc::sipc_port`

Port number.

The documentation for this struct was generated from the following file:

- [include/rtdm/rtipc.h](#)

# Chapter 7

## File Documentation

### 7.1 include/rtdm/rtdm.h File Reference

Real-Time Driver Model for RT-Socket-CAN, CAN device profile header.

```
#include <linux/net.h>
#include <linux/socket.h>
#include <linux/if.h>
#include <rtdm/rtdm.h>
```

#### Data Structures

- struct [can\\_bittime\\_std](#)  
*Standard bit-time parameters according to Bosch.*
- struct [can\\_bittime\\_btr](#)  
*Hardware-specific BTR bit-times.*
- struct [can\\_bittime](#)  
*Custom CAN bit-time definition.*
- struct [can\\_filter](#)  
*Filter for reception of CAN messages.*
- struct [sockaddr\\_can](#)  
*Socket address structure for the CAN address family.*
- struct [can\\_frame](#)  
*Raw CAN frame.*

#### Macros

- #define [AF\\_CAN](#) 29  
*CAN address family.*
- #define [PF\\_CAN](#) [AF\\_CAN](#)  
*CAN protocol family.*
- #define [SOL\\_CAN\\_RAW](#) 103  
*CAN socket levels.*

#### **CAN ID masks**

*Bit masks for masking CAN IDs*

- #define `CAN_EFF_MASK` 0x1FFFFFFF  
*Bit mask for extended CAN IDs.*
- #define `CAN_SFF_MASK` 0x000007FF  
*Bit mask for standard CAN IDs.*

## CAN ID flags

Flags within a CAN ID indicating special CAN frame attributes

- #define `CAN_EFF_FLAG` 0x80000000  
*Extended frame.*
- #define `CAN_RTR_FLAG` 0x40000000  
*Remote transmission frame.*
- #define `CAN_ERR_FLAG` 0x20000000  
*Error frame (see [Errors](#)), not valid in struct `can_filter`.*
- #define `CAN_INV_FILTER CAN_ERR_FLAG`  
*Invert CAN filter definition, only valid in struct `can_filter`.*

## Particular CAN protocols

Possible protocols for the `PF_CAN` protocol family

Currently only the RAW protocol is supported.

- #define `CAN_RAW` 1  
*Raw protocol of `PF_CAN`, applicable to socket type `SOCK_RAW`.*

## CAN controller modes

Special CAN controllers modes, which can be or'ed together.

Note

*These modes are hardware-dependent. Please consult the hardware manual of the CAN controller for more detailed information.*

- #define `CAN_CTRLMODE_LISTENONLY` 0x1
- #define `CAN_CTRLMODE_LOOPBACK` 0x2
- #define `CAN_CTRLMODE_3_SAMPLES` 0x4

## Timestamp switches

Arguments to pass to `RTCAN_RTIOC_TAKE_TIMESTAMP`

- #define `RTCAN_TAKE_NO_TIMESTAMPS` 0  
*Switch off taking timestamps.*
- #define `RTCAN_TAKE_TIMESTAMPS` 1  
*Do take timestamps.*

## RAW socket options

Setting and getting CAN RAW socket options.

- #define `CAN_RAW_FILTER` 0x1  
*CAN filter definition.*
- #define `CAN_RAW_ERR_FILTER` 0x2  
*CAN error mask.*
- #define `CAN_RAW_LOOPBACK` 0x3  
*CAN TX loopback.*
- #define `CAN_RAW_RECV_OWN_MSGS` 0x4  
*CAN receive own messages.*

## IOCTLs

### CAN device IOCTLs

- #define [SIOCGIFINDEX](#) defined\_by\_kernel\_header\_file  
Get CAN interface index by name.
- #define [SIOCSCANBAUDRATE](#) \_IOW(RTIOC\_TYPE\_CAN, 0x01, struct ifreq)  
Set baud rate.
- #define [SIOCGCANBAUDRATE](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x02, struct ifreq)  
Get baud rate.
- #define [SIOCSCANCUSTOMBITTIME](#) \_IOW(RTIOC\_TYPE\_CAN, 0x03, struct ifreq)  
Set custom bit time parameter.
- #define [SIOCGCANCUSTOMBITTIME](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x04, struct ifreq)  
Get custom bit-time parameters.
- #define [SIOCSCANMODE](#) \_IOW(RTIOC\_TYPE\_CAN, 0x05, struct ifreq)  
Set operation mode of CAN controller.
- #define [SIOCGCANSTATE](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x06, struct ifreq)  
Get current state of CAN controller.
- #define [SIOCSCANCTRLMODE](#) \_IOW(RTIOC\_TYPE\_CAN, 0x07, struct ifreq)  
Set special controller modes.
- #define [SIOCGCANCTRLMODE](#) \_IOWR(RTIOC\_TYPE\_CAN, 0x08, struct ifreq)  
Get special controller modes.
- #define [RTCAN\\_RTIOC\\_TAKE\\_TIMESTAMP](#) \_IOW(RTIOC\_TYPE\_CAN, 0x09, int)  
Enable or disable storing a high precision timestamp upon reception of a CAN frame.
- #define [RTCAN\\_RTIOC\\_RCV\\_TIMEOUT](#) \_IOW(RTIOC\_TYPE\_CAN, 0x0A, nanosecs\_rel\_t)  
Specify a reception timeout for a socket.
- #define [RTCAN\\_RTIOC\\_SND\\_TIMEOUT](#) \_IOW(RTIOC\_TYPE\_CAN, 0x0B, nanosecs\_rel\_t)  
Specify a transmission timeout for a socket.

### Error mask

Error class (mask) in `can_id` field of struct [can\\_frame](#) to be used with [CAN\\_RAW\\_ERR\\_FILTER](#).

**Note:** Error reporting is hardware dependent and most CAN controllers report less detailed error conditions than the SJA1000.

**Note:** In case of a bus-off error condition ([CAN\\_ERR\\_BUSOFF](#)), the CAN controller is **not** restarted automatically. It is the application's responsibility to react appropriately, e.g. calling [CAN\\_MODE\\_START](#).

**Note:** Bus error interrupts ([CAN\\_ERR\\_BUSERROR](#)) are enabled when an application is calling a [Recv](#) function on a socket listening on bus errors (using [CAN\\_RAW\\_ERR\\_FILTER](#)). After one bus error has occurred, the interrupt will be disabled to allow the application time for error processing and to efficiently avoid bus error interrupt flooding.

- #define [CAN\\_ERR\\_TX\\_TIMEOUT](#) 0x00000001U  
TX timeout (netdevice driver)
- #define [CAN\\_ERR\\_LOSTARB](#) 0x00000002U  
Lost arbitration (see [data\[0\]](#))
- #define [CAN\\_ERR\\_CRTL](#) 0x00000004U  
Controller problems (see [data\[1\]](#))
- #define [CAN\\_ERR\\_PROT](#) 0x00000008U  
Protocol violations (see [data\[2\]](#), [data\[3\]](#))
- #define [CAN\\_ERR\\_TRX](#) 0x00000010U  
Transceiver status (see [data\[4\]](#))
- #define [CAN\\_ERR\\_ACK](#) 0x00000020U  
Received no ACK on transmission.
- #define [CAN\\_ERR\\_BUSOFF](#) 0x00000040U  
Bus off.
- #define [CAN\\_ERR\\_BUSERROR](#) 0x00000080U  
Bus error (may flood!)

- #define `CAN_ERR_RESTARTED` 0x0000100U  
*Controller restarted.*
- #define `CAN_ERR_MASK` 0x1FFFFFFFU  
*Omit EFF, RTR, ERR flags.*

### Arbitration lost error

Error in the `data[0]` field of struct `can_frame`.

- #define `CAN_ERR_LOSTARB_UNSPEC` 0x00  
*unspecified*

### Controller problems

Error in the `data[1]` field of struct `can_frame`.

- #define `CAN_ERR_CRTL_UNSPEC` 0x00  
*unspecified*
- #define `CAN_ERR_CRTL_RX_OVERFLOW` 0x01  
*RX buffer overflow.*
- #define `CAN_ERR_CRTL_TX_OVERFLOW` 0x02  
*TX buffer overflow.*
- #define `CAN_ERR_CRTL_RX_WARNING` 0x04  
*reached warning level for RX errors*
- #define `CAN_ERR_CRTL_TX_WARNING` 0x08  
*reached warning level for TX errors*
- #define `CAN_ERR_CRTL_RX_PASSIVE` 0x10  
*reached passive level for RX errors*
- #define `CAN_ERR_CRTL_TX_PASSIVE` 0x20  
*reached passive level for TX errors*

### Protocol error type

Error in the `data[2]` field of struct `can_frame`.

- #define `CAN_ERR_PROT_UNSPEC` 0x00  
*unspecified*
- #define `CAN_ERR_PROT_BIT` 0x01  
*single bit error*
- #define `CAN_ERR_PROT_FORM` 0x02  
*frame format error*
- #define `CAN_ERR_PROT_STUFF` 0x04  
*bit stuffing error*
- #define `CAN_ERR_PROT_BIT0` 0x08  
*unable to send dominant bit*
- #define `CAN_ERR_PROT_BIT1` 0x10  
*unable to send recessive bit*
- #define `CAN_ERR_PROT_OVERLOAD` 0x20  
*bus overload*
- #define `CAN_ERR_PROT_ACTIVE` 0x40  
*active error announcement*
- #define `CAN_ERR_PROT_TX` 0x80  
*error occurred on transmission*

### Protocol error location

Error in the `data[4]` field of struct `can_frame`.

- #define `CAN_ERR_PROT_LOC_UNSPEC` 0x00



- *unspecified*  
#define CAN\_ERR\_PROT\_LOC\_SOF 0x03
- *start of frame*  
#define CAN\_ERR\_PROT\_LOC\_ID28\_21 0x02  
*ID bits 28 - 21 (SFF: 10 - 3)*
- #define CAN\_ERR\_PROT\_LOC\_ID20\_18 0x06  
*ID bits 20 - 18 (SFF: 2 - 0)*
- #define CAN\_ERR\_PROT\_LOC\_SRTR 0x04  
*substitute RTR (SFF: RTR)*
- #define CAN\_ERR\_PROT\_LOC\_IDE 0x05  
*identifier extension*
- #define CAN\_ERR\_PROT\_LOC\_ID17\_13 0x07  
*ID bits 17-13.*
- #define CAN\_ERR\_PROT\_LOC\_ID12\_05 0x0F  
*ID bits 12-5.*
- #define CAN\_ERR\_PROT\_LOC\_ID04\_00 0x0E  
*ID bits 4-0.*
- #define CAN\_ERR\_PROT\_LOC\_RTR 0x0C  
*RTR.*
- #define CAN\_ERR\_PROT\_LOC\_RES1 0x0D  
*reserved bit 1*
- #define CAN\_ERR\_PROT\_LOC\_RES0 0x09  
*reserved bit 0*
- #define CAN\_ERR\_PROT\_LOC\_DLC 0x0B  
*data length code*
- #define CAN\_ERR\_PROT\_LOC\_DATA 0x0A  
*data section*
- #define CAN\_ERR\_PROT\_LOC\_CRC\_SEQ 0x08  
*CRC sequence.*
- #define CAN\_ERR\_PROT\_LOC\_CRC\_DEL 0x18  
*CRC delimiter.*
- #define CAN\_ERR\_PROT\_LOC\_ACK 0x19  
*ACK slot.*
- #define CAN\_ERR\_PROT\_LOC\_ACK\_DEL 0x1B  
*ACK delimiter.*
- #define CAN\_ERR\_PROT\_LOC\_EOF 0x1A  
*end of frame*
- #define CAN\_ERR\_PROT\_LOC\_INTERM 0x12  
*intermission*
- #define CAN\_ERR\_TRX\_UNSPEC 0x00  
*0000 0000*
- #define CAN\_ERR\_TRX\_CANH\_NO\_WIRE 0x04  
*0000 0100*
- #define CAN\_ERR\_TRX\_CANH\_SHORT\_TO\_BAT 0x05  
*0000 0101*
- #define CAN\_ERR\_TRX\_CANH\_SHORT\_TO\_VCC 0x06  
*0000 0110*
- #define CAN\_ERR\_TRX\_CANH\_SHORT\_TO\_GND 0x07  
*0000 0111*
- #define CAN\_ERR\_TRX\_CANL\_NO\_WIRE 0x40  
*0100 0000*
- #define CAN\_ERR\_TRX\_CANL\_SHORT\_TO\_BAT 0x50  
*0101 0000*
- #define CAN\_ERR\_TRX\_CANL\_SHORT\_TO\_VCC 0x60  
*0110 0000*
- #define CAN\_ERR\_TRX\_CANL\_SHORT\_TO\_GND 0x70  
*0111 0000*
- #define CAN\_ERR\_TRX\_CANL\_SHORT\_TO\_CANH 0x80  
*1000 0000*

## Typedefs

- typedef uint32\_t `can_id_t`  
*Type of CAN id (see `CAN_xxx_MASK` and `CAN_xxx_FLAG`)*
- typedef `can_id_t` `can_err_mask_t`  
*Type of CAN error mask.*
- typedef uint32\_t `can_baudrate_t`  
*Baudrate definition in bits per second.*
- typedef enum `CAN_BITTIME_TYPE` `can_bittime_type_t`  
*See `CAN_BITTIME_TYPE`.*
- typedef enum `CAN_MODE` `can_mode_t`  
*See `CAN_MODE`.*
- typedef int `can_ctrlmode_t`  
*See `CAN_CTRLMODE`.*
- typedef enum `CAN_STATE` `can_state_t`  
*See `CAN_STATE`.*
- typedef struct `can_filter` `can_filter_t`  
*Filter for reception of CAN messages.*
- typedef struct `can_frame` `can_frame_t`  
*Raw CAN frame.*

## Enumerations

- enum `CAN_BITTIME_TYPE` { `CAN_BITTIME_STD`, `CAN_BITTIME_BTR` }  
*Supported CAN bit-time types.*

### **CAN operation modes**

*Modes into which CAN controllers can be set*

- enum `CAN_MODE` { `CAN_MODE_STOP` = 0, `CAN_MODE_START`, `CAN_MODE_SLEEP` }

### **CAN controller states**

*States a CAN controller can be in.*

- enum `CAN_STATE` {  
`CAN_STATE_ERROR_ACTIVE` = 0, `CAN_STATE_ERROR_WARNING` = 1, `CAN_STATE_ERROR_PASSIVE`  
= 2, `CAN_STATE_BUS_OFF`,  
`CAN_STATE_SCANNING_BAUDRATE`, `CAN_STATE_STOPPED`, `CAN_STATE_SLEEPING`  
}

### 7.1.1 Detailed Description

Real-Time Driver Model for RT-Socket-CAN, CAN device profile header.

Note

Copyright (C) 2006 Wolfgang Grandegger [wg@grandegger.com](mailto:wg@grandegger.com)

Copyright (C) 2005, 2006 Sebastian Smolorz [Sebastian.Smolorz@stud.uni-hannover.de](mailto:Sebastian.Smolorz@stud.uni-hannover.de)

This RTDM CAN device profile header is based on:

include/linux/can.h, include/linux/socket.h, net/can/pf\_can.h in linux-can.patch, a CAN socket framework for Linux

Copyright (C) 2004, 2005, Robert Schwebel, Benedikt Spranger, Marc Kleine-Budde, Pengutronix

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU

General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 7.2 include/rtdm/rtdm.h File Reference

Real-Time Driver Model for Xenomai, user API header.

```
#include <linux/types.h>
#include <linux/fcntl.h>
#include <linux/ioctl.h>
#include <linux/sched.h>
#include <linux/socket.h>
#include <asm/xenomai/wrappers.h>
```

### Data Structures

- struct [rtdm\\_device\\_info](#)  
*Device information.*

### Macros

#### API Versioning

- #define [RTDM\\_API\\_VER](#) 8  
*Common user and driver API version.*
- #define [RTDM\\_API\\_MIN\\_COMPAT\\_VER](#) 6  
*Minimum API revision compatible with the current release.*

#### RTDM\_TIMEOUT\_XXX

*Special timeout values*

- #define [RTDM\\_TIMEOUT\\_INFINITE](#) 0  
*Block forever.*
- #define [RTDM\\_TIMEOUT\\_NONE](#) (-1)

*Any negative timeout means non-blocking.*

## RTDM\_CLASS\_xxx

*Device classes*

- #define RTDM\_CLASS\_PARPORT 1
- #define RTDM\_CLASS\_SERIAL 2
- #define RTDM\_CLASS\_CAN 3
- #define RTDM\_CLASS\_NETWORK 4
- #define RTDM\_CLASS\_RTMAC 5
- #define RTDM\_CLASS\_TESTING 6
- #define RTDM\_CLASS\_RTIPC 7
- #define RTDM\_CLASS\_EXPERIMENTAL 224
- #define RTDM\_CLASS\_MAX 255

## Device Naming

*Maximum length of device names (excluding the final null character)*

- #define RTDM\_MAX\_DEVNAME\_LEN 31

## RTDM\_PURGE\_xxx\_BUFFER

*Flags selecting buffers to be purged*

- #define RTDM\_PURGE\_RX\_BUFFER 0x0001
- #define RTDM\_PURGE\_TX\_BUFFER 0x0002

## Common IOCTLs

*The following IOCTLs are common to all device profiles.*

- #define RTIOC\_DEVICE\_INFO \_IOR(RTIOC\_TYPE\_COMMON, 0x00, struct rtdm\_device\_info)  
*Retrieve information about a device or socket.*
- #define RTIOC\_PURGE \_IOW(RTIOC\_TYPE\_COMMON, 0x10, int)  
*Purge internal device or socket buffers.*

## Typedefs

- typedef uint64\_t nanosecs\_abs\_t  
*RTDM type for representing absolute dates.*
- typedef int64\_t nanosecs\_rel\_t  
*RTDM type for representing relative intervals.*
- typedef struct rtdm\_device\_info rtdm\_device\_info\_t  
*Device information.*

### 7.2.1 Detailed Description

Real-Time Driver Model for Xenomai, user API header.

Note

Copyright (C) 2005, 2006 Jan Kiszka [jan.kiszka@web.de](mailto:jan.kiszka@web.de)  
Copyright (C) 2005 Joerg Langenberg [joerg.langenberg@gmx.net](mailto:joerg.langenberg@gmx.net)

Xenomai is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

Xenomai is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with Xenomai; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 7.3 include/rtdm/rtdm\_driver.h File Reference

Real-Time Driver Model for Xenomai, driver API header.

```
#include <asm/atomic.h>
#include <linux/list.h>
#include <nucleus/xenomai.h>
#include <nucleus/heap.h>
#include <nucleus/pod.h>
#include <nucleus/synch.h>
#include <nucleus/select.h>
#include <nucleus/vfile.h>
#include <rtdm/rtdm.h>
#include <nucleus/assert.h>
```

### Data Structures

- struct [rtdm\\_operations](#)  
*Device operations.*
- struct [rtdm\\_dev\\_context](#)  
*Device context.*
- struct [rtdm\\_device](#)  
*RTDM device.*

### Macros

- #define [rtdm\\_irq\\_get\\_arg](#)(irq\_handle, type) ((type \*)irq\_handle->cookie)  
*Retrieve IRQ handler argument.*

### Device Flags

*Static flags describing a RTDM device*

- #define [RTDM\\_EXCLUSIVE](#) 0x0001  
*If set, only a single instance of the device can be requested by an application.*
- #define [RTDM\\_NAMED\\_DEVICE](#) 0x0010  
*If set, the device is addressed via a clear-text name.*
- #define [RTDM\\_PROTOCOL\\_DEVICE](#) 0x0020  
*If set, the device is addressed via a combination of protocol ID and socket type.*

- #define `RTDM_DEVICE_TYPE_MASK` 0x00F0  
Mask selecting the device type.

## Context Flags

Dynamic flags describing the state of an open RTDM device (bit numbers)

- #define `RTDM_CREATED_IN_NRT` 0  
Set by RTDM if the device instance was created in non-real-time context.
- #define `RTDM_CLOSING` 1  
Set by RTDM when the device is being closed.
- #define `RTDM_USER_CONTEXT_FLAG` 8 /\* first user-definable flag \*/  
Lowest bit number the driver developer can use freely.

## Driver Versioning

Current revisions of RTDM structures, encoding of driver versions. See [API Versioning](#) for the interface revision.

- #define `RTDM_DEVICE_STRUCT_VER` 5  
Version of struct `rtdm_device`.
- #define `RTDM_CONTEXT_STRUCT_VER` 3  
Version of struct `rtdm_dev_context`.
- #define `RTDM_SECURE_DEVICE` 0x80000000  
Flag indicating a secure variant of RTDM (not supported here)
- #define `RTDM_DRIVER_VER`(major, minor, patch) (((major & 0xFF) << 16) | ((minor & 0xFF) << 8) | (patch & 0xFF))  
Version code constructor for driver revisions.
- #define `RTDM_DRIVER_MAJOR_VER`(ver) (((ver) >> 16) & 0xFF)  
Get major version number from driver revision code.
- #define `RTDM_DRIVER_MINOR_VER`(ver) (((ver) >> 8) & 0xFF)  
Get minor version number from driver revision code.
- #define `RTDM_DRIVER_PATCH_VER`(ver) ((ver) & 0xFF)  
Get patch version number from driver revision code.

## Global Lock across Scheduler Invocation

- #define `RTDM_EXECUTE_ATOMICALLY`(code\_block)  
Execute code block atomically.

## RTDM\_IRQTYPE\_xxx

Interrupt registrations flags

- #define `RTDM_IRQTYPE_SHARED` XN\_ISR\_SHARED  
Enable IRQ-sharing with other real-time drivers.
- #define `RTDM_IRQTYPE_EDGE` XN\_ISR\_EDGE  
Mark IRQ as edge-triggered, relevant for correct handling of shared edge-triggered IRQs.

## RTDM\_IRQ\_xxx

Return flags of interrupt handlers

- #define `RTDM_IRQ_NONE` XN\_ISR\_NONE  
Unhandled interrupt.
- #define `RTDM_IRQ_HANDLED` XN\_ISR\_HANDLED  
Denote handled interrupt.

## Task Priority Range

Maximum and minimum task priorities

- #define **RTDM\_TASK\_LOWEST\_PRIORITY** XNSCHED\_LOW\_PRIO
- #define **RTDM\_TASK\_HIGHEST\_PRIORITY** XNSCHED\_HIGH\_PRIO

## Task Priority Modification

Raise or lower task priorities by one level

- #define **RTDM\_TASK\_RAISE\_PRIORITY** (+1)
- #define **RTDM\_TASK\_LOWER\_PRIORITY** (-1)

## Typedefs

- typedef int(\* **rtdm\_irq\_handler\_t**)(rtdm\_irq\_t \*irq\_handle)  
*Interrupt handler.*
- typedef void(\* **rtdm\_nrtsig\_handler\_t**)(rtdm\_nrtsig\_t nrt\_sig, void \*arg)  
*Non-real-time signal handler.*
- typedef void(\* **rtdm\_timer\_handler\_t**)(rtdm\_timer\_t \*timer)  
*Timer handler.*
- typedef void(\* **rtdm\_task\_proc\_t**)(void \*arg)  
*Real-time task procedure.*

## Operation Handler Prototypes

- typedef int(\* **rtdm\_open\_handler\_t**)(struct **rtdm\_dev\_context** \*context, rtdm\_user\_info\_t \*user\_info, int oflag)  
*Named device open handler.*
- typedef int(\* **rtdm\_socket\_handler\_t**)(struct **rtdm\_dev\_context** \*context, rtdm\_user\_info\_t \*user\_info, int protocol)  
*Socket creation handler for protocol devices.*
- typedef int(\* **rtdm\_close\_handler\_t**)(struct **rtdm\_dev\_context** \*context, rtdm\_user\_info\_t \*user\_info)  
*Close handler.*
- typedef int(\* **rtdm\_ioctl\_handler\_t**)(struct **rtdm\_dev\_context** \*context, rtdm\_user\_info\_t \*user\_info, unsigned int request, void \_\_user \*arg)  
*IOCTL handler.*
- typedef int(\* **rtdm\_select\_bind\_handler\_t**)(struct **rtdm\_dev\_context** \*context, rtdm\_selector\_t \*selector, enum **rtdm\_selecttype** type, unsigned fd\_index)  
*Select binding handler.*
- typedef ssize\_t(\* **rtdm\_read\_handler\_t**)(struct **rtdm\_dev\_context** \*context, rtdm\_user\_info\_t \*user\_info, void \*buf, size\_t nbyte)  
*Read handler.*
- typedef ssize\_t(\* **rtdm\_write\_handler\_t**)(struct **rtdm\_dev\_context** \*context, rtdm\_user\_info\_t \*user\_info, const void \*buf, size\_t nbyte)  
*Write handler.*
- typedef ssize\_t(\* **rtdm\_rcvmsg\_handler\_t**)(struct **rtdm\_dev\_context** \*context, rtdm\_user\_info\_t \*user\_info, struct user\_msghdr \*msg, int flags)  
*Receive message handler.*
- typedef ssize\_t(\* **rtdm\_sendmsg\_handler\_t**)(struct **rtdm\_dev\_context** \*context, rtdm\_user\_info\_t \*user\_info, const struct user\_msghdr \*msg, int flags)  
*Transmit message handler.*

## Enumerations

### **RTDM\_SELECTTYPE\_XXX**

*Event types select can bind to*

- enum `rtm_selecttype` { `RTDM_SELECTTYPE_READ` = `XNSELECT_READ`, `RTDM_SELECTTYPE_WRITE` = `XNSELECT_WRITE`, `RTDM_SELECTTYPE_EXCEPT` = `XNSELECT_EXCEPT` }

## RTDM\_TIMERMODE\_xxx

*Timer operation modes*

- enum `rtm_timer_mode` { `RTDM_TIMERMODE_RELATIVE` = `XN_RELATIVE`, `RTDM_TIMERMODE_ABSOLUTE` = `XN_ABSOLUTE`, `RTDM_TIMERMODE_REALTIME` = `XN_REALTIME` }

## Functions

- static void \* `rtm_context_to_private` (struct `rtm_dev_context` \*context)  
*Locate the driver private area associated to a device context structure.*
- static struct `rtm_dev_context` \* `rtm_private_to_context` (void \*dev\_private)  
*Locate a device context structure from its driver private area.*
- int `rtm_dev_register` (struct `rtm_device` \*device)  
*Register a RTDM device.*
- int `rtm_dev_unregister` (struct `rtm_device` \*device, unsigned int poll\_delay)  
*Unregisters a RTDM device.*
- struct `rtm_dev_context` \* `rtm_context_get` (int fd)  
*Retrieve and lock a device context.*
- int `rtm_select_bind` (int fd, `rtm_selector_t` \*selector, enum `rtm_selecttype` type, unsigned fd\_index)  
*Bind a selector to specified event types of a given file descriptor.*
- int `rtm_irq_request` (`rtm_irq_t` \*irq\_handle, unsigned int irq\_no, `rtm_irq_handler_t` handler, unsigned long flags, const char \*device\_name, void \*arg)  
*Register an interrupt handler.*
- void `rtm_timer_destroy` (`rtm_timer_t` \*timer)  
*Destroy a timer.*
- int `rtm_timer_start` (`rtm_timer_t` \*timer, `nanosecs_abs_t` expiry, `nanosecs_rel_t` interval, enum `rtm_timer_mode` mode)  
*Start a timer.*
- void `rtm_timer_stop` (`rtm_timer_t` \*timer)  
*Stop a timer.*
- int `rtm_task_init` (`rtm_task_t` \*task, const char \*name, `rtm_task_proc_t` task\_proc, void \*arg, int priority, `nanosecs_rel_t` period)  
*Initialise and start a real-time task.*
- void `rtm_task_busy_sleep` (`nanosecs_rel_t` delay)  
*Busy-wait a specified amount of time.*
- void `rtm_toseq_init` (`rtm_toseq_t` \*timeout\_seq, `nanosecs_rel_t` timeout)  
*Initialise a timeout sequence.*
- void `rtm_event_init` (`rtm_event_t` \*event, unsigned long pending)  
*Initialise an event.*
- int `rtm_event_select_bind` (`rtm_event_t` \*event, `rtm_selector_t` \*selector, enum `rtm_selecttype` type, unsigned fd\_index)  
*Bind a selector to an event.*
- int `rtm_event_wait` (`rtm_event_t` \*event)  
*Wait on event occurrence.*
- int `rtm_event_timedwait` (`rtm_event_t` \*event, `nanosecs_rel_t` timeout, `rtm_toseq_t` \*timeout\_seq)  
*Wait on event occurrence with timeout.*
- void `rtm_event_signal` (`rtm_event_t` \*event)



- *Signal an event occurrence.*
- void `rtm_event_clear` (`rtm_event_t *event`)
  - *Clear event state.*
- void `rtm_sem_init` (`rtm_sem_t *sem`, unsigned long value)
  - *Initialise a semaphore.*
- int `rtm_sem_select_bind` (`rtm_sem_t *sem`, `rtm_selector_t *selector`, enum `rtm_selecttype` type, unsigned `fd_index`)
  - *Bind a selector to a semaphore.*
- int `rtm_sem_down` (`rtm_sem_t *sem`)
  - *Decrement a semaphore.*
- int `rtm_sem_timeddown` (`rtm_sem_t *sem`, `nanosecs_rel_t` timeout, `rtm_toseq_t *timeout_seq`)
  - *Decrement a semaphore with timeout.*
- void `rtm_sem_up` (`rtm_sem_t *sem`)
  - *Increment a semaphore.*
- void `rtm_mutex_init` (`rtm_mutex_t *mutex`)
  - *Initialise a mutex.*
- int `rtm_mutex_lock` (`rtm_mutex_t *mutex`)
  - *Request a mutex.*
- int `rtm_mutex_timedlock` (`rtm_mutex_t *mutex`, `nanosecs_rel_t` timeout, `rtm_toseq_t *timeout_seq`)
  - *Request a mutex with timeout.*
- int `rtm_ratelimit` (`struct rtm_ratelimit_state *rs`, `const char *func`)
  - *Enforces a rate limit.*

## Spinlock with Preemption Deactivation

- `#define RTDM_LOCK_UNLOCKED` `RTHAL_SPIN_LOCK_UNLOCKED`
  - *Static lock initialisation.*
- `#define rtm_lock_init(lock)` `rthal_spin_lock_init(lock)`
  - *Dynamic lock initialisation.*
- `#define rtm_lock_get(lock)` `rthal_spin_lock(lock)`
  - *Acquire lock from non-preemptible contexts.*
- `#define rtm_lock_put(lock)`
  - *Release lock without preemption restoration.*
- `#define rtm_lock_get_irqsave(lock, context)`
  - *Acquire lock and disable preemption.*
- `#define rtm_lock_put_irqrestore(lock, context)`
  - *Release lock and restore preemption state.*
- `#define rtm_lock_irqsave(context)` `rthal_local_irq_save(context)`
  - *Disable preemption locally.*
- `#define rtm_lock_irqrestore(context)` `rthal_local_irq_restore(context)`
  - *Restore preemption state.*
- `typedef rthal_spinlock_t rtm_lock_t`
  - *Lock variable.*
- `typedef unsigned long rtm_lockctx_t`
  - *Variable to save the context while holding a lock.*

## 7.3.1 Detailed Description

Real-Time Driver Model for Xenomai, driver API header.

Note

Copyright (C) 2005-2007 Jan Kiszka [jan.kiszka@web.de](mailto:jan.kiszka@web.de)  
Copyright (C) 2005 Joerg Langenberg [joerg.langenberg@gmx.net](mailto:joerg.langenberg@gmx.net)  
Copyright (C) 2008 Gilles Chanteperdrix [gilles.chanteperdrix@xenomai.org](mailto:gilles.chanteperdrix@xenomai.org)

Xenomai is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

Xenomai is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with Xenomai; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 7.4 include/rtdm/rtipc.h File Reference

This file is part of the Xenomai project.

```
#include <linux/net.h>
#include <linux/socket.h>
#include <linux/if.h>
#include <nucleus/types.h>
#include <rtdm/rtdm.h>
```

### Data Structures

- struct [rtipc\\_port\\_label](#)  
*Port label information structure.*
- struct [sockaddr\\_ipc](#)  
*Socket address structure for the RTIPC address family.*

### Macros

#### XDDP socket options

*Setting and getting XDDP socket options.*

- #define [XDDP\\_LABEL](#) 1  
*XDDP label assignment.*
- #define [XDDP\\_POOLSZ](#) 2  
*XDDP local pool size configuration.*
- #define [XDDP\\_BUFSZ](#) 3  
*XDDP streaming buffer size configuration.*
- #define [XDDP\\_MONITOR](#) 4  
*XDDP monitoring callback.*

#### XDDP events

*Specific events occurring on XDDP channels, which can be monitored via the [XDDP\\_MONITOR](#) socket option.*

- #define `XDDP_EVTIN` 1  
*Monitor writes to the non real-time endpoint.*
- #define `XDDP_EVTOUT` 2  
*Monitor reads from the non real-time endpoint.*
- #define `XDDP_EVTDOWN` 3  
*Monitor close from the non real-time endpoint.*
- #define `XDDP_EVTNOBUF` 4  
*Monitor memory shortage for non real-time datagrams.*

### IDDP socket options

Setting and getting IDDP socket options.

- #define `IDDP_LABEL` 1  
*IDDP label assignment.*
- #define `IDDP_POOLSZ` 2  
*IDDP local pool size configuration.*

### BUFP socket options

Setting and getting BUFP socket options.

- #define `BUFP_LABEL` 1  
*BUFP label assignment.*
- #define `BUFP_BUFSZ` 2  
*BUFP buffer size configuration.*

### Socket level options

Setting and getting supported standard socket level options.

- #define `SO_SNDTIMEO` `defined_by_kernel_header_file`  
*IPPROTO\_IDDP and IPPROTO\_BUFP protocols support the standard SO\_SNDTIMEO socket option, from the SOL\_SOCKET level.*
- #define `SO_RCVTIMEO` `defined_by_kernel_header_file`  
*All RTIPC protocols support the standard SO\_RCVTIMEO socket option, from the SOL\_SOCKET level.*

## Typedefs

- typedef int16\_t `rtipc_port_t`  
*Port number type for the RTIPC address family.*

## Enumerations

### RTIPC protocol list

protocols for the PF\_RTIPC protocol family

- enum { `IPPROTO_IPC` = 0, `IPPROTO_XDDP` = 1, `IPPROTO_IDDP` = 2, `IPPROTO_BUFP` = 3 }

## Functions

### Supported operations

Standard socket operations supported by the RTIPC protocols.

- int `socket_AF_RTIPC` (int domain=AF\_RTIPC, int type=SOCK\_DGRAM, int protocol)  
*Create an endpoint for communication in the AF\_RTIPC domain.*

- int [close\\_\\_AF\\_RTIPC](#) (int sockfd)  
*Close a RTIPC socket descriptor.*
- int [bind\\_\\_AF\\_RTIPC](#) (int sockfd, const struct [sockaddr\\_ipc](#) \*addr, socklen\_t addrlen)  
*Bind a RTIPC socket to a port.*
- int [connect\\_\\_AF\\_RTIPC](#) (int sockfd, const struct [sockaddr\\_ipc](#) \*addr, socklen\_t addrlen)  
*Initiate a connection on a RTIPC socket.*
- int [setsockopt\\_\\_AF\\_RTIPC](#) (int sockfd, int level, int optname, const void \*optval, socklen\_t optlen)  
*Set options on RTIPC sockets.*
- int [getsockopt\\_\\_AF\\_RTIPC](#) (int sockfd, int level, int optname, void \*optval, socklen\_t \*optlen)  
*Get options on RTIPC sockets.*
- ssize\_t [sendmsg\\_\\_AF\\_RTIPC](#) (int sockfd, const struct user\_msghdr \*msg, int flags)  
*Send a message on a RTIPC socket.*
- ssize\_t [recvmsg\\_\\_AF\\_RTIPC](#) (int sockfd, struct user\_msghdr \*msg, int flags)  
*Receive a message from a RTIPC socket.*
- int [getsockname\\_\\_AF\\_RTIPC](#) (int sockfd, struct [sockaddr\\_ipc](#) \*addr, socklen\_t \*addrlen)  
*Get socket name.*
- int [getpeername\\_\\_AF\\_RTIPC](#) (int sockfd, struct [sockaddr\\_ipc](#) \*addr, socklen\_t \*addrlen)  
*Get socket peer.*

## 7.4.1 Detailed Description

This file is part of the Xenomai project.

Note

Copyright (C) 2009 Philippe Gerum [rpm@xenomai.org](mailto:rpm@xenomai.org)

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 7.5 include/rtdm/rtserial.h File Reference

Real-Time Driver Model for Xenomai, serial device profile header.

```
#include <rtdm/rtserial.h>
```

### Data Structures

- struct [rtser\\_config](#)  
*Serial device configuration.*
- struct [rtser\\_status](#)  
*Serial device status.*
- struct [rtser\\_event](#)  
*Additional information about serial device events.*

## Macros

- #define **RTSER\_RTIOC\_BREAK\_CTL** \_IOR(RTIOC\_TYPE\_SERIAL, 0x06, int)  
*Set or clear break on UART output line.*

### **RTSER\_DEF\_BAUD**

*Default baud rate*

- #define **RTSER\_DEF\_BAUD** 9600

### **RTSER\_xxx\_PARITY**

*Number of parity bits*

- #define **RTSER\_NO\_PARITY** 0x00
- #define **RTSER\_ODD\_PARITY** 0x01
- #define **RTSER\_EVEN\_PARITY** 0x03
- #define **RTSER\_DEF\_PARITY** RTSER\_NO\_PARITY

### **RTSER\_xxx\_BITS**

*Number of data bits*

- #define **RTSER\_5\_BITS** 0x00
- #define **RTSER\_6\_BITS** 0x01
- #define **RTSER\_7\_BITS** 0x02
- #define **RTSER\_8\_BITS** 0x03
- #define **RTSER\_DEF\_BITS** RTSER\_8\_BITS

### **RTSER\_xxx\_STOPB**

*Number of stop bits*

- #define **RTSER\_1\_STOPB** 0x00  
*valid only in combination with 5 data bits*
- #define **RTSER\_1\_5\_STOPB** 0x01  
*valid only in combination with 5 data bits*
- #define **RTSER\_2\_STOPB** 0x01  
*valid only in combination with 5 data bits*
- #define **RTSER\_DEF\_STOPB** RTSER\_1\_STOPB  
*valid only in combination with 5 data bits*

### **RTSER\_xxx\_HAND**

*Handshake mechanisms*

- #define **RTSER\_NO\_HAND** 0x00
- #define **RTSER\_RTSCTS\_HAND** 0x01
- #define **RTSER\_DEF\_HAND** RTSER\_NO\_HAND

### **RTSER\_RS485\_xxx**

*RS485 mode with automatic RTS handling*

- #define **RTSER\_RS485\_DISABLE** 0x00
- #define **RTSER\_RS485\_ENABLE** 0x01
- #define **RTSER\_DEF\_RS485** RTSER\_RS485\_DISABLE

### **RTSER\_FIFO\_xxx**

*Reception FIFO interrupt threshold*

- #define **RTSER\_FIFO\_DEPTH\_1** 0x00

- #define RTSER\_FIFO\_DEPTH\_4 0x40
- #define RTSER\_FIFO\_DEPTH\_8 0x80
- #define RTSER\_FIFO\_DEPTH\_14 0xC0
- #define RTSER\_DEF\_FIFO\_DEPTH RTSER\_FIFO\_DEPTH\_1

## RTSER\_TIMEOUT\_xxx

*Special timeout values, see also [RTDM\\_TIMEOUT\\_xxx](#)*

- #define RTSER\_TIMEOUT\_INFINITE [RTDM\\_TIMEOUT\\_INFINITE](#)
- #define RTSER\_TIMEOUT\_NONE [RTDM\\_TIMEOUT\\_NONE](#)
- #define RTSER\_DEF\_TIMEOUT [RTDM\\_TIMEOUT\\_INFINITE](#)

## RTSER\_xxx\_TIMESTAMP\_HISTORY

*Timestamp history control*

- #define RTSER\_RX\_TIMESTAMP\_HISTORY 0x01
- #define RTSER\_DEF\_TIMESTAMP\_HISTORY 0x00

## RTSER\_EVENT\_xxx

*Events bits*

- #define RTSER\_EVENT\_RXPEND 0x01
- #define RTSER\_EVENT\_ERRPEND 0x02
- #define RTSER\_EVENT\_MODEMHI 0x04
- #define RTSER\_EVENT\_MODEMLO 0x08
- #define RTSER\_EVENT\_TXEMPTY 0x10
- #define RTSER\_DEF\_EVENT\_MASK 0x00

## RTSER\_SET\_xxx

*Configuration mask bits*

- #define RTSER\_SET\_BAUD 0x0001
- #define RTSER\_SET\_PARITY 0x0002
- #define RTSER\_SET\_DATA\_BITS 0x0004
- #define RTSER\_SET\_STOP\_BITS 0x0008
- #define RTSER\_SET\_HANDSHAKE 0x0010
- #define RTSER\_SET\_FIFO\_DEPTH 0x0020
- #define RTSER\_SET\_TIMEOUT\_RX 0x0100
- #define RTSER\_SET\_TIMEOUT\_TX 0x0200
- #define RTSER\_SET\_TIMEOUT\_EVENT 0x0400
- #define RTSER\_SET\_TIMESTAMP\_HISTORY 0x0800
- #define RTSER\_SET\_EVENT\_MASK 0x1000
- #define RTSER\_SET\_RS485 0x2000

## RTSER\_LSR\_xxx

*Line status bits*

- #define RTSER\_LSR\_DATA 0x01
- #define RTSER\_LSR\_OVERRUN\_ERR 0x02
- #define RTSER\_LSR\_PARITY\_ERR 0x04
- #define RTSER\_LSR\_FRAMING\_ERR 0x08
- #define RTSER\_LSR\_BREAK\_IND 0x10
- #define RTSER\_LSR\_THR\_EMPTY 0x20
- #define RTSER\_LSR\_TRANSM\_EMPTY 0x40
- #define RTSER\_LSR\_FIFO\_ERR 0x80
- #define RTSER\_SOFT\_OVERRUN\_ERR 0x0100

## RTSER\_MSR\_xxx

*Modem status bits*

- #define **RTSER\_MSR\_DCTS** 0x01
- #define **RTSER\_MSR\_DDSR** 0x02
- #define **RTSER\_MSR\_TERI** 0x04
- #define **RTSER\_MSR\_DDCD** 0x08
- #define **RTSER\_MSR\_CTS** 0x10
- #define **RTSER\_MSR\_DSR** 0x20
- #define **RTSER\_MSR\_RI** 0x40
- #define **RTSER\_MSR\_DCD** 0x80

## RTSER\_MCR\_xxx

*Modem control bits*

- #define **RTSER\_MCR\_DTR** 0x01
- #define **RTSER\_MCR\_RTS** 0x02
- #define **RTSER\_MCR\_OUT1** 0x04
- #define **RTSER\_MCR\_OUT2** 0x08
- #define **RTSER\_MCR\_LOOP** 0x10

## Sub-Classes of RTDM\_CLASS\_SERIAL

- #define **RTDM\_SUBCLASS\_16550A** 0

## IOCTLs

*Serial device IOCTLs*

- #define **RTSER\_RTIOC\_GET\_CONFIG** \_IOR(RTIOC\_TYPE\_SERIAL, 0x00, struct rtser\_config)  
*Get serial device configuration.*
- #define **RTSER\_RTIOC\_SET\_CONFIG** \_IOW(RTIOC\_TYPE\_SERIAL, 0x01, struct rtser\_config)  
*Set serial device configuration.*
- #define **RTSER\_RTIOC\_GET\_STATUS** \_IOR(RTIOC\_TYPE\_SERIAL, 0x02, struct rtser\_status)  
*Get serial device status.*
- #define **RTSER\_RTIOC\_GET\_CONTROL** \_IOR(RTIOC\_TYPE\_SERIAL, 0x03, int)  
*Get serial device's modem control register.*
- #define **RTSER\_RTIOC\_SET\_CONTROL** \_IOW(RTIOC\_TYPE\_SERIAL, 0x04, int)  
*Set serial device's modem control register.*
- #define **RTSER\_RTIOC\_WAIT\_EVENT** \_IOR(RTIOC\_TYPE\_SERIAL, 0x05, struct rtser\_event)  
*Wait on serial device events according to previously set mask.*

## RTSER\_BREAK\_xxx

*Break control*

- #define **RTSER\_BREAK\_CLR** 0x00  
*Serial device configuration.*
- #define **RTSER\_BREAK\_SET** 0x01  
*Serial device configuration.*
- #define **RTIOC\_TYPE\_SERIAL** RTDM\_CLASS\_SERIAL  
*Serial device configuration.*
- typedef struct **rtser\_config** **rtser\_config\_t**  
*Serial device configuration.*
- typedef struct **rtser\_status** **rtser\_status\_t**  
*Serial device status.*
- typedef struct **rtser\_event** **rtser\_event\_t**  
*Additional information about serial device events.*

## 7.5.1 Detailed Description

Real-Time Driver Model for Xenomai, serial device profile header.

Note

Copyright (C) 2005-2007 Jan Kiszka [jan.kiszka@web.de](mailto:jan.kiszka@web.de)

Xenomai is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

Xenomai is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with Xenomai; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 7.6 include/rtdm/rtesting.h File Reference

Real-Time Driver Model for Xenomai, testing device profile header.

```
#include <rtdm/rtdm.h>
```

Macros

### Sub-Classes of RTDM\_CLASS\_TESTING

- #define **RTDM\_SUBCLASS\_TIMERBENCH** 0  
*subclass name: "timerbench"*
- #define **RTDM\_SUBCLASS\_IRQBENCH** 1  
*subclass name: "irqbench"*
- #define **RTDM\_SUBCLASS\_SWITCHTEST** 2  
*subclass name: "switchtest"*
- #define **RTDM\_SUBCLASS\_RTDMTTEST** 3  
*subclass name: "rtdm"*

### IOCTLs

*Testing device IOCTLs*

- #define **RTTST\_RTIOC\_INTERM\_BENCH\_RES** \_IOWR(RTIOC\_TYPE\_TESTING, 0x00, struct rttst\_interm\_bench\_res)
- #define **RTTST\_RTIOC\_TMBENCH\_START** \_IOW(RTIOC\_TYPE\_TESTING, 0x10, struct rttst\_tmbench\_config)
- #define **RTTST\_RTIOC\_TMBENCH\_STOP** \_IOWR(RTIOC\_TYPE\_TESTING, 0x11, struct rttst\_overall\_bench\_res)
- #define **RTTST\_RTIOC\_IRQBENCH\_START** \_IOW(RTIOC\_TYPE\_TESTING, 0x20, struct rttst\_irqbench\_config)
- #define **RTTST\_RTIOC\_IRQBENCH\_STOP** \_IO(RTIOC\_TYPE\_TESTING, 0x21)
- #define **RTTST\_RTIOC\_IRQBENCH\_GET\_STATS** \_IOR(RTIOC\_TYPE\_TESTING, 0x22, struct rttst\_irqbench\_stats)
- #define **RTTST\_RTIOC\_IRQBENCH\_WAIT\_IRQ** \_IO(RTIOC\_TYPE\_TESTING, 0x23)
- #define **RTTST\_RTIOC\_IRQBENCH\_REPLY\_IRQ** \_IO(RTIOC\_TYPE\_TESTING, 0x24)
- #define **RTTST\_RTIOC\_SWTEST\_SET\_TASKS\_COUNT** \_IOW(RTIOC\_TYPE\_TESTING, 0x30, unsigned long)



- #define `RTTST_RTIOC_SWTEST_SET_CPU` `_IOW(RTTST_RTIOC_TYPE_TESTING, 0x31, unsigned long)`
- #define `RTTST_RTIOC_SWTEST_REGISTER_UTASK` `_IOW(RTTST_RTIOC_TYPE_TESTING, 0x32, struct rttst_swtest_task)`
- #define `RTTST_RTIOC_SWTEST_CREATE_KTASK` `_IOWR(RTTST_RTIOC_TYPE_TESTING, 0x33, struct rttst_swtest_task)`
- #define `RTTST_RTIOC_SWTEST_PEND` `_IOR(RTTST_RTIOC_TYPE_TESTING, 0x34, struct rttst_swtest_task)`
- #define `RTTST_RTIOC_SWTEST_SWITCH_TO` `_IOR(RTTST_RTIOC_TYPE_TESTING, 0x35, struct rttst_swtest_dir)`
- #define `RTTST_RTIOC_SWTEST_GET_SWITCHES_COUNT` `_IOR(RTTST_RTIOC_TYPE_TESTING, 0x36, unsigned long)`
- #define `RTTST_RTIOC_SWTEST_GET_LAST_ERROR` `_IOR(RTTST_RTIOC_TYPE_TESTING, 0x37, struct rttst_swtest_error)`
- #define `RTTST_RTIOC_SWTEST_SET_PAUSE` `_IOW(RTTST_RTIOC_TYPE_TESTING, 0x38, unsigned long)`
- #define `RTTST_RTIOC_RTDM_DEFER_CLOSE` `_IOW(RTTST_RTIOC_TYPE_TESTING, 0x40, unsigned long)`

## 7.6.1 Detailed Description

Real-Time Driver Model for Xenomai, testing device profile header.

Note

Copyright (C) 2005 Jan Kiszka [jan.kiszka@web.de](mailto:jan.kiszka@web.de)

Xenomai is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

Xenomai is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with Xenomai; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 7.7 ksrc/skins/rtdm/device.c File Reference

Real-Time Driver Model for Xenomai, device management.

```
#include <linux/module.h>
#include <linux/delay.h>
#include "rtdm/internal.h"
```

### Functions

- int `rtdm_dev_register` (struct `rtdm_device` \*device)  
*Register a RTDM device.*
- int `rtdm_dev_unregister` (struct `rtdm_device` \*device, unsigned int poll\_delay)  
*Unregisters a RTDM device.*

## 7.7.1 Detailed Description

Real-Time Driver Model for Xenomai, device management.

Note

Copyright (C) 2005 Jan Kiszka [jan.kiszka@web.de](mailto:jan.kiszka@web.de)  
Copyright (C) 2005 Joerg Langenberg [joerg.langenberg@gmx.net](mailto:joerg.langenberg@gmx.net)

Xenomai is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

Xenomai is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with Xenomai; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 7.8 ksrc/skins/rtdm/drvlib.c File Reference

Real-Time Driver Model for Xenomai, driver library.

```
#include <linux/bitops.h>
#include <linux/delay.h>
#include <linux/mman.h>
#include <asm/page.h>
#include <asm/io.h>
#include <asm/pgtable.h>
#include <linux/highmem.h>
#include <linux/err.h>
#include <rtdm/rtdm_driver.h>
```

### Functions

- [nanosecs\\_abs\\_t rtdm\\_clock\\_read](#) (void)  
*Get system time.*
- [nanosecs\\_abs\\_t rtdm\\_clock\\_read\\_monotonic](#) (void)  
*Get monotonic time.*
- int [rtdm\\_task\\_init](#) (rtdm\_task\_t \*task, const char \*name, [rtdm\\_task\\_proc\\_t](#) task\_proc, void \*arg, int priority, [nanosecs\\_rel\\_t](#) period)  
*Intialise and start a real-time task.*
- void [rtdm\\_task\\_destroy](#) (rtdm\_task\_t \*task)  
*Destroy a real-time task.*
- void [rtdm\\_task\\_set\\_priority](#) (rtdm\_task\_t \*task, int priority)  
*Adjust real-time task priority.*
- int [rtdm\\_task\\_set\\_period](#) (rtdm\_task\_t \*task, [nanosecs\\_rel\\_t](#) period)  
*Adjust real-time task period.*
- int [rtdm\\_task\\_wait\\_period](#) (void)  
*Wait on next real-time task period.*
- int [rtdm\\_task\\_unblock](#) (rtdm\_task\_t \*task)  
*Activate a blocked real-time task.*

- `rtm_task_t * rtdm_task_current` (void)  
*Get current real-time task.*
- `int rtdm_task_sleep` (`nanosecs_rel_t` delay)  
*Sleep a specified amount of time.*
- `int rtdm_task_sleep_until` (`nanosecs_abs_t` wakeup\_time)  
*Sleep until a specified absolute time.*
- `int rtdm_task_sleep_abs` (`nanosecs_abs_t` wakeup\_time, enum `rtdm_timer_mode` mode)  
*Sleep until a specified absolute time.*
- `void rtdm_task_join_nrt` (`rtm_task_t *task`, unsigned int poll\_delay)  
*Wait on a real-time task to terminate.*
- `void rtdm_task_busy_sleep` (`nanosecs_rel_t` delay)  
*Busy-wait a specified amount of time.*
- `int rtdm_timer_init` (`rtm_timer_t *timer`, `rtdm_timer_handler_t` handler, const char \*name)  
*Initialise a timer.*
- `void rtdm_timer_destroy` (`rtm_timer_t *timer`)  
*Destroy a timer.*
- `int rtdm_timer_start` (`rtm_timer_t *timer`, `nanosecs_abs_t` expiry, `nanosecs_rel_t` interval, enum `rtdm_timer_mode` mode)  
*Start a timer.*
- `void rtdm_timer_stop` (`rtm_timer_t *timer`)  
*Stop a timer.*
- `int rtdm_timer_start_in_handler` (`rtm_timer_t *timer`, `nanosecs_abs_t` expiry, `nanosecs_rel_t` interval, enum `rtdm_timer_mode` mode)  
*Start a timer from inside a timer handler.*
- `void rtdm_timer_stop_in_handler` (`rtm_timer_t *timer`)  
*Stop a timer from inside a timer handler.*
- `int rtdm_irq_request` (`rtm_irq_t *irq_handle`, unsigned int irq\_no, `rtdm_irq_handler_t` handler, unsigned long flags, const char \*device\_name, void \*arg)  
*Register an interrupt handler.*
- `int rtdm_irq_free` (`rtm_irq_t *irq_handle`)  
*Release an interrupt handler.*
- `int rtdm_irq_enable` (`rtm_irq_t *irq_handle`)  
*Enable interrupt line.*
- `int rtdm_irq_disable` (`rtm_irq_t *irq_handle`)  
*Disable interrupt line.*
- `int rtdm_nrtsig_init` (`rtm_nrtsig_t *nrt_sig`, `rtdm_nrtsig_handler_t` handler, void \*arg)  
*Register a non-real-time signal handler.*
- `void rtdm_nrtsig_destroy` (`rtm_nrtsig_t *nrt_sig`)  
*Release a non-realtime signal handler.*
- `void rtdm_nrtsig_pend` (`rtm_nrtsig_t *nrt_sig`)  
*Trigger non-real-time signal.*
- `int rtdm_mmap_to_user` (`rtm_user_info_t *user_info`, void \*src\_addr, size\_t len, int prot, void \*\*pptr, struct `vm_operations_struct *vm_ops`, void \*vm\_private\_data)  
*Map a kernel memory range into the address space of the user.*
- `int rtdm_iomap_to_user` (`rtm_user_info_t *user_info`, `phys_addr_t` src\_addr, size\_t len, int prot, void \*\*pptr, struct `vm_operations_struct *vm_ops`, void \*vm\_private\_data)  
*Map an I/O memory range into the address space of the user.*
- `int rtdm_munmap` (`rtm_user_info_t *user_info`, void \*ptr, size\_t len)  
*Unmap a user memory range.*
- `int rtdm_ratelimit` (struct `rtdm_ratelimit_state *rs`, const char \*func)  
*Enforces a rate limit.*

- void `rtm_printk_ratelimited` (const char \*format,...)  
*Real-time safe rate-limited message printing on kernel console.*
- void `rtm_printk` (const char \*format,...)  
*Real-time safe message printing on kernel console.*
- void \* `rtm_malloc` (size\_t size)  
*Allocate memory block in real-time context.*
- void `rtm_free` (void \*ptr)  
*Release real-time memory block.*
- int `rtm_read_user_ok` (rtm\_user\_info\_t \*user\_info, const void \_\_user \*ptr, size\_t size)  
*Check if read access to user-space memory block is safe.*
- int `rtm_rw_user_ok` (rtm\_user\_info\_t \*user\_info, const void \_\_user \*ptr, size\_t size)  
*Check if read/write access to user-space memory block is safe.*
- int `rtm_copy_from_user` (rtm\_user\_info\_t \*user\_info, void \*dst, const void \_\_user \*src, size\_t size)  
*Copy user-space memory block to specified buffer.*
- int `rtm_safe_copy_from_user` (rtm\_user\_info\_t \*user\_info, void \*dst, const void \_\_user \*src, size\_t size)  
*Check if read access to user-space memory block and copy it to specified buffer.*
- int `rtm_copy_to_user` (rtm\_user\_info\_t \*user\_info, void \_\_user \*dst, const void \*src, size\_t size)  
*Copy specified buffer to user-space memory block.*
- int `rtm_safe_copy_to_user` (rtm\_user\_info\_t \*user\_info, void \_\_user \*dst, const void \*src, size\_t size)  
*Check if read/write access to user-space memory block is safe and copy specified buffer to it.*
- int `rtm_strncpy_from_user` (rtm\_user\_info\_t \*user\_info, char \*dst, const char \_\_user \*src, size\_t count)  
*Copy user-space string to specified buffer.*
- int `rtm_in_rt_context` (void)  
*Test if running in a real-time task.*
- int `rtm_rt_capable` (rtm\_user\_info\_t \*user\_info)  
*Test if the caller is capable of running in real-time context.*

## Timeout Sequence Management

- void `rtm_toseq_init` (rtm\_toseq\_t \*timeout\_seq, `nanosecs_rel_t` timeout)  
*Initialise a timeout sequence.*
- `EXPORT_SYMBOL_GPL` (`rtm_toseq_init`)  
*Initialise a timeout sequence.*

## Event Services

- void `rtm_event_init` (rtm\_event\_t \*event, unsigned long pending)  
*Initialise an event.*
- `EXPORT_SYMBOL_GPL` (`rtm_event_init`)  
*Initialise an event.*
- void `rtm_event_destroy` (rtm\_event\_t \*event)  
*Destroy an event.*
- void `rtm_event_pulse` (rtm\_event\_t \*event)  
*Signal an event occurrence to currently listening waiters.*
- void `rtm_event_signal` (rtm\_event\_t \*event)  
*Signal an event occurrence.*
- `EXPORT_SYMBOL_GPL` (`rtm_event_signal`)  
*Initialise an event.*
- int `rtm_event_wait` (rtm\_event\_t \*event)  
*Wait on event occurrence.*

- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_event\\_wait\)](#)  
*Initialise an event.*
- `int rtdm_event_timedwait (rtdm_event_t *event, nanosecs_rel_t timeout, rtdm_toseq_t *timeout_seq)`  
*Wait on event occurrence with timeout.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_event\\_timedwait\)](#)  
*Initialise an event.*
- `void rtdm_event_clear (rtdm_event_t *event)`  
*Clear event state.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_event\\_clear\)](#)  
*Initialise an event.*
- `int rtdm_event_select_bind (rtdm_event_t *event, rtdm_selector_t *selector, enum rtdm_selecttype type, unsigned fd_index)`  
*Bind a selector to an event.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_event\\_select\\_bind\)](#)  
*Initialise an event.*

## Semaphore Services

- `void rtdm_sem_init (rtdm_sem_t *sem, unsigned long value)`  
*Initialise a semaphore.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_sem\\_init\)](#)  
*Initialise a semaphore.*
- `void rtdm_sem_destroy (rtdm_sem_t *sem)`  
*Destroy a semaphore.*
- `int rtdm_sem_down (rtdm_sem_t *sem)`  
*Decrement a semaphore.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_sem\\_down\)](#)  
*Initialise a semaphore.*
- `int rtdm_sem_timeddown (rtdm_sem_t *sem, nanosecs_rel_t timeout, rtdm_toseq_t *timeout_seq)`  
*Decrement a semaphore with timeout.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_sem\\_timeddown\)](#)  
*Initialise a semaphore.*
- `void rtdm_sem_up (rtdm_sem_t *sem)`  
*Increment a semaphore.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_sem\\_up\)](#)  
*Initialise a semaphore.*
- `int rtdm_sem_select_bind (rtdm_sem_t *sem, rtdm_selector_t *selector, enum rtdm_selecttype type, unsigned fd_index)`  
*Bind a selector to a semaphore.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_sem\\_select\\_bind\)](#)  
*Initialise a semaphore.*

## Mutex Services

- `void rtdm_mutex_init (rtdm_mutex_t *mutex)`  
*Initialise a mutex.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_mutex\\_init\)](#)  
*Initialise a mutex.*
- `void rtdm_mutex_destroy (rtdm_mutex_t *mutex)`  
*Destroy a mutex.*
- `void rtdm_mutex_unlock (rtdm_mutex_t *mutex)`  
*Release a mutex.*
- `int rtdm_mutex_lock (rtdm_mutex_t *mutex)`  
*Request a mutex.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_mutex\\_lock\)](#)  
*Initialise a mutex.*
- `int rtdm_mutex_timedlock (rtdm_mutex_t *mutex, nanosecs_rel_t timeout, rtdm_toseq_t *timeout_seq)`  
*Request a mutex with timeout.*
- [EXPORT\\_SYMBOL\\_GPL \(rtdm\\_mutex\\_timedlock\)](#)  
*Initialise a mutex.*

## 7.8.1 Detailed Description

Real-Time Driver Model for Xenomai, driver library.

Note

Copyright (C) 2005-2007 Jan Kiszka [jan.kiszka@web.de](mailto:jan.kiszka@web.de)  
Copyright (C) 2005 Joerg Langenberg [joerg.langenberg@gmx.net](mailto:joerg.langenberg@gmx.net)  
Copyright (C) 2008 Gilles Chanteperdrix [gilles.chanteperdrix@xenomai.org](mailto:gilles.chanteperdrix@xenomai.org)

Xenomai is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

Xenomai is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with Xenomai; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 7.9 ksrc/skins/rtdm/module.c File Reference

Real-Time Driver Model for Xenomai.

```
#include <nucleus/pod.h>
#include <rtdm/syscall.h>
#include "rtdm/internal.h"
```

### 7.9.1 Detailed Description

Real-Time Driver Model for Xenomai.

Note

Copyright (C) 2005, 2006 Jan Kiszka [jan.kiszka@web.de](mailto:jan.kiszka@web.de)  
Copyright (C) 2005 Joerg Langenberg [joerg.langenberg@gmx.net](mailto:joerg.langenberg@gmx.net)

Xenomai is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

Xenomai is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with Xenomai; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

## 7.10 ksrc/skins/rtdm/core.c File Reference

Real-Time Driver Model for Xenomai, device operation multiplexing.

```

#include <nucleus/pod.h>
#include <nucleus/ppd.h>
#include <nucleus/heap.h>
#include <rtm/syscall.h>
#include "rtm/internal.h"
#include <linux/workqueue.h>

```

## Functions

- struct `rtm_dev_context` \* `rtm_context_get` (int fd)  
*Retrieve and lock a device context.*
- int `rtm_select_bind` (int fd, `rtm_selector_t` \*selector, enum `rtm_selecttype` type, unsigned fd\_index)  
*Bind a selector to specified event types of a given file descriptor.*
- void `rtm_context_lock` (struct `rtm_dev_context` \*context)  
*Increment context reference counter.*
- void `rtm_context_unlock` (struct `rtm_dev_context` \*context)  
*Decrement context reference counter.*
- void `rtm_context_put` (struct `rtm_dev_context` \*context)  
*Release a device context obtained via `rtm_context_get()`*
- int `rtm_open` (const char \*path, int oflag,...)  
*Open a device.*
- int `rtm_socket` (int protocol\_family, int socket\_type, int protocol)  
*Create a socket.*
- int `rtm_close` (int fd)  
*Close a device or socket.*
- int `rtm_ioctl` (int fd, int request,...)  
*Issue an IOCTL.*
- `ssize_t` `rtm_read` (int fd, void \*buf, `size_t` nbyte)  
*Read from device.*
- `ssize_t` `rtm_write` (int fd, const void \*buf, `size_t` nbyte)  
*Write to device.*
- `ssize_t` `rtm_recvmsg` (int fd, struct `user_msghdr` \*msg, int flags)  
*Receive message from socket.*
- `ssize_t` `rtm_recvfrom` (int fd, void \*buf, `size_t` len, int flags, struct `sockaddr` \*from, `socklen_t` \*fromlen)  
*Receive message from socket.*
- `ssize_t` `rtm_recv` (int fd, void \*buf, `size_t` len, int flags)  
*Receive message from socket.*
- `ssize_t` `rtm_sendmsg` (int fd, const struct `user_msghdr` \*msg, int flags)  
*Transmit message to socket.*
- `ssize_t` `rtm_sendto` (int fd, const void \*buf, `size_t` len, int flags, const struct `sockaddr` \*to, `socklen_t` tolen)  
*Transmit message to socket.*
- `ssize_t` `rtm_send` (int fd, const void \*buf, `size_t` len, int flags)  
*Transmit message to socket.*
- int `rtm_bind` (int fd, const struct `sockaddr` \*my\_addr, `socklen_t` addrlen)  
*Bind to local address.*
- int `rtm_connect` (int fd, const struct `sockaddr` \*serv\_addr, `socklen_t` addrlen)  
*Connect to remote address.*

- int `rtm_listen` (int fd, int backlog)  
*Listen for incoming connection requests.*
- int `rtm_accept` (int fd, struct sockaddr \*addr, socklen\_t \*addrlen)  
*Accept a connection requests.*
- int `rtm_shutdown` (int fd, int how)  
*Shut down parts of a connection.*
- int `rtm_getsockopt` (int fd, int level, int optname, void \*optval, socklen\_t \*optlen)  
*Get socket option.*
- int `rtm_setsockopt` (int fd, int level, int optname, const void \*optval, socklen\_t optlen)  
*Set socket option.*
- int `rtm_getsockname` (int fd, struct sockaddr \*name, socklen\_t \*namelen)  
*Get local socket address.*
- int `rtm_getpeername` (int fd, struct sockaddr \*name, socklen\_t \*namelen)  
*Get socket destination address.*
- int `rt_dev_open` (const char \*path, int oflag,...)  
*Open a device.*
- int `rt_dev_socket` (int protocol\_family, int socket\_type, int protocol)  
*Create a socket.*
- int `rt_dev_close` (int fd)  
*Close a device or socket.*
- int `rt_dev_ioctl` (int fd, int request,...)  
*Issue an IOCTL.*
- ssize\_t `rt_dev_read` (int fd, void \*buf, size\_t nbyte)  
*Read from device.*
- ssize\_t `rt_dev_write` (int fd, const void \*buf, size\_t nbyte)  
*Write to device.*
- ssize\_t `rt_dev_recvmsg` (int fd, struct user\_msghdr \*msg, int flags)  
*Receive message from socket.*
- ssize\_t `rt_dev_recvfrom` (int fd, void \*buf, size\_t len, int flags, struct sockaddr \*from, socklen\_t \*fromlen)  
*Receive message from socket.*
- ssize\_t `rt_dev_recv` (int fd, void \*buf, size\_t len, int flags)  
*Receive message from socket.*
- ssize\_t `rt_dev_sendmsg` (int fd, const struct user\_msghdr \*msg, int flags)  
*Transmit message to socket.*
- ssize\_t `rt_dev_sendto` (int fd, const void \*buf, size\_t len, int flags, const struct sockaddr \*to, socklen\_t tolen)  
*Transmit message to socket.*
- ssize\_t `rt_dev_send` (int fd, const void \*buf, size\_t len, int flags)  
*Transmit message to socket.*
- int `rt_dev_bind` (int fd, const struct sockaddr \*my\_addr, socklen\_t addrlen)  
*Bind to local address.*
- int `rt_dev_connect` (int fd, const struct sockaddr \*serv\_addr, socklen\_t addrlen)  
*Connect to remote address.*
- int `rt_dev_listen` (int fd, int backlog)  
*Listen for incoming connection requests.*
- int `rt_dev_accept` (int fd, struct sockaddr \*addr, socklen\_t \*addrlen)  
*Accept a connection requests.*
- int `rt_dev_shutdown` (int fd, int how)  
*Shut down parts of a connection.*
- int `rt_dev_getsockopt` (int fd, int level, int optname, void \*optval, socklen\_t \*optlen)



*Get socket option.*

- int [rt\\_dev\\_setsockopt](#) (int fd, int level, int optname, const void \*optval, socklen\_t optlen)

*Set socket option.*

- int [rt\\_dev\\_getsockname](#) (int fd, struct sockaddr \*name, socklen\_t \*namelen)

*Get local socket address.*

- int [rt\\_dev\\_getpeername](#) (int fd, struct sockaddr \*name, socklen\_t \*namelen)

*Get socket destination address.*

## 7.10.1 Detailed Description

Real-Time Driver Model for Xenomai, device operation multiplexing.

Note

Copyright (C) 2005 Jan Kiszka [jan.kiszka@web.de](mailto:jan.kiszka@web.de)

Copyright (C) 2005 Joerg Langenberg [joerg.langenberg@gmx.net](mailto:joerg.langenberg@gmx.net)

Xenomai is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

Xenomai is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with Xenomai; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.



# Chapter 8

## Example Documentation

### 8.1 bufp-label.c

```
/*
 * BUFp-based client/server demo, using the read(2)/write(2)
 * system calls to exchange data over a socket.
 *
 * In this example, two sockets are created. A server thread (reader)
 * is bound to a real-time port and receives a stream of bytes sent to
 * this port from a client thread (writer).
 *
 * See Makefile in this directory for build directives.
 */
#include <sys/mman.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <pthread.h>
#include <errno.h>
#include <rtdk.h>
#include <rtdm/rtipc.h>

pthread_t svtid, cltid;

#define BUFp_PORT_LABEL "bufp-demo"

static const char *msg[] = {
    "Surfing With The Alien",
    "Lords of Karma",
    "Banana Mango",
    "Psycho Monkey",
    "Luminous Flesh Giants",
    "Moroccan Sunset",
    "Satch Boogie",
    "Flying In A Blue Dream",
    "Ride",
    "Summer Song",
    "Speed Of Light",
    "Crystal Planet",
    "Raspberry Jam Delta-V",
    "Champagne?",
    "Clouds Race Across The Sky",
    "Engines Of Creation"
};

static void fail(const char *reason)
{
    perror(reason);
    exit(EXIT_FAILURE);
}

static void *server(void *arg)
{
    struct rtipc_port_label plabel;
    struct sockaddr_ipc saddr;
    char buf[128];
    size_t bufsz;
    int ret, s;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_BUFp);
```

```

if (s < 0)
    fail("socket");

/*
 * Set a 16k buffer for the server endpoint. This
 * configuration must be done prior to binding the socket to a
 * port.
 */
bufsz = 16384; /* bytes */
ret = setsockopt(s, SOL_BUFPP, BUFPP_BUFSZ,
                &bufsz, sizeof(bufsz));
if (ret)
    fail("setsockopt");

/*
 * Set a port label. This name will be registered when
 * binding, in addition to the port number (if given).
 */
strcpy(plabel.label, BUFPP_PORT_LABEL);
ret = setsockopt(s, SOL_BUFPP, BUFPP_LABEL,
                &plabel, sizeof(plabel));
if (ret)
    fail("setsockopt");

/*
 * Bind the socket to the port. Assign that port a label, so
 * that peers may use a descriptive information to locate
 * it. Labeled ports will appear in the
 * /proc/xenomai/registry/rtpc/bufp directory once the socket
 * is bound.
 *
 * saddr.sipc_port specifies the port number to use. If -1 is
 * passed, the BUFPP driver will auto-select an idle port.
 */
saddr.sipc_family = AF_RTIPC;
saddr.sipc_port = -1;
ret = bind(s, (struct sockaddr *)&saddr, sizeof(saddr));
if (ret)
    fail("bind");

for (;;) {
    ret = read(s, buf, sizeof(buf));
    if (ret < 0) {
        close(s);
        fail("read");
    }
    rt_printf("%s: received %d bytes, \"%s\"\n",
              __FUNCTION__, ret, buf);
}

return NULL;
}

static void *client(void *arg)
{
    struct rtpc_port_label plabel;
    struct sockaddr_ipc svaddr;
    int ret, s, n = 0, len;
    struct timespec ts;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_BUFPP);
    if (s < 0)
        fail("socket");

    /*
     * Set the port label. This name will be used to find the peer
     * when connecting, instead of the port number. The label must
     * be set _after_ the socket is bound to the port, so that
     * BUFPP does not try to register this label for the client
     * port as well (like the server thread did).
     */
    strcpy(plabel.label, BUFPP_PORT_LABEL);
    ret = setsockopt(s, SOL_BUFPP, BUFPP_LABEL,
                    &plabel, sizeof(plabel));
    if (ret)
        fail("setsockopt");

    memset(&svaddr, 0, sizeof(svaddr));
    svaddr.sipc_family = AF_RTIPC;
    svaddr.sipc_port = -1; /* Tell BUFPP to search by label. */
    ret = connect(s, (struct sockaddr *)&svaddr, sizeof(svaddr));
    if (ret)
        fail("connect");

    for (;;) {
        len = strlen(msg[n]);
        ret = write(s, msg[n], len);

```

```

        if (ret < 0) {
            close(s);
            fail("write");
        }
        rt_printf("%s: sent %d bytes, \"%.*s\"\n",
            __FUNCTION__, ret, ret, msg[n]);
        n = (n + 1) % (sizeof(msg) / sizeof(msg[0]));
        /*
         * We run in full real-time mode (i.e. primary mode),
         * so we have to let the system breathe between two
         * iterations.
         */
        ts.tv_sec = 0;
        ts.tv_nsec = 500000000; /* 500 ms */
        clock_nanosleep(CLOCK_REALTIME, 0, &ts, NULL);
    }

    return NULL;
}

static void cleanup_upon_sig(int sig)
{
    pthread_cancel(svtid);
    pthread_cancel(cltid);
    signal(sig, SIG_DFL);
    pthread_join(svtid, NULL);
    pthread_join(cltid, NULL);
}

int main(int argc, char **argv)
{
    struct sched_param svparam = {.sched_priority = 71 };
    struct sched_param clparam = {.sched_priority = 70 };
    pthread_attr_t svattr, clattr;
    sigset_t mask, oldmask;

    mlockall(MCL_CURRENT | MCL_FUTURE);

    sigemptyset(&mask);
    sigaddset(&mask, SIGINT);
    signal(SIGINT, cleanup_upon_sig);
    sigaddset(&mask, SIGTERM);
    signal(SIGTERM, cleanup_upon_sig);
    sigaddset(&mask, SIGHUP);
    signal(SIGHUP, cleanup_upon_sig);
    pthread_sigmask(SIG_BLOCK, &mask, &oldmask);

    /*
     * This is a real-time compatible printf() package from
     * Xenomai's RT Development Kit (RTDK), that does NOT cause
     * any transition to secondary mode.
     */
    rt_print_auto_init(1);

    pthread_attr_init(&svattr);
    pthread_attr_setdetachstate(&svattr, PTHREAD_CREATE_JOINABLE);
    pthread_attr_setinheritsched(&svattr, PTHREAD_EXPLICIT_SCHED);
    pthread_attr_setschedpolicy(&svattr, SCHED_FIFO);
    pthread_attr_setschedparam(&svattr, &svparam);

    errno = pthread_create(&svtid, &svattr, &server, NULL);
    if (errno)
        fail("pthread_create");

    pthread_attr_init(&clattr);
    pthread_attr_setdetachstate(&clattr, PTHREAD_CREATE_JOINABLE);
    pthread_attr_setinheritsched(&clattr, PTHREAD_EXPLICIT_SCHED);
    pthread_attr_setschedpolicy(&clattr, SCHED_FIFO);
    pthread_attr_setschedparam(&clattr, &clparam);

    errno = pthread_create(&cltid, &clattr, &client, NULL);
    if (errno)
        fail("pthread_create");

    sigsuspend(&oldmask);

    return 0;
}

```

## 8.2 bufp-readwrite.c

```

/*
 * BUFp-based client/server demo, using the read(2)/write(2)

```

```

* system calls to exchange data over a socket.
*
* In this example, two sockets are created. A server thread (reader)
* is bound to a real-time port and receives a stream of bytes sent to
* this port from a client thread (writer).
*
* See Makefile in this directory for build directives.
*/
#include <sys/mman.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <pthread.h>
#include <errno.h>
#include <rtdk.h>
#include <rtdm/rtipc.h>

pthread_t svtid, cltid;

#define BUFP_SVSPORT 12

static const char *msg[] = {
    "Surfing With The Alien",
    "Lords of Karma",
    "Banana Mango",
    "Psycho Monkey",
    "Luminous Flesh Giants",
    "Moroccan Sunset",
    "Satch Boogie",
    "Flying In A Blue Dream",
    "Ride",
    "Summer Song",
    "Speed Of Light",
    "Crystal Planet",
    "Raspberry Jam Delta-V",
    "Champagne?",
    "Clouds Race Across The Sky",
    "Engines Of Creation"
};

static void fail(const char *reason)
{
    perror(reason);
    exit(EXIT_FAILURE);
}

static void *server(void *arg)
{
    struct sockaddr_ipc saddr;
    char buf[128];
    size_t bufsz;
    int ret, s;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_BUFP);
    if (s < 0)
        fail("socket");

    /*
     * Set a 16k buffer for the server endpoint. This
     * configuration must be done prior to binding the socket to a
     * port.
     */
    bufsz = 16384; /* bytes */
    ret = setsockopt(s, SOL_BUFP, BUFP_BUFSZ,
        &bufsz, sizeof(bufsz));
    if (ret)
        fail("setsockopt");

    saddr.sipc_family = AF_RTIPC;
    saddr.sipc_port = BUFP_SVSPORT;
    ret = bind(s, (struct sockaddr *)&saddr, sizeof(saddr));
    if (ret)
        fail("bind");

    for (;;) {
        ret = read(s, buf, sizeof(buf));
        if (ret < 0) {
            close(s);
            fail("read");
        }
        rt_printf("%s: received %d bytes, \\\"%.*s\\\"\\n",
            __FUNCTION__, ret, ret, buf);
    }

    return NULL;
}

```

```

}

static void *client(void *arg)
{
    struct sockaddr_ipc svaddr;
    int ret, s, n = 0, len;
    struct timespec ts;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_BUF);
    if (s < 0)
        fail("socket");

    memset(&svaddr, 0, sizeof(svaddr));
    svaddr.sipc_family = AF_RTIPC;
    svaddr.sipc_port = BUFP_SVPORT;
    ret = connect(s, (struct sockaddr *)&svaddr, sizeof(svaddr));
    if (ret)
        fail("connect");

    for (;;) {
        len = strlen(msg[n]);
        ret = write(s, msg[n], len);
        if (ret < 0) {
            close(s);
            fail("write");
        }
        rt_printf("%s: sent %d bytes, \"%s\"\n",
            __FUNCTION__, ret, ret, msg[n]);
        n = (n + 1) % (sizeof(msg) / sizeof(msg[0]));
        /*
         * We run in full real-time mode (i.e. primary mode),
         * so we have to let the system breathe between two
         * iterations.
         */
        ts.tv_sec = 0;
        ts.tv_nsec = 500000000; /* 500 ms */
        clock_nanosleep(CLOCK_REALTIME, 0, &ts, NULL);
    }

    return NULL;
}

static void cleanup_upon_sig(int sig)
{
    pthread_cancel(svtid);
    pthread_cancel(cltid);
    signal(sig, SIG_DFL);
    pthread_join(svtid, NULL);
    pthread_join(cltid, NULL);
}

int main(int argc, char **argv)
{
    struct sched_param svparam = {.sched_priority = 71 };
    struct sched_param clparam = {.sched_priority = 70 };
    pthread_attr_t svattr, clattr;
    sigset_t mask, oldmask;

    mlockall(MCL_CURRENT | MCL_FUTURE);

    sigemptyset(&mask);
    sigaddset(&mask, SIGINT);
    signal(SIGINT, cleanup_upon_sig);
    sigaddset(&mask, SIGTERM);
    signal(SIGTERM, cleanup_upon_sig);
    sigaddset(&mask, SIGHUP);
    signal(SIGHUP, cleanup_upon_sig);
    pthread_sigmask(SIG_BLOCK, &mask, &oldmask);

    /*
     * This is a real-time compatible printf() package from
     * Xenomai's RT Development Kit (RTDK), that does NOT cause
     * any transition to secondary mode.
     */
    rt_print_auto_init(1);

    pthread_attr_init(&svattr);
    pthread_attr_setdetachstate(&svattr, PTHREAD_CREATE_JOINABLE);
    pthread_attr_setinheritsched(&svattr, PTHREAD_EXPLICIT_SCHED);
    pthread_attr_setschedpolicy(&svattr, SCHED_FIFO);
    pthread_attr_setschedparam(&svattr, &svparam);

    errno = pthread_create(&svtid, &svattr, &server, NULL);
    if (errno)
        fail("pthread_create");

    pthread_attr_init(&clattr);

```

```

pthread_attr_setdetachstate(&clattr, PTHREAD_CREATE_JOINABLE);
pthread_attr_setinheritsched(&clattr, PTHREAD_EXPLICIT_SCHED);
pthread_attr_setschedpolicy(&clattr, SCHED_FIFO);
pthread_attr_setschedparam(&clattr, &clparam);

errno = pthread_create(&cltid, &clattr, &client, NULL);
if (errno)
    fail("pthread_create");

sigsuspend(&oldmask);

return 0;
}

```

## 8.3 cross-link.c

```

/*
 * cross-link.c
 *
 * Userspace test program (Xenomai native skin) for RTDM-based UART drivers
 * Copyright 2005 by Joerg Langenberg <joergel75@gmx.net>
 *
 * Updates by Jan Kiszka <jan.kiszka@web.de>
 *
 * This program is free software; you can redistribute it and/or modify
 * it under the terms of the GNU General Public License as published by
 * the Free Software Foundation; either version 2 of the License, or
 * (at your option) any later version.
 *
 * This program is distributed in the hope that it will be useful,
 * but WITHOUT ANY WARRANTY; without even the implied warranty of
 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
 * GNU General Public License for more details.
 *
 * You should have received a copy of the GNU General Public License
 * along with this program; if not, write to the Free Software
 * Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
 */
#include <stdio.h>
#include <signal.h>
#include <unistd.h>
#include <sys/mman.h>

#include <native/task.h>
#include <native/timer.h>

#include <rtdm/rtserial.h>

#define MAIN_PREFIX "main : "
#define WTASK_PREFIX "write_task: "
#define RTASK_PREFIX "read_task: "

#define WRITE_FILE "rtser0"
#define READ_FILE "rtser1"

int read_fd = -1;
int write_fd = -1;

#define STATE_FILE_OPENED 1
#define STATE_TASK_CREATED 2

unsigned int read_state = 0;
unsigned int write_state = 0;

/* --s-ms-us-ns */
RTIME write_task_period_ns = 100000000llu;
RT_TASK write_task;
RT_TASK read_task;

static const struct rtser_config read_config = {
    .config_mask = 0xFFFF,
    .baud_rate = 115200,
    .parity = RTSER_DEF_PARITY,
    .data_bits = RTSER_DEF_BITS,
    .stop_bits = RTSER_DEF_STOPB,
    .handshake = RTSER_DEF_HAND,
    .fifo_depth = RTSER_DEF_FIFO_DEPTH,
    .rx_timeout = RTSER_DEF_TIMEOUT,
    .tx_timeout = RTSER_DEF_TIMEOUT,
    .event_timeout = 100000000, /* 1 s */
    .timestamp_history = RTSER_RX_TIMESTAMP_HISTORY,
    .event_mask = RTSER_EVENT_RXPEND,
};

```



```

static const struct rtser_config write_config = {
    .config_mask      = RTSER_SET_BAUD | RTSER_SET_TIMESTAMP_HISTORY,
    .baud_rate        = 115200,
    .timestamp_history = RTSER_DEF_TIMESTAMP_HISTORY,
    /* the rest implicitly remains default */
};

static int close_file( int fd, char *name)
{
    int err, i=0;

    do {
        i++;
        err = rt_dev_close(fd);
        switch (err) {
            case -EAGAIN:
                printf(MAIN_PREFIX "%s -> EAGAIN (%d times)\n",
                    name, i);
                rt_task_sleep(50000); /* wait 50us */
                break;
            case 0:
                printf(MAIN_PREFIX "%s -> closed\n", name);
                break;
            default:
                printf(MAIN_PREFIX "%s -> %s\n", name,
                    strerror(-err));
                break;
        }
    } while (err == -EAGAIN && i < 10);

    return err;
}

void cleanup_all(void)
{
    if (read_state & STATE_FILE_OPENED) {
        close_file(read_fd, READ_FILE " (read)");
        read_state &= ~STATE_FILE_OPENED;
    }

    if (write_state & STATE_FILE_OPENED) {
        close_file(write_fd, WRITE_FILE " (write)");
        write_state &= ~STATE_FILE_OPENED;
    }

    if (write_state & STATE_TASK_CREATED) {
        printf(MAIN_PREFIX "delete write_task\n");
        rt_task_delete(&write_task);
        write_state &= ~STATE_TASK_CREATED;
    }

    if (read_state & STATE_TASK_CREATED) {
        printf(MAIN_PREFIX "delete read_task\n");
        rt_task_delete(&read_task);
        read_state &= ~STATE_TASK_CREATED;
    }
}

void catch_signal(int sig)
{
    cleanup_all();
    printf(MAIN_PREFIX "exit\n");
    return;
}

void write_task_proc(void *arg)
{
    int err;
    RTIME write_time;
    ssize_t sz = sizeof(RTIME);
    ssize_t written = 0;

    err = rt_task_set_periodic(NULL, TM_NOW,
        rt_timer_ns2ticks(write_task_period_ns));
    if (err) {
        printf(WTASK_PREFIX "error on set periodic, %s\n",
            strerror(-err));
        goto exit_write_task;
    }

    while (1) {
        err = rt_task_wait_period(NULL);
        if (err) {
            printf(WTASK_PREFIX
                "error on rt_task_wait_period, %s\n",
                strerror(-err));
        }
    }
}

```

```

        }
        break;
    }

    write_time = rt_timer_read();

    written = rt_dev_write(write_fd, &write_time, sz);
    if (written < 0 ) {
        printf(WTASK_PREFIX "error on rt_dev_write, %s\n",
               strerror(-err));
        break;
    } else if (written != sz) {
        printf(WTASK_PREFIX "only %d / %d byte transmitted\n",
               written, sz);
        break;
    }
}

exit_write_task:
if ((write_state & STATE_FILE_OPENED) &&
    close_file(write_fd, WRITE_FILE " (write)") == 0)
    write_state &= ~STATE_FILE_OPENED;

printf(WTASK_PREFIX "exit\n");
}

void read_task_proc(void *arg)
{
    int err;
    int nr = 0;
    RTIME read_time = 0;
    RTIME write_time = 0;
    RTIME irq_time = 0;
    ssize_t sz = sizeof(RTIME);
    ssize_t read = 0;
    struct rtser_event rx_event;

    printf(" Nr | write->irq | irq->read | write->read | \n");
    printf("-----\n");

    /*
     * We are in secondary mode now due to printf, the next
     * blocking Xenomai or driver call will switch us back
     * (here: RTSER_RTIOC_WAIT_EVENT).
     */

    while (1) {
        /* waiting for event */
        err = rt_dev_ioctl(read_fd, RTSER_RTIOC_WAIT_EVENT, &rx_event);
        if (err) {
            printf(RTASK_PREFIX
                   "error on RTSER_RTIOC_WAIT_EVENT, %s\n",
                   strerror(-err));
            if (err == -ETIMEDOUT)
                continue;
            break;
        }

        irq_time = rx_event.rxpnd_timestamp;
        read = rt_dev_read(read_fd, &write_time, sz);
        if (read == sz) {
            read_time = rt_timer_read();
            printf("%3d |%16llu |%16llu |%16llu\n", nr,
                   irq_time - write_time,
                   read_time - irq_time,
                   read_time - write_time);
            nr++;
        } else if (read < 0 ) {
            printf(RTASK_PREFIX "error on rt_dev_read, code %s\n",
                   strerror(-err));
            break;
        } else {
            printf(RTASK_PREFIX "only %d / %d byte received \n",
                   read, sz);
            break;
        }
    }

    if ((read_state & STATE_FILE_OPENED) &&
        close_file(read_fd, READ_FILE " (read)") == 0)
        read_state &= ~STATE_FILE_OPENED;

    printf(RTASK_PREFIX "exit\n");
}

int main(int argc, char* argv[])
{
    int err = 0;

```

```

signal(SIGTERM, catch_signal);
signal(SIGINT, catch_signal);

/* no memory-swapping for this programm */
mlockall(MCL_CURRENT | MCL_FUTURE);

/* open rtser0 */
write_fd = rt_dev_open( WRITE_FILE, 0);
if (write_fd < 0) {
    printf(MAIN_PREFIX "can't open %s (write), %s\n", WRITE_FILE,
           strerror(-write_fd));
    goto error;
}
write_state |= STATE_FILE_OPENED;
printf(MAIN_PREFIX "write-file opened\n");

/* writing write-config */
err = rt_dev_ioctl(write_fd, RTSER_RTIOC_SET_CONFIG, &write_config);
if (err) {
    printf(MAIN_PREFIX "error while RTSER_RTIOC_SET_CONFIG, %s\n",
           strerror(-err));
    goto error;
}
printf(MAIN_PREFIX "write-config written\n");

/* open rtser1 */
read_fd = rt_dev_open( READ_FILE, 0 );
if (read_fd < 0) {
    printf(MAIN_PREFIX "can't open %s (read), %s\n", READ_FILE,
           strerror(-read_fd));
    goto error;
}
read_state |= STATE_FILE_OPENED;
printf(MAIN_PREFIX "read-file opened\n");

/* writing read-config */
err = rt_dev_ioctl(read_fd, RTSER_RTIOC_SET_CONFIG, &read_config);
if (err) {
    printf(MAIN_PREFIX "error while rt_dev_ioctl, %s\n",
           strerror(-err));
    goto error;
}
printf(MAIN_PREFIX "read-config written\n");

/* create write_task */
err = rt_task_create(&write_task, "write_task", 0, 50, 0);
if (err) {
    printf(MAIN_PREFIX "failed to create write_task, %s\n",
           strerror(-err));
    goto error;
}
write_state |= STATE_TASK_CREATED;
printf(MAIN_PREFIX "write-task created\n");

/* create read_task */
err = rt_task_create(&read_task, "read_task", 0, 51, 0);
if (err) {
    printf(MAIN_PREFIX "failed to create read_task, %s\n",
           strerror(-err));
    goto error;
}
read_state |= STATE_TASK_CREATED;
printf(MAIN_PREFIX "read-task created\n");

/* start write_task */
printf(MAIN_PREFIX "starting write-task\n");
err = rt_task_start(&write_task, &write_task_proc, NULL);
if (err) {
    printf(MAIN_PREFIX "failed to start write_task, %s\n",
           strerror(-err));
    goto error;
}

/* start read_task */
printf(MAIN_PREFIX "starting read-task\n");
err = rt_task_start(&read_task, &read_task_proc, NULL);
if (err) {
    printf(MAIN_PREFIX "failed to start read_task, %s\n",
           strerror(-err));
    goto error;
}

pause();
return 0;

```

error:

```

        cleanup_all();
        return err;
}

```

## 8.4 iddp-label.c

```

/*
 * IDDP-based client/server demo, using the write(2)/recvfrom(2)
 * system calls to exchange data over a socket.
 *
 * In this example, two sockets are created. A server thread (reader)
 * is bound to a labeled real-time port and receives datagrams sent to
 * this port from a client thread (writer). The client thread attaches
 * to the port opened by the server using a labeled connection
 * request. The client socket is bound to a different port, only to
 * provide a valid peer name; this is optional.
 *
 * ASCII labels can be attached to bound ports, in order to connect
 * sockets to them in a more descriptive way than using plain numeric
 * port values.
 *
 * See Makefile in this directory for build directives.
 */
#include <sys/mman.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <pthread.h>
#include <errno.h>
#include <rtdk.h>
#include <rtdm/rtipc.h>

pthread_t svtid, cltid;

#define IDDP_CLPORT 27

#define IDDP_PORT_LABEL "iddp-demo"

static const char *msg[] = {
    "Surfing With The Alien",
    "Lords of Karma",
    "Banana Mango",
    "Psycho Monkey",
    "Luminous Flesh Giants",
    "Moroccan Sunset",
    "Satch Boogie",
    "Flying In A Blue Dream",
    "Ride",
    "Summer Song",
    "Speed Of Light",
    "Crystal Planet",
    "Raspberry Jam Delta-V",
    "Champagne?",
    "Clouds Race Across The Sky",
    "Engines Of Creation"
};

static void fail(const char *reason)
{
    perror(reason);
    exit(EXIT_FAILURE);
}

static void *server(void *arg)
{
    struct sockaddr_ipc saddr, claddr;
    struct rtipc_port_label plabel;
    socklen_t addrlen;
    char buf[128];
    int ret, s;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_IDDP);
    if (s < 0)
        fail("socket");

    /*
     * We will use Xenomai's system heap for datagram, so no
     * IDDP_POOLSZ required here.
     */

    /*

```

```

    * Set a port label. This name will be registered when
    * binding, in addition to the port number (if given).
    */
strcpy(plabel.label, IDDP_PORT_LABEL);
ret = setsockopt(s, SOL_IDDP, IDDP_LABEL,
                &plabel, sizeof(plabel));
if (ret)
    fail("setsockopt");

/*
 * Bind the socket to the port. Assign that port a label, so
 * that peers may use a descriptive information to locate
 * it. Labeled ports will appear in the
 * /proc/xenomai/registry/rtipc/iddp directory once the socket
 * is bound.
 */
saddr.sipc_port specifies the port number to use. If -1 is
 * passed, the IDDP driver will auto-select an idle port.
 */
saddr.sipc_family = AF_RTIPC;
saddr.sipc_port = -1; /* Pick next free */
ret = bind(s, (struct sockaddr *)&saddr, sizeof(saddr));
if (ret)
    fail("bind");

for (;;) {
    addrlen = sizeof(saddr);
    ret = recvfrom(s, buf, sizeof(buf), 0,
                  (struct sockaddr *)&claddr, &addrlen);
    if (ret < 0) {
        close(s);
        fail("recvfrom");
    }
    rt_printf("%s: received %d bytes, \"%.*s\" from port %d\n",
              __FUNCTION__, ret, ret, buf, claddr.sipc_port);
}

return NULL;
}

static void *client(void *arg)
{
    struct sockaddr_ipc svaddr, claddr;
    struct rtipc_port_label plabel;
    int ret, s, n = 0, len;
    struct timespec ts;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_IDDP);
    if (s < 0)
        fail("socket");

    /*
     * Set a name on the client socket. This is strictly optional,
     * and only done here for the purpose of getting back a
     * different port number in recvfrom().
     */
    claddr.sipc_family = AF_RTIPC;
    claddr.sipc_port = IDDP_CLPORT;
    ret = bind(s, (struct sockaddr *)&claddr, sizeof(claddr));
    if (ret)
        fail("bind");

    /*
     * Set the port label. This name will be used to find the peer
     * when connecting, instead of the port number. The label must
     * be set _after_ the socket is bound to the port, so that
     * IDDP does not try to register this label for the client
     * port as well (like the server thread did).
     */
    strcpy(plabel.label, IDDP_PORT_LABEL);
    ret = setsockopt(s, SOL_IDDP, IDDP_LABEL,
                    &plabel, sizeof(plabel));
    if (ret)
        fail("setsockopt");

    memset(&svaddr, 0, sizeof(svaddr));
    svaddr.sipc_family = AF_RTIPC;
    svaddr.sipc_port = -1; /* Tell IDDP to search by label. */
    ret = connect(s, (struct sockaddr *)&svaddr, sizeof(svaddr));
    if (ret)
        fail("connect");

    for (;;) {
        len = strlen(msg[n]);
        /* Send to default destination we connected to. */
        ret = write(s, msg[n], len);
        if (ret < 0) {

```

```

        close(s);
        fail("sendto");
    }
    rt_printf("%s: sent %d bytes, \".*s\"\n",
        __FUNCTION__, ret, msg[n]);
    n = (n + 1) % (sizeof(msg) / sizeof(msg[0]));
    /*
     * We run in full real-time mode (i.e. primary mode),
     * so we have to let the system breathe between two
     * iterations.
     */
    ts.tv_sec = 0;
    ts.tv_nsec = 500000000; /* 500 ms */
    clock_nanosleep(CLOCK_REALTIME, 0, &ts, NULL);
}

return NULL;
}

static void cleanup_upon_sig(int sig)
{
    pthread_cancel(svtid);
    pthread_cancel(cltid);
    signal(sig, SIG_DFL);
    pthread_join(svtid, NULL);
    pthread_join(cltid, NULL);
}

int main(int argc, char **argv)
{
    struct sched_param svparam = {.sched_priority = 71 };
    struct sched_param clparam = {.sched_priority = 70 };
    pthread_attr_t svattr, clattr;
    sigset_t mask, oldmask;

    mlockall(MCL_CURRENT | MCL_FUTURE);

    sigemptyset(&mask);
    sigaddset(&mask, SIGINT);
    signal(SIGINT, cleanup_upon_sig);
    sigaddset(&mask, SIGTERM);
    signal(SIGTERM, cleanup_upon_sig);
    sigaddset(&mask, SIGHUP);
    signal(SIGHUP, cleanup_upon_sig);
    pthread_sigmask(SIG_BLOCK, &mask, &oldmask);

    /*
     * This is a real-time compatible printf() package from
     * Xenomai's RT Development Kit (RTDK), that does NOT cause
     * any transition to secondary mode.
     */
    rt_print_auto_init(1);

    pthread_attr_init(&svattr);
    pthread_attr_setdetachstate(&svattr, PTHREAD_CREATE_JOINABLE);
    pthread_attr_setinheritsched(&svattr, PTHREAD_EXPLICIT_SCHED);
    pthread_attr_setschedpolicy(&svattr, SCHED_FIFO);
    pthread_attr_setschedparam(&svattr, &svparam);

    errno = pthread_create(&svtid, &svattr, &server, NULL);
    if (errno)
        fail("pthread_create");

    pthread_attr_init(&clattr);
    pthread_attr_setdetachstate(&clattr, PTHREAD_CREATE_JOINABLE);
    pthread_attr_setinheritsched(&clattr, PTHREAD_EXPLICIT_SCHED);
    pthread_attr_setschedpolicy(&clattr, SCHED_FIFO);
    pthread_attr_setschedparam(&clattr, &clparam);

    errno = pthread_create(&cltid, &clattr, &client, NULL);
    if (errno)
        fail("pthread_create");

    sigsuspend(&oldmask);

    return 0;
}

```

## 8.5 iddp-sendrecv.c

```

/*
 * IDDP-based client/server demo, using the sendto(2)/recvfrom(2)
 * system calls to exchange data over a socket.

```

```

*
* In this example, two sockets are created. A server thread (reader)
* is bound to a real-time port and receives datagrams sent to this
* port from a client thread (writer). The client socket is bound to a
* different port, only to provide a valid peer name; this is
* optional.
*
* See Makefile in this directory for build directives.
*/
#include <sys/mman.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <pthread.h>
#include <errno.h>
#include <rtdk.h>
#include <rtdm/rtipc.h>

pthread_t svtid, cltid;

#define IDDP_SVSPORT 12
#define IDDP_CLPORT 13

static const char *msg[] = {
    "Surfing With The Alien",
    "Lords of Karma",
    "Banana Mango",
    "Psycho Monkey",
    "Luminous Flesh Giants",
    "Moroccan Sunset",
    "Satch Boogie",
    "Flying In A Blue Dream",
    "Ride",
    "Summer Song",
    "Speed Of Light",
    "Crystal Planet",
    "Raspberry Jam Delta-V",
    "Champagne?",
    "Clouds Race Across The Sky",
    "Engines Of Creation"
};

static void fail(const char *reason)
{
    perror(reason);
    exit(EXIT_FAILURE);
}

static void *server(void *arg)
{
    struct sockaddr_ipc saddr, claddr;
    socklen_t addrlen;
    char buf[128];
    size_t poolsz;
    int ret, s;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_IDDP);
    if (s < 0)
        fail("socket");

    /*
     * Set a local 32k pool for the server endpoint. Memory needed
     * to convey datagrams will be pulled from this pool, instead
     * of Xenomai's system pool.
     */
    poolsz = 32768; /* bytes */
    ret = setsockopt(s, SOL_IDDP, IDDP_POOLSZ,
                    &poolsz, sizeof(poolsz));
    if (ret)
        fail("setsockopt");

    saddr.sipc_family = AF_RTIPC;
    saddr.sipc_port = IDDP_SVSPORT;
    ret = bind(s, (struct sockaddr *)&saddr, sizeof(saddr));
    if (ret)
        fail("bind");

    for (;;) {
        addrlen = sizeof(saddr);
        ret = recvfrom(s, buf, sizeof(buf), 0,
                     (struct sockaddr *)&claddr, &addrlen);
        if (ret < 0) {
            close(s);
            fail("recvfrom");
        }
    }
}

```

```

        rt_printf("%s: received %d bytes, \"%.*s\" from port %d\n",
                __FUNCTION__, ret, ret, buf, claddr.sipc_port);
    }

    return NULL;
}

static void *client(void *arg)
{
    struct sockaddr_ipc svaddr, claddr;
    int ret, s, n = 0, len;
    struct timespec ts;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_IDDP);
    if (s < 0)
        fail("socket");

    claddr.sipc_family = AF_RTIPC;
    claddr.sipc_port = IDDP_CLPORT;
    ret = bind(s, (struct sockaddr *)&claddr, sizeof(claddr));
    if (ret)
        fail("bind");

    svaddr.sipc_family = AF_RTIPC;
    svaddr.sipc_port = IDDP_SVPORT;
    for (;;) {
        len = strlen(msg[n]);
        ret = sendto(s, msg[n], len, 0,
                    (struct sockaddr *)&svaddr, sizeof(svaddr));
        if (ret < 0) {
            close(s);
            fail("sendto");
        }
        rt_printf("%s: sent %d bytes, \"%.*s\"\\n",
                __FUNCTION__, ret, ret, msg[n]);
        n = (n + 1) % (sizeof(msg) / sizeof(msg[0]));
        /*
         * We run in full real-time mode (i.e. primary mode),
         * so we have to let the system breathe between two
         * iterations.
         */
        ts.tv_sec = 0;
        ts.tv_nsec = 500000000; /* 500 ms */
        clock_nanosleep(CLOCK_REALTIME, 0, &ts, NULL);
    }

    return NULL;
}

static void cleanup_upon_sig(int sig)
{
    pthread_cancel(svtid);
    pthread_cancel(cltid);
    signal(sig, SIG_DFL);
    pthread_join(svtid, NULL);
    pthread_join(cltid, NULL);
}

int main(int argc, char **argv)
{
    struct sched_param svparam = {.sched_priority = 71 };
    struct sched_param clparam = {.sched_priority = 70 };
    pthread_attr_t svattr, clattr;
    sigset_t mask, oldmask;

    mlockall(MCL_CURRENT | MCL_FUTURE);

    sigemptyset(&mask);
    sigaddset(&mask, SIGINT);
    signal(SIGINT, cleanup_upon_sig);
    sigaddset(&mask, SIGTERM);
    signal(SIGTERM, cleanup_upon_sig);
    sigaddset(&mask, SIGHUP);
    signal(SIGHUP, cleanup_upon_sig);
    pthread_sigmask(SIG_BLOCK, &mask, &oldmask);

    /*
     * This is a real-time compatible printf() package from
     * Xenomai's RT Development Kit (RTDK), that does NOT cause
     * any transition to secondary mode.
     */
    rt_print_auto_init(1);

    pthread_attr_init(&svattr);
    pthread_attr_setdetachstate(&svattr, PTHREAD_CREATE_JOINABLE);
    pthread_attr_setinheritsched(&svattr, PTHREAD_EXPLICIT_SCHED);
    pthread_attr_setschedpolicy(&svattr, SCHED_FIFO);

```



```

pthread_attr_setschedparam(&svattr, &svparam);

errno = pthread_create(&svtid, &svattr, &server, NULL);
if (errno)
    fail("pthread_create");

pthread_attr_init(&clattr);
pthread_attr_setdetachstate(&clattr, PTHREAD_CREATE_JOINABLE);
pthread_attr_setinheritsched(&clattr, PTHREAD_EXPLICIT_SCHED);
pthread_attr_setschedpolicy(&clattr, SCHED_FIFO);
pthread_attr_setschedparam(&clattr, &clparam);

errno = pthread_create(&cltid, &clattr, &client, NULL);
if (errno)
    fail("pthread_create");

sigsuspend(&oldmask);

return 0;
}

```

## 8.6 rtcan\_rtt.c

```

/*
 * Round-Trip-Time Test - sends and receives messages and measures the
 * time in between.
 *
 * Copyright (C) 2006 Wolfgang Grandegger <wg@grandegger.com>
 *
 * Based on RTnet's examples/xenomai/posix/rtt-sender.c.
 *
 * Copyright (C) 2002 Ulrich Marx <marx@kammer.uni-hannover.de>
 *      2002 Marc Kleine-Budde <kleine-budde@gmx.de>
 *      2006 Jan Kiszka <jan.kiszka@web.de>
 *
 * This program is free software; you can redistribute it and/or modify
 * it under the terms of the GNU General Public License as published by
 * the Free Software Foundation; either version 2 of the License, or
 * (at your option) any later version.
 *
 * This program is distributed in the hope that it will be useful,
 * but WITHOUT ANY WARRANTY; without even the implied warranty of
 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
 * GNU General Public License for more details.
 *
 * You should have received a copy of the GNU General Public License
 * along with this program; if not, write to the Free Software
 * Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
 *
 * The program sends out CAN messages periodically and copies the current
 * time-stamp to the payload. At reception, that time-stamp is compared
 * with the current time to determine the round-trip time. The jitter
 * values are printed out regularly. Concurrent tests can be carried out
 * by starting the program with different message identifiers. It is also
 * possible to use this program on a remote system as simple repeater to
 * loopback messages.
 */

#include <errno.h>
#include <mqueue.h>
#include <signal.h>
#include <pthread.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <unistd.h>
#include <limits.h>
#include <getopt.h>
#include <netinet/in.h>
#include <net/if.h>
#include <sys/ioctl.h>
#include <sys/mman.h>

#ifdef __XENO__
#include <rtm/rtcan.h>
#else
#include <linux/can.h>
#include <linux/can/raw.h>
#endif

#define NSEC_PER_SEC 1000000000

```

```

static unsigned int cycle = 10000; /* 10 ms */
static canid_t can_id = 0x1;

static pthread_t txthread, rxthread;
static int txsock, rxsock;
static mqd_t mq;
static int txcount, rxcount;
static int overruns;
static int repeater;

struct rtt_stat {
    long long rtt;
    long long rtt_min;
    long long rtt_max;
    long long rtt_sum;
    long long rtt_sum_last;
    int counts_per_sec;
};

static void print_usage(char *prg)
{
    fprintf(stderr,
        "Usage: %s [Options] <tx-can-interface> <rx-can-interface>\n"
        "Options:\n"
        " -h, --help      This help\n"
        " -r, --repeater  Repeater, send back received messages\n"
        " -i, --id=ID     CAN Identifier (default = 0x1)\n"
        " -c, --cycle     Cycle time in us (default = 10000us)\n",
        prg);
}

void *transmitter(void *arg)
{
    struct sched_param param = { .sched_priority = 80 };
    struct timespec next_period;
    struct timespec time;
    struct can_frame frame;
    long long *rtt_time = (long long *)&frame.data;

    /* Pre-fill CAN frame */
    frame.can_id = can_id;
    frame.can_dlc = sizeof(*rtt_time);

#ifdef __XENO__
    pthread_set_name_np(pthread_self(), "rtcan_rtt_transmitter");
#endif
    pthread_setschedparam(pthread_self(), SCHED_FIFO, &param);

    clock_gettime(CLOCK_MONOTONIC, &next_period);

    while(1) {
        next_period.tv_nsec += cycle * 1000;
        while (next_period.tv_nsec >= NSEC_PER_SEC) {
            next_period.tv_nsec -= NSEC_PER_SEC;
            next_period.tv_sec++;
        }

        clock_nanosleep(CLOCK_MONOTONIC, TIMER_ABSTIME, &next_period, NULL);

        if (rxcount != txcount) {
            overruns++;
            continue;
        }

        clock_gettime(CLOCK_MONOTONIC, &time);
        *rtt_time = (long long)time.tv_sec * NSEC_PER_SEC + time.tv_nsec;

        /* Transmit the message containing the local time */
        if (send(txsock, (void *)&frame, sizeof(struct can_frame), 0) < 0) {
            if (errno == EBADF)
                printf("terminating transmitter thread\n");
            else
                perror("send failed");
            return NULL;
        }
        txcount++;
    }
}

void *receiver(void *arg)
{
    struct sched_param param = { .sched_priority = 82 };
    struct timespec time;
    struct can_frame frame;
    long long *rtt_time = (long long *)&frame.data;
    struct rtt_stat rtt_stat = {0, 100000000000000000LL, -100000000000000000LL,

```

```

        0, 0, 0};

#ifdef __XENO__
    pthread_set_name_np(pthread_self(), "rtcan_rtt_receiver");
#endif
pthread_setschedparam(pthread_self(), SCHED_FIFO, &param);

rtt_stat.counts_per_sec = 1000000 / cycle;

while (1) {
    if (recv(rxsock, (void *)&frame, sizeof(struct can_frame), 0) < 0) {
        if (errno == EBADF)
            printf("terminating receiver thread\n");
        else
            perror("recv failed");
        return NULL;
    }
    if (repeater) {
        /* Transmit the message back as is */
        if (send(txsock, (void *)&frame, sizeof(struct can_frame), 0) < 0) {
            if (errno == EBADF)
                printf("terminating transmitter thread\n");
            else
                perror("send failed");
            return NULL;
        }
        txcount++;
    } else {
        clock_gettime(CLOCK_MONOTONIC, &time);
        if (rxcount > 0) {
            rtt_stat.rtt = ((long long)time.tv_sec * 1000000000LL +
                time.tv_nsec - *rtt_time);
            rtt_stat.rtt_sum += rtt_stat.rtt;
            if (rtt_stat.rtt < rtt_stat.rtt_min)
                rtt_stat.rtt_min = rtt_stat.rtt;
            if (rtt_stat.rtt > rtt_stat.rtt_max)
                rtt_stat.rtt_max = rtt_stat.rtt;
        }
        rxcount++;

        if ((rxcount % rtt_stat.counts_per_sec) == 0) {
            mq_send(mq, (char *)&rtt_stat, sizeof(rtt_stat), 0);
            rtt_stat.rtt_sum_last = rtt_stat.rtt_sum;
        }
    }
}

void catch_signal(int sig)
{
    mq_close(mq);
}

int main(int argc, char *argv[])
{
    struct sched_param param = { .sched_priority = 1 };
    pthread_attr_t thattr;
    struct mq_attr mqattr;
    struct sockaddr_can rxaddr, txaddr;
    struct can_filter rxfilter[1];
    struct rtt_stat rtt_stat;
    char mqname[32];
    char *txdev, *rxdev;
    struct ifreq ifr;
    int ret, opt;

    struct option long_options[] = {
        { "id", required_argument, 0, 'i'},
        { "cycle", required_argument, 0, 'c'},
        { "repeater", no_argument, 0, 'r'},
        { "help", no_argument, 0, 'h'},
        { 0, 0, 0, 0},
    };

    while ((opt = getopt_long(argc, argv, "hri:c:",
        long_options, NULL)) != -1) {
        switch (opt) {
            case 'c':
                cycle = atoi(optarg);
                break;

            case 'i':
                can_id = strtoul(optarg, NULL, 0);
                break;

            case 'r':

```

```

        repeater = 1;
        break;

default:
    fprintf(stderr, "Unknown option %c\n", opt);
case 'h':
    print_usage(argv[0]);
    exit(-1);
}
}

printf("%d %d\n", optind, argc);
if (optind + 2 != argc) {
    print_usage(argv[0]);
    exit(0);
}

txdev = argv[optind];
rxdev = argv[optind + 1];

/* Create and configure RX socket */
if ((rxsock = socket(PF_CAN, SOCK_RAW, CAN_RAW)) < 0) {
    perror("RX socket failed");
    return -1;
}

strncpy(ifr.ifr_name, rxdev, IFNAMSIZ);
printf("RX rxsock=%d, ifr_name=%s\n", rxsock, ifr.ifr_name);

if (ioctl(rxsock, SIOCGIFINDEX, &ifr) < 0) {
    perror("RX ioctl SIOCGIFINDEX failed");
    goto failure1;
}

/* We only want to receive our own messages */
rxfilter[0].can_id = can_id;
rxfilter[0].can_mask = 0x3fff;
if (setsockopt(rxsock, SOL_CAN_RAW, CAN_RAW_FILTER,
               &rxfilter, sizeof(struct can_filter)) < 0) {
    perror("RX setsockopt CAN_RAW_FILTER failed");
    goto failure1;
}
memset(&rxaddr, 0, sizeof(rxaddr));
rxaddr.can_ifindex = ifr.ifr_ifindex;
rxaddr.can_family = AF_CAN;
if (bind(rxsock, (struct sockaddr *)&rxaddr, sizeof(rxaddr)) < 0) {
    perror("RX bind failed\n");
    goto failure1;
}

/* Create and configure TX socket */
if (strcmp(rxdev, txdev) == 0) {
    txsock = rxsock;
} else {
    if ((txsock = socket(PF_CAN, SOCK_RAW, 0)) < 0) {
        perror("TX socket failed");
        goto failure1;
    }

    strncpy(ifr.ifr_name, txdev, IFNAMSIZ);
    printf("TX txsock=%d, ifr_name=%s\n", txsock, ifr.ifr_name);

    if (ioctl(txsock, SIOCGIFINDEX, &ifr) < 0) {
        perror("TX ioctl SIOCGIFINDEX failed");
        goto failure2;
    }

    /* Suppress definition of a default receive filter list */
    if (setsockopt(txsock, SOL_CAN_RAW, CAN_RAW_FILTER, NULL, 0) < 0) {
        perror("TX setsockopt CAN_RAW_FILTER failed");
        goto failure2;
    }

    memset(&txaddr, 0, sizeof(txaddr));
    txaddr.can_ifindex = ifr.ifr_ifindex;
    txaddr.can_family = AF_CAN;

    if (bind(txsock, (struct sockaddr *)&txaddr, sizeof(txaddr)) < 0) {
        perror("TX bind failed\n");
        goto failure2;
    }
}

signal(SIGTERM, catch_signal);
signal(SIGINT, catch_signal);
signal(SIGHUP, catch_signal);

```

```

mlockall(MCL_CURRENT|MCL_FUTURE);

printf("Round-Trip-Time test %s -> %s with CAN ID 0x%x\n",
      argv[optind], argv[optind + 1], can_id);
printf("Cycle time: %d us\n", cycle);
printf("All RTT timing figures are in us.\n");

/* Create statistics message queue */
snprintf(mqname, sizeof(mqname), "/rtcan_rtt-%d", getpid());
mqattr.mq_flags = 0;
mqattr.mq_maxmsg = 100;
mqattr.mq_msgsize = sizeof(struct rtt_stat);
mq = mq_open(mqname, O_RDWR | O_CREAT | O_EXCL, 0600, &mqattr);
if (mq == (mqd_t)-1) {
    perror("opening mqueue failed");
    goto failure2;
}

/* Create receiver RT-thread */
pthread_attr_init(&thattr);
pthread_attr_setdetachstate(&thattr, PTHREAD_CREATE_JOINABLE);
pthread_attr_setstacksize(&thattr, PTHREAD_STACK_MIN);
ret = pthread_create(&rxthread, &thattr, &receiver, NULL);
if (ret) {
    fprintf(stderr, "%s: pthread_create(receiver) failed\n",
            strerror(-ret));
    goto failure3;
}

if (!repeater) {
    /* Create transitter RT-thread */
    ret = pthread_create(&txthread, &thattr, &transmitter, NULL);
    if (ret) {
        fprintf(stderr, "%s: pthread_create(transmitter) failed\n",
                strerror(-ret));
        goto failure4;
    }
}

pthread_setschedparam(pthread_self(), SCHED_FIFO, &param);

if (repeater)
    printf("Messages\n");
else
    printf("Messages RTTlast RTT_avg RTT_min RTT_max Overruns\n");

while (1) {
    long long rtt_avg;

    ret = mq_receive(mq, (char *)&rtt_stat, sizeof(rtt_stat), NULL);
    if (ret != sizeof(rtt_stat)) {
        if (ret < 0) {
            if (errno == EBADF)
                printf("terminating mq_receive\n");
            else
                perror("mq_receive failed");
        } else
            fprintf(stderr,
                    "mq_receive returned invalid length %d\n", ret);
        break;
    }

    if (repeater) {
        printf("%8d\n", rxcount);
    } else {
        rtt_avg = ((rtt_stat.rtt_sum - rtt_stat.rtt_sum_last) /
                  rtt_stat.counts_per_sec);
        printf("%8d %7ld %7ld %7ld %8d\n", rxcount,
              (long)(rtt_stat.rtt / 1000), (long)(rtt_avg / 1000),
              (long)(rtt_stat.rtt_min / 1000),
              (long)(rtt_stat.rtt_max / 1000),
              overruns);
    }
}

/* This call also leaves primary mode, required for socket cleanup. */
printf("shutting down\n");

/* Important: First close the sockets! */
while ((close(rxsock) < 0) && (errno == EAGAIN)) {
    printf("RX socket busy - waiting...\n");
    sleep(1);
}
while ((close(txsock) < 0) && (errno == EAGAIN)) {
    printf("TX socket busy - waiting...\n");
    sleep(1);
}

```

```

pthread_join(txthread, NULL);
pthread_kill(rxthread, SIGHUP);
pthread_join(rxthread, NULL);

return 0;

failure4:
pthread_kill(rxthread, SIGHUP);
pthread_join(rxthread, NULL);
failure3:
mq_close(mq);
failure2:
close(txsock);
failure1:
close(rxsock);

return 1;
}

```

## 8.7 rtcanconfig.c

```

/*
 * Program to configuring the CAN controller
 *
 * Copyright (C) 2006 Wolfgang Grandegger <wg@grandegger.com>
 *
 * Copyright (C) 2005, 2006 Sebastian Smolorz
 * <Sebastian.Smolorz@stud.uni-hannover.de>
 *
 *
 * This program is free software; you can redistribute it and/or modify
 * it under the terms of the GNU General Public License as published by
 * the Free Software Foundation; either version 2 of the License, or
 * (at your option) any later version.
 *
 * This program is distributed in the hope that it will be useful,
 * but WITHOUT ANY WARRANTY; without even the implied warranty of
 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
 * GNU General Public License for more details.
 *
 * You should have received a copy of the GNU General Public License
 * along with this program; if not, write to the Free Software
 * Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
 */

#include <stdio.h>
#include <stdlib.h>
#include <signal.h>
#include <unistd.h>
#include <string.h>
#include <time.h>
#include <errno.h>
#include <getopt.h>
#include <sys/mman.h>

#include <rtcm/rtdm/rtdm.h>

static void print_usage(char *prg)
{
    fprintf(stderr,
            "Usage: %s <can-interface> [Options] [up|down|start|stop|sleep]\n"
            "Options:\n"
            " -v, --verbose          be verbose\n"
            " -h, --help            this help\n"
            " -c, --ctrlmode=CTRLMODE listenonly, loopback or none\n"
            " -b, --baudrate=BPS    baudrate in bits/sec\n"
            " -B, --bittime=BTR0:BTR1 BTR or standard bit-time\n"
            " -B, --bittime=BRP:PROP_SEG:PHASE_SEG1:PHASE_SEG2:SJW:SAM\n",
            prg);
}

can_baudrate_t string_to_baudrate(char *str)
{
    can_baudrate_t baudrate;
    if (sscanf(str, "%i", &baudrate) != 1)
        return -1;
    return baudrate;
}

int string_to_mode(char *str)
{
    if ( !strcmp(str, "up") || !strcmp(str, "start") )

```

```

        return CAN_MODE_START;
    else if ( !strcmp(str, "down") || !strcmp(str, "stop") )
        return CAN_MODE_STOP;
    else if ( !strcmp(str, "sleep") )
        return CAN_MODE_SLEEP;
    return -EINVAL;
}

int string_to_ctrlmode(char *str)
{
    if ( !strcmp(str, "listenonly") )
        return CAN_CTRLMODE_LISTENONLY;
    else if ( !strcmp(str, "loopback") )
        return CAN_CTRLMODE_LOOPBACK;
    else if ( !strcmp(str, "none") )
        return 0;

    return -1;
}

int main(int argc, char *argv[])
{
    char    ifname[16];
    int     can_fd = -1;
    int     new_baudrate = -1;
    int     new_mode = -1;
    int     new_ctrlmode = 0, set_ctrlmode = 0;
    int     verbose = 0;
    int     bittime_count = 0, bittime_data[6];
    struct  ifreq ifr;
    can_baudrate_t *baudrate;
    can_ctrlmode_t *ctrlmode;
    can_mode_t *mode;
    struct  can_bittime *bittime;
    int     opt, ret;
    char*   ptr;

    struct option long_options[] = {
        { "help", no_argument, 0, 'h' },
        { "verbose", no_argument, 0, 'v' },
        { "baudrate", required_argument, 0, 'b' },
        { "bittime", required_argument, 0, 'B' },
        { "ctrlmode", required_argument, 0, 'c' },
        { 0, 0, 0, 0 },
    };

    while ((opt = getopt_long(argc, argv, "hvb:B:c:",
                             long_options, NULL)) != -1) {
        switch (opt) {
            case 'h':
                print_usage(argv[0]);
                exit(0);

            case 'v':
                verbose = 1;
                break;

            case 'b':
                new_baudrate = string_to_baudrate(optarg);
                if (new_baudrate == -1) {
                    print_usage(argv[0]);
                    exit(0);
                }
                break;

            case 'B':
                ptr = optarg;
                while (1) {
                    bittime_data[bittime_count++] = strtoul(ptr, NULL, 0);
                    if (!(ptr = strchr(ptr, ',')))
                        break;
                    ptr++;
                }
                if (bittime_count != 2 && bittime_count != 6) {
                    print_usage(argv[0]);
                    exit(0);
                }
                break;

            case 'c':
                ret = string_to_ctrlmode(optarg);
                if (ret == -1) {
                    print_usage(argv[0]);
                    exit(0);
                }
                new_ctrlmode |= ret;
                set_ctrlmode = 1;
        }
    }
}

```

```

        break;

        break;

    default:
        fprintf(stderr, "Unknown option %c\n", opt);
        break;
    }
}

/* Get CAN interface name */
if (optind != argc - 1 && optind != argc - 2) {
    print_usage(argv[0]);
    return 0;
}

strncpy(iframe, argv[optind], IFNAMSIZ);
strncpy(ifr.ifr_name, iframe, IFNAMSIZ);

if (optind == argc - 2) { /* Get mode setting */
    new_mode = string_to_mode(argv[optind + 1]);
    if (verbose)
        printf("mode: %s (%#x)\n", argv[optind + 1], new_mode);
    if (new_mode < 0) {
        print_usage(argv[0]);
        return 0;
    }
}

can_fd = rt_dev_socket(PF_CAN, SOCK_RAW, CAN_RAW);
if (can_fd < 0) {
    fprintf(stderr, "Cannot open RTDM CAN socket. Maybe driver not loaded? \n");
    return can_fd;
}

ret = rt_dev_ioctl(can_fd, SIOCGIFINDEX, &ifr);
if (ret) {
    fprintf(stderr, "Can't get interface index for %s, code = %d\n", iframe, ret);
    return ret;
}

if (new_baudrate != -1) {
    if (verbose)
        printf("baudrate: %d\n", new_baudrate);
    baudrate = (can_baudrate_t *)&ifr.ifr_ifru;
    *baudrate = new_baudrate;
    ret = rt_dev_ioctl(can_fd, SIOCSANBAUDRATE, &ifr);
    if (ret) {
        goto abort;
    }
}

if (bittime_count) {
    bittime = (struct can_bittime *)&ifr.ifr_ifru;
    if (bittime_count == 2) {
        bittime->type = CAN_BITTIME_BTR;
        bittime->btr.btr0 = bittime_data[0];
        bittime->btr.btr1 = bittime_data[1];
        if (verbose)
            printf("bit-time: btr0=0x%02x btr1=0x%02x\n",
                bittime->btr.btr0, bittime->btr.btr1);
    } else {
        bittime->type = CAN_BITTIME_STD;
        bittime->std.brp = bittime_data[0];
        bittime->std.prop_seg = bittime_data[1];
        bittime->std.phase_seg1 = bittime_data[2];
        bittime->std.phase_seg2 = bittime_data[3];
        bittime->std.sjw = bittime_data[4];
        bittime->std.sam = bittime_data[5];
        if (verbose)
            printf("bit-time: brp=%d prop_seg=%d phase_seg1=%d "
                "phase_seg2=%d sjw=%d sam=%d\n",
                bittime->std.brp,
                bittime->std.prop_seg,
                bittime->std.phase_seg1,
                bittime->std.phase_seg2,
                bittime->std.sjw,
                bittime->std.sam);
    }

    ret = rt_dev_ioctl(can_fd, SIOCSANCUSTOMBITTIME, &ifr);
    if (ret) {
        goto abort;
    }
}
}

```



```

    if (set_ctrlmode != 0) {
        ctrlmode = (can_ctrlmode_t *)&ifr.ifr_ifru;
        *ctrlmode = new_ctrlmode;
        if (verbose)
            printf("ctrlmode: %#x\n", new_ctrlmode);
        ret = rt_dev_ioctl(can_fd, SIOCSANCTRLMODE, &ifr);
        if (ret) {
            goto abort;
        }
    }

    if (new_mode != -1) {
        mode = (can_mode_t *)&ifr.ifr_ifru;
        *mode = new_mode;
        ret = rt_dev_ioctl(can_fd, SIOCSANMODE, &ifr);
        if (ret) {
            goto abort;
        }
    }

    rt_dev_close(can_fd);
    return 0;

abort:
    rt_dev_close(can_fd);
    return ret;
}

```

## 8.8 rtcanrecvc

```

/*
 * Program to receive CAN messages
 *
 * Copyright (C) 2006 Wolfgang Grandegger <wg@grandegger.com>
 *
 * This program is free software; you can redistribute it and/or modify
 * it under the terms of the GNU General Public License as published by
 * the Free Software Foundation; either version 2 of the License, or
 * (at your option) any later version.
 *
 * This program is distributed in the hope that it will be useful,
 * but WITHOUT ANY WARRANTY; without even the implied warranty of
 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
 * GNU General Public License for more details.
 *
 * You should have received a copy of the GNU General Public License
 * along with this program; if not, write to the Free Software
 * Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
 */
#include <stdio.h>
#include <stdlib.h>
#include <signal.h>
#include <unistd.h>
#include <time.h>
#include <errno.h>
#include <getopt.h>
#include <sys/mman.h>

#include <native/task.h>
#include <native/pipe.h>

#include <rtdm/rtdm.h>

static void print_usage(char *prg)
{
    fprintf(stderr,
        "Usage: %s [<can-interface>] [Options]\n"
        "Options:\n"
        " -f --filter=id:mask[:id:mask]... apply filter\n"
        " -e --error=mask      receive error messages\n"
        " -t, --timeout=MS     timeout in ms\n"
        " -T, --timestamp      with absolute timestamp\n"
        " -R, --timestamp-rel  with relative timestamp\n"
        " -v, --verbose        be verbose\n"
        " -p, --print=MODULO   print every MODULO message\n"
        " -h, --help          this help\n",
        prg);
}

extern int optind, opterr, optopt;

```

```

static int s = -1, verbose = 0, print = 1;
static nanosecs_rel_t timeout = 0, with_timestamp = 0, timestamp_rel = 0;

RT_TASK rt_task_desc;

#define BUF_SIZ 255
#define MAX_FILTER 16

struct sockaddr_can recv_addr;
struct can_filter recv_filter[MAX_FILTER];
static int filter_count = 0;

int add_filter(u_int32_t id, u_int32_t mask)
{
    if (filter_count >= MAX_FILTER)
        return -1;
    recv_filter[filter_count].can_id = id;
    recv_filter[filter_count].can_mask = mask;
    printf("Filter #%d: id=0x%08x mask=0x%08x\n", filter_count, id, mask);
    filter_count++;
    return 0;
}

void cleanup(void)
{
    int ret;

    if (verbose)
        printf("Cleaning up...\n");

    if (s >= 0) {
        ret = rt_dev_close(s);
        s = -1;
        if (ret) {
            fprintf(stderr, "rt_dev_close: %s\n", strerror(-ret));
        }
        exit(EXIT_SUCCESS);
    }
}

void cleanup_and_exit(int sig)
{
    if (verbose)
        printf("Signal %d received\n", sig);
    cleanup();
    exit(0);
}

void rt_task(void)
{
    int i, ret, count = 0;
    struct can_frame frame;
    struct sockaddr_can addr;
    socklen_t addrlen = sizeof(addr);
    struct msghdr msg;
    struct iovec iov;
    nanosecs_abs_t timestamp, timestamp_prev = 0;

    if (with_timestamp) {
        msg.msg_iov = &iov;
        msg.msg_iovlen = 1;
        msg.msg_name = (void *)&addr;
        msg.msg_namelen = sizeof(struct sockaddr_can);
        msg.msg_control = (void *)&timestamp;
        msg.msg_controllen = sizeof(nanosecs_abs_t);
    }

    while (1) {
        if (with_timestamp) {
            iov.iov_base = (void *)&frame;
            iov.iov_len = sizeof(can_frame_t);
            ret = rt_dev_recvmmsg(s, &msg, 0);
        } else
            ret = rt_dev_recvfrom(s, (void *)&frame, sizeof(can_frame_t), 0,
                                (struct sockaddr *)&addr, &addrlen);

        if (ret < 0) {
            switch (ret) {
                case -ETIMEDOUT:
                    if (verbose)
                        printf("rt_dev_recv: timed out");
                    continue;
                case -EBADF:
                    if (verbose)
                        printf("rt_dev_recv: aborted because socket was closed");
                    break;
                default:
                    fprintf(stderr, "rt_dev_recv: %s\n", strerror(-ret));
            }
        }
    }
}

```

```

    }
    break;
}

if (print && (count % print) == 0) {
    printf("#%d: (%d) ", count, addr.can_ifindex);
    if (with_timestamp && msg.msg_controlen) {
        if (timestamp_rel) {
            printf("%lldns ", (long long)(timestamp - timestamp_prev));
            timestamp_prev = timestamp;
        } else
            printf("%lldns ", (long long)timestamp);
    }
    if (frame.can_id & CAN_ERR_FLAG)
        printf("!0x%08x!", frame.can_id & CAN_ERR_MASK);
    else if (frame.can_id & CAN_EFF_FLAG)
        printf("<0x%08x>", frame.can_id & CAN_EFF_MASK);
    else
        printf("<0x%03x>", frame.can_id & CAN_SFF_MASK);

    printf(" [%d]", frame.can_dlc);
    if (!(frame.can_id & CAN_RTR_FLAG))
        for (i = 0; i < frame.can_dlc; i++) {
            printf(" %02x", frame.data[i]);
        }
    if (frame.can_id & CAN_ERR_FLAG) {
        printf(" ERROR ");
        if (frame.can_id & CAN_ERR_BUSOFF)
            printf("bus-off");
        if (frame.can_id & CAN_ERR_CRTL)
            printf("controller problem");
    } else if (frame.can_id & CAN_RTR_FLAG)
        printf(" remote request");
    printf("\n");
}
count++;
}
}

int main(int argc, char **argv)
{
    int opt, ret;
    u_int32_t id, mask;
    u_int32_t err_mask = 0;
    struct ifreq ifr;
    char *ptr;
    char name[32];

    struct option long_options[] = {
        { "help", no_argument, 0, 'h' },
        { "verbose", no_argument, 0, 'v' },
        { "filter", required_argument, 0, 'f' },
        { "error", required_argument, 0, 'e' },
        { "timeout", required_argument, 0, 't' },
        { "timestamp", no_argument, 0, 'T' },
        { "timestamp-rel", no_argument, 0, 'R' },
        { 0, 0, 0, 0 },
    };

    mlockall(MCL_CURRENT | MCL_FUTURE);

    signal(SIGTERM, cleanup_and_exit);
    signal(SIGINT, cleanup_and_exit);

    while ((opt = getopt_long(argc, argv, "hve:f:t:p:RT",
                             long_options, NULL)) != -1) {
        switch (opt) {
            case 'h':
                print_usage(argv[0]);
                exit(0);

            case 'p':
                print = strtoul(optarg, NULL, 0);
                break;

            case 'v':
                verbose = 1;
                break;

            case 'e':
                err_mask = strtoul(optarg, NULL, 0);
                break;

            case 'f':
                ptr = optarg;
                while (1) {
                    id = strtoul(ptr, NULL, 0);

```

```

        ptr = strchr(ptr, ':');
        if (!ptr) {
            fprintf(stderr, "filter must be applied in the form id:mask[:id:mask]...\n");
            exit(1);
        }
        ptr++;
        mask = strtoul(ptr, NULL, 0);
        ptr = strchr(ptr, ':');
        add_filter(id, mask);
        if (!ptr)
            break;
        ptr++;
    }
    break;

case 't':
    timeout = (nanosecs_rel_t)strtoul(optarg, NULL, 0) * 1000000;
    break;

case 'R':
    timestamp_rel = 1;
case 'T':
    with_timestamp = 1;
    break;

default:
    fprintf(stderr, "Unknown option %c\n", opt);
    break;
}
}

ret = rt_dev_socket(PF_CAN, SOCK_RAW, CAN_RAW);
if (ret < 0) {
    fprintf(stderr, "rt_dev_socket: %s\n", strerror(-ret));
    return -1;
}
s = ret;

if (argv[optind] == NULL) {
    if (verbose)
        printf("interface all\n");

    ifr.ifr_ifindex = 0;
} else {
    if (verbose)
        printf("interface %s\n", argv[optind]);

    strncpy(ifr.ifr_name, argv[optind], IFNAMSIZ);
    if (verbose)
        printf("s=%d, ifr_name=%s\n", s, ifr.ifr_name);

    ret = rt_dev_ioctl(s, SIOCGIFINDEX, &ifr);
    if (ret < 0) {
        fprintf(stderr, "rt_dev_ioctl GET_IFINDEX: %s\n", strerror(-ret));
        goto failure;
    }
}

if (err_mask) {
    ret = rt_dev_setsockopt(s, SOL_CAN_RAW, CAN_RAW_ERR_FILTER,
                           &err_mask, sizeof(err_mask));

    if (ret < 0) {
        fprintf(stderr, "rt_dev_setsockopt: %s\n", strerror(-ret));
        goto failure;
    }
    if (verbose)
        printf("Using err_mask=%#x\n", err_mask);
}

if (filter_count) {
    ret = rt_dev_setsockopt(s, SOL_CAN_RAW, CAN_RAW_FILTER,
                           &recv_filter, filter_count *
                           sizeof(struct can_filter));

    if (ret < 0) {
        fprintf(stderr, "rt_dev_setsockopt: %s\n", strerror(-ret));
        goto failure;
    }
}

recv_addr.can_family = AF_CAN;
recv_addr.can_ifindex = ifr.ifr_ifindex;
ret = rt_dev_bind(s, (struct sockaddr *)&recv_addr,
                 sizeof(struct sockaddr_can));
if (ret < 0) {
    fprintf(stderr, "rt_dev_bind: %s\n", strerror(-ret));
    goto failure;
}
}

```

```

    if (timeout) {
        if (verbose)
            printf("Timeout: %lld ns\n", (long long)timeout);
        ret = rt_dev_ioctl(s, RTCAN_RTIOC_RCV_TIMEOUT, &timeout);
        if (ret) {
            fprintf(stderr, "rt_dev_ioctl RCV_TIMEOUT: %s\n", strerror(-ret));
            goto failure;
        }
    }

    if (with_timestamp) {
        ret = rt_dev_ioctl(s, RTCAN_RTIOC_TAKE_TIMESTAMP, RTCAN_TAKE_TIMESTAMPS);
        if (ret) {
            fprintf(stderr, "rt_dev_ioctl TAKE_TIMESTAMP: %s\n", strerror(-ret));
            goto failure;
        }
    }

    snprintf(name, sizeof(name), "rtcanrcv-%d", getpid());
    ret = rt_task_shadow(&rt_task_desc, name, 0, 0);
    if (ret) {
        fprintf(stderr, "rt_task_shadow: %s\n", strerror(-ret));
        goto failure;
    }

    rt_task();
    /* never returns */

failure:
    cleanup();
    return -1;
}

```

## 8.9 rtcansend.c

```

/*
 * Program to send CAN messages
 *
 * Copyright (C) 2006 Wolfgang Grandegger <wg@grandegger.com>
 *
 * This program is free software; you can redistribute it and/or modify
 * it under the terms of the GNU General Public License as published by
 * the Free Software Foundation; either version 2 of the License, or
 * (at your option) any later version.
 *
 * This program is distributed in the hope that it will be useful,
 * but WITHOUT ANY WARRANTY; without even the implied warranty of
 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
 * GNU General Public License for more details.
 *
 * You should have received a copy of the GNU General Public License
 * along with this program; if not, write to the Free Software
 * Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
 */
#include <stdio.h>
#include <stdlib.h>
#include <signal.h>
#include <unistd.h>
#include <time.h>
#include <errno.h>
#include <getopt.h>
#include <sys/mman.h>

#include <native/task.h>
#include <native/timer.h>
#include <native/pipe.h>

#include <rtcm/rtcan.h>

extern int optind, opterr, optopt;

static void print_usage(char *prg)
{
    fprintf(stderr,
            "Usage: %s <can-interface> [Options] <can-msg>\n"
            "<can-msg> can consist of up to 8 bytes given as a space separated list\n"
            "Options:\n"
            " -i, --identifier=ID   CAN Identifier (default = 1)\n"
            " -r --rtr              send remote request\n"
            " -e --extended         send extended frame\n"
            " -l --loop=COUNT     send message COUNT times\n"
            " -c, --count           message count in data[0-3]\n"
    );
}

```

```

    " -d, --delay=MS      delay in ms (default = 1ms)\n"
    " -s, --send         use send instead of sendto\n"
    " -t, --timeout=MS   timeout in ms\n"
    " -L, --loopback=0|1 switch local loopback off or on\n"
    " -v, --verbose       be verbose\n"
    " -p, --print=MODULO print every MODULO message\n"
    " -h, --help         this help\n",
    prg);
}

RT_TASK rt_task_desc;

static int s=-1, dlc=0, rtr=0, extended=0, verbose=0, loops=1;
static SRTIME delay=1000000;
static int count=0, print=1, use_send=0, loopback=-1;
static nanosecs_rel_t timeout = 0;
static struct can_frame frame;
static struct sockaddr_can to_addr;

void cleanup(void)
{
    int ret;

    if (verbose)
        printf("Cleaning up...\n");

    usleep(100000);

    if (s >= 0) {
        ret = rt_dev_close(s);
        s = -1;
        if (ret) {
            fprintf(stderr, "rt_dev_close: %s\n", strerror(-ret));
        }
        exit(EXIT_SUCCESS);
    }
}

void cleanup_and_exit(int sig)
{
    if (verbose)
        printf("Signal %d received\n", sig);
    cleanup();
    exit(0);
}

void rt_task(void)
{
    int i, j, ret;

    for (i = 0; i < loops; i++) {
        rt_task_sleep(rt_timer_ns2ticks(delay));
        if (count)
            memcpy(&frame.data[0], &i, sizeof(i));
        /* Note: sendto avoids the definition of a receive filter list */
        if (use_send)
            ret = rt_dev_send(s, (void *)&frame, sizeof(can_frame_t), 0);
        else
            ret = rt_dev_sendto(s, (void *)&frame, sizeof(can_frame_t), 0,
                               (struct sockaddr *)&to_addr, sizeof(to_addr));
        if (ret < 0) {
            switch (ret) {
                case -ETIMEDOUT:
                    if (verbose)
                        printf("rt_dev_send(to): timed out");
                    break;
                case -EBADF:
                    if (verbose)
                        printf("rt_dev_send(to): aborted because socket was closed");
                    break;
                default:
                    fprintf(stderr, "rt_dev_send: %s\n", strerror(-ret));
                    break;
            }
            i = loops;          /* abort */
            break;
        }
        if (verbose && (i % print) == 0) {
            if (frame.can_id & CAN_EFF_FLAG)
                printf("<0x%08x>", frame.can_id & CAN_EFF_MASK);
            else
                printf("<0x%03x>", frame.can_id & CAN_SFF_MASK);
            printf(" [%d]", frame.can_dlc);
            for (j = 0; j < frame.can_dlc; j++) {
                printf(" %02x", frame.data[j]);
            }
        }
    }
}

```

```

        }
        printf("\n");
    }
}

int main(int argc, char **argv)
{
    int i, opt, ret;
    struct ifreq ifr;
    char name[32];

    struct option long_options[] = {
        { "help", no_argument, 0, 'h' },
        { "identifier", required_argument, 0, 'i' },
        { "rtr", no_argument, 0, 'r' },
        { "extended", no_argument, 0, 'e' },
        { "verbose", no_argument, 0, 'v' },
        { "count", no_argument, 0, 'c' },
        { "print", required_argument, 0, 'p' },
        { "loop", required_argument, 0, 'l' },
        { "delay", required_argument, 0, 'd' },
        { "send", no_argument, 0, 's' },
        { "timeout", required_argument, 0, 't' },
        { "loopback", required_argument, 0, 'L' },
        { 0, 0, 0, 0 },
    };

    mlockall(MCL_CURRENT | MCL_FUTURE);

    signal(SIGTERM, cleanup_and_exit);
    signal(SIGINT, cleanup_and_exit);

    frame.can_id = 1;

    while ((opt = getopt_long(argc, argv, "hvi:l:red:t:cp:sL:",
                             long_options, NULL)) != -1) {
        switch (opt) {
            case 'h':
                print_usage(argv[0]);
                exit(0);

            case 'p':
                print = strtoul(optarg, NULL, 0);

            case 'v':
                verbose = 1;
                break;

            case 'c':
                count = 1;
                break;

            case 'l':
                loops = strtoul(optarg, NULL, 0);
                break;

            case 'i':
                frame.can_id = strtoul(optarg, NULL, 0);
                break;

            case 'r':
                rtr = 1;
                break;

            case 'e':
                extended = 1;
                break;

            case 'd':
                delay = strtoul(optarg, NULL, 0) * 1000000LL;
                break;

            case 's':
                use_send = 1;
                break;

            case 't':
                timeout = strtoul(optarg, NULL, 0) * 1000000LL;
                break;

            case 'L':
                loopback = strtoul(optarg, NULL, 0);
                break;

            default:
                fprintf(stderr, "Unknown option %c\n", opt);
        }
    }
}

```

```

        break;
    }
}

if (optind == argc) {
    print_usage(argv[0]);
    exit(0);
}

if (argv[optind] == NULL) {
    fprintf(stderr, "No Interface supplied\n");
    exit(-1);
}

if (verbose)
    printf("interface %s\n", argv[optind]);

ret = rt_dev_socket(PF_CAN, SOCK_RAW, CAN_RAW);
if (ret < 0) {
    fprintf(stderr, "rt_dev_socket: %s\n", strerror(-ret));
    return -1;
}
s = ret;

if (loopback >= 0) {
    ret = rt_dev_setsockopt(s, SOL_CAN_RAW, CAN_RAW_LOOPBACK,
                           &loopback, sizeof(loopback));

    if (ret < 0) {
        fprintf(stderr, "rt_dev_setsockopt: %s\n", strerror(-ret));
        goto failure;
    }
    if (verbose)
        printf("Using loopback=%d\n", loopback);
}

strncpy(ifr.ifr_name, argv[optind], IFNAMSIZ);
if (verbose)
    printf("s=%d, ifr_name=%s\n", s, ifr.ifr_name);

ret = rt_dev_ioctl(s, SIOCGIFINDEX, &ifr);
if (ret < 0) {
    fprintf(stderr, "rt_dev_ioctl: %s\n", strerror(-ret));
    goto failure;
}

memset(&to_addr, 0, sizeof(to_addr));
to_addr.can_ifindex = ifr.ifr_ifindex;
to_addr.can_family = AF_CAN;
if (use_send) {
    /* Suppress definiton of a default receive filter list */
    ret = rt_dev_setsockopt(s, SOL_CAN_RAW, CAN_RAW_FILTER, NULL, 0);
    if (ret < 0) {
        fprintf(stderr, "rt_dev_setsockopt: %s\n", strerror(-ret));
        goto failure;
    }

    ret = rt_dev_bind(s, (struct sockaddr *)&to_addr, sizeof(to_addr));
    if (ret < 0) {
        fprintf(stderr, "rt_dev_bind: %s\n", strerror(-ret));
        goto failure;
    }
}

if (count)
    frame.can_dlc = sizeof(int);
else {
    for (i = optind + 1; i < argc; i++) {
        frame.data[dlc] = strtoul(argv[i], NULL, 0);
        dlc++;
        if (dlc == 8)
            break;
    }
    frame.can_dlc = dlc;
}

if (rtr)
    frame.can_id |= CAN_RTR_FLAG;

if (extended)
    frame.can_id |= CAN_EFF_FLAG;

if (timeout) {
    if (verbose)
        printf("Timeout: %lld ns\n", (long long)timeout);
    ret = rt_dev_ioctl(s, RTCAN_RTIOC_SND_TIMEOUT, &timeout);
    if (ret) {
        fprintf(stderr, "rt_dev_ioctl SND_TIMEOUT: %s\n", strerror(-ret));
    }
}

```



```

        goto failure;
    }
}

snprintf(name, sizeof(name), "rtcansend-%d", getpid());
ret = rt_task_shadow(&rt_task_desc, name, 1, 0);
if (ret) {
    fprintf(stderr, "rt_task_shadow: %s\n", strerror(-ret));
    goto failure;
}

rt_task();

cleanup();
return 0;

failure:
cleanup();
return -1;
}

```

## 8.10 xddp-echo.c

```

/*
 * XDDP-based RT/NRT threads communication demo.
 *
 * Real-time Xenomai threads and regular Linux threads may want to
 * exchange data in a way that does not require the former to leave
 * the real-time domain (i.e. secondary mode). Message pipes - as
 * implemented by the RTDM-based XDDP protocol - are provided for this
 * purpose.
 *
 * On the Linux domain side, pseudo-device files named /dev/rtp<minor>
 * give regular POSIX threads access to non real-time communication
 * endpoints, via the standard character-based I/O interface. On the
 * Xenomai domain side, sockets may be bound to XDDP ports, which act
 * as proxies to send and receive data to/from the associated
 * pseudo-device files. Ports and pseudo-device minor numbers are
 * paired, meaning that e.g. port 7 will proxy the traffic for
 * /dev/rtp7. Therefore, port numbers may range from 0 to
 * CONFIG_XENO_OPT_PIPE_NRDEV - 1.
 *
 * All data sent through a bound/connected XDDP socket via sendto(2) or
 * write(2) will be passed to the peer endpoint in the Linux domain,
 * and made available for reading via the standard read(2) system
 * call. Conversely, all data sent using write(2) through the non
 * real-time endpoint will be conveyed to the real-time socket
 * endpoint, and made available to the recvfrom(2) or read(2) system
 * calls.
 *
 * Both threads can use the bi-directional data path to send and
 * receive datagrams in a FIFO manner, as illustrated by the simple
 * echoing process implemented by this program.
 *
 * realtime_thread----->-----+
 * => get socket |
 * => bind socket to port 0 | v
 * => write traffic to NRT domain via sendto() |
 * => read traffic from NRT domain via recvfrom() <--|--+
 * | |
 * regular_thread-----+ |
 * => open /dev/rtp0 | | ^
 * => read traffic from RT domain via read() | |
 * => echo traffic back to RT domain via write() +--+
 *
 * See Makefile in this directory for build directives.
 *
 * NOTE: XDDP is a replacement for the legacy RT_PIPE interface
 * available from the native skin until Xenomai 3.
 */
#include <sys/mman.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <malloc.h>
#include <pthread.h>
#include <fcntl.h>
#include <errno.h>
#include <rtdk.h>
#include <rtdm/rtipc.h>

```

```

pthread_t rt, nrt;

#define XDDP_PORT 0 /* [0..CONFIG-XENO_OPT_PIPE_NRDEV - 1] */

static const char *msg[] = {
    "Surfing With The Alien",
    "Lords of Karma",
    "Banana Mango",
    "Psycho Monkey",
    "Luminous Flesh Giants",
    "Moroccan Sunset",
    "Satch Boogie",
    "Flying In A Blue Dream",
    "Ride",
    "Summer Song",
    "Speed Of Light",
    "Crystal Planet",
    "Raspberry Jam Delta-V",
    "Champagne?",
    "Clouds Race Across The Sky",
    "Engines Of Creation"
};

static void fail(const char *reason)
{
    perror(reason);
    exit(EXIT_FAILURE);
}

static void *realtime_thread(void *arg)
{
    struct sockaddr_ipc saddr;
    int ret, s, n = 0, len;
    struct timespec ts;
    size_t poolsz;
    char buf[128];

    /*
     * Get a datagram socket to bind to the RT endpoint. Each
     * endpoint is represented by a port number within the XDDP
     * protocol namespace.
     */
    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_XDDP);
    if (s < 0) {
        perror("socket");
        exit(EXIT_FAILURE);
    }

    /*
     * Set a local 16k pool for the RT endpoint. Memory needed to
     * convey datagrams will be pulled from this pool, instead of
     * Xenomai's system pool.
     */
    poolsz = 16384; /* bytes */
    ret = setsockopt(s, SOL_XDDP, XDDP_POOLSZ,
                    &poolsz, sizeof(poolsz));
    if (ret)
        fail("setsockopt");

    /*
     * Bind the socket to the port, to setup a proxy to channel
     * traffic to/from the Linux domain.
     *
     * saddr.sipc_port specifies the port number to use.
     */
    memset(&saddr, 0, sizeof(saddr));
    saddr.sipc_family = AF_RTIPC;
    saddr.sipc_port = XDDP_PORT;
    ret = bind(s, (struct sockaddr *)&saddr, sizeof(saddr));
    if (ret)
        fail("bind");

    for (;;) {
        len = strlen(msg[n]);
        /*
         * Send a datagram to the NRT endpoint via the proxy.
         * We may pass a NULL destination address, since a
         * bound socket is assigned a default destination
         * address matching the binding address (unless
         * connect(2) was issued before bind(2), in which case
         * the former would prevail).
         */
        ret = sendto(s, msg[n], len, 0, NULL, 0);
        if (ret != len)
            fail("sendto");

        rt_printf("%s: sent %d bytes, \".*s\"\\n",

```

```

        __FUNCTION__, ret, ret, msg[n]);

    /* Read back packets echoed by the regular thread */
    ret = recvfrom(s, buf, sizeof(buf), 0, NULL, 0);
    if (ret <= 0)
        fail("recvfrom");

    rt_printf(" => \"%s\" echoed by peer\n", ret, buf);

    n = (n + 1) % (sizeof(msg) / sizeof(msg[0]));
    /*
     * We run in full real-time mode (i.e. primary mode),
     * so we have to let the system breathe between two
     * iterations.
     */
    ts.tv_sec = 0;
    ts.tv_nsec = 500000000; /* 500 ms */
    clock_nanosleep(CLOCK_REALTIME, 0, &ts, NULL);
}

return NULL;
}

static void *regular_thread(void *arg)
{
    char buf[128], *devname;
    int fd, ret;

    if (asprintf(&devname, "/dev/rtp%d", XDDP_PORT) < 0)
        fail("asprintf");

    fd = open(devname, O_RDWR);
    free(devname);
    if (fd < 0)
        fail("open");

    for (;;) {
        /* Get the next message from realtime_thread. */
        ret = read(fd, buf, sizeof(buf));
        if (ret <= 0)
            fail("read");

        /* Echo the message back to realtime_thread. */
        ret = write(fd, buf, ret);
        if (ret <= 0)
            fail("write");
    }

    return NULL;
}

static void cleanup_upon_sig(int sig)
{
    pthread_cancel(rt);
    pthread_cancel(nrt);
    signal(sig, SIG_DFL);
    pthread_join(rt, NULL);
    pthread_join(nrt, NULL);
}

int main(int argc, char **argv)
{
    struct sched_param rtparam = { .sched_priority = 42 };
    pthread_attr_t rtattr, regattr;
    sigset_t mask, oldmask;

    mlockall(MCL_CURRENT | MCL_FUTURE);

    sigemptyset(&mask);
    sigaddset(&mask, SIGINT);
    signal(SIGINT, cleanup_upon_sig);
    sigaddset(&mask, SIGTERM);
    signal(SIGTERM, cleanup_upon_sig);
    sigaddset(&mask, SIGHUP);
    signal(SIGHUP, cleanup_upon_sig);
    pthread_sigmask(SIG_BLOCK, &mask, &oldmask);

    /*
     * This is a real-time compatible printf() package from
     * Xenomai's RT Development Kit (RTDK), that does NOT cause
     * any transition to secondary (i.e. non real-time) mode when
     * writing output.
     */
    rt_print_auto_init(1);

    pthread_attr_init(&rtattr);
    pthread_attr_setdetachstate(&rtattr, PTHREAD_CREATE_JOINABLE);
}

```

```

pthread_attr_setinheritsched(&rtattr, PTHREAD_EXPLICIT_SCHED);
pthread_attr_setschedpolicy(&rtattr, SCHED_FIFO);
pthread_attr_setschedparam(&rtattr, &rtparam);

errno = pthread_create(&rt, &rtattr, &realtime_thread, NULL);
if (errno)
    fail("pthread_create");

pthread_attr_init(&regattr);
pthread_attr_setdetachstate(&regattr, PTHREAD_CREATE_JOINABLE);
pthread_attr_setinheritsched(&regattr, PTHREAD_EXPLICIT_SCHED);
pthread_attr_setschedpolicy(&regattr, SCHED_OTHER);

errno = pthread_create(&nrt, &regattr, &regular_thread, NULL);
if (errno)
    fail("pthread_create");

sigsuspend(&oldmask);

return 0;
}

```

## 8.11 xddp-label.c

```

/*
 * XDDP-based RT/NRT threads communication demo.
 *
 * Real-time Xenomai threads and regular Linux threads may want to
 * exchange data in a way that does not require the former to leave
 * the real-time domain (i.e. secondary mode). Message pipes - as
 * implemented by the RTDM-based XDDP protocol - are provided for this
 * purpose.
 *
 * On the Linux domain side, pseudo-device files named /dev/rtp<minor>
 * give regular POSIX threads access to non real-time communication
 * endpoints, via the standard character-based I/O interface. On the
 * Xenomai domain side, sockets may be bound to XDDP ports, which act
 * as proxies to send and receive data to/from the associated
 * pseudo-device files. Ports and pseudo-device minor numbers are
 * paired, meaning that e.g. port 7 will proxy the traffic for
 * /dev/rtp7. Therefore, port numbers may range from 0 to
 * CONFIG_XENO_OPT_PIPE_NRDEV - 1.
 *
 * All data sent through a bound/connected XDDP socket via sendto(2) or
 * write(2) will be passed to the peer endpoint in the Linux domain,
 * and made available for reading via the standard read(2) system
 * call. Conversely, all data sent using write(2) through the non
 * real-time endpoint will be conveyed to the real-time socket
 * endpoint, and made available to the recvfrom(2) or read(2) system
 * calls.
 *
 * ASCII labels can be attached to bound ports, in order to connect
 * sockets to them in a more descriptive way than using plain numeric
 * port values.
 *
 * The example code below illustrates the following process:
 *
 * realtime_thread1----->-----+
 * => get socket                      |
 * => bind socket to port "xddp-demo  |
 * => read traffic from NRT domain via recvfrom() <---+
 *
 * realtime_thread2-----+-----|
 * => get socket              | |
 * => connect socket to port "xddp-demo" | |
 * => write traffic to NRT domain via sendto() v |
 *
 * regular_thread-----+-----|
 * => open /proc/xenomai/registry/rtipc/xddp/xddp-demo | |
 * => read traffic from RT domain via read() | |
 * => mirror traffic to RT domain via write() +---+
 *
 * See Makefile in this directory for build directives.
 *
 * NOTE: XDDP is a replacement for the legacy RT_PIPE interface
 * available from the native skin until Xenomai 3.
 */
#include <sys/mman.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>

```

```

#include <malloc.h>
#include <pthread.h>
#include <fcntl.h>
#include <errno.h>
#include <rtdk.h>
#include <rtdm/rtipc.h>

pthread_t rt1, rt2, nrt;

#define XDDP_PORT_LABEL "xddp-demo"

static const char *msg[] = {
    "Surfing With The Alien",
    "Lords of Karma",
    "Banana Mango",
    "Psycho Monkey",
    "Luminous Flesh Giants",
    "Moroccan Sunset",
    "Satch Boogie",
    "Flying In A Blue Dream",
    "Ride",
    "Summer Song",
    "Speed Of Light",
    "Crystal Planet",
    "Raspberry Jam Delta-V",
    "Champagne?",
    "Clouds Race Across The Sky",
    "Engines Of Creation"
};

static void fail(const char *reason)
{
    perror(reason);
    exit(EXIT_FAILURE);
}

static void *realtime_thread1(void *arg)
{
    struct rtipc_port_label plabel;
    struct sockaddr_ipc saddr;
    char buf[128];
    int ret, s;

    /*
     * Get a datagram socket to bind to the RT endpoint. Each
     * endpoint is represented by a port number within the XDDP
     * protocol namespace.
     */
    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_XDDP);
    if (s < 0) {
        perror("socket");
        exit(EXIT_FAILURE);
    }

    /*
     * Set a port label. This name will be registered when
     * binding, in addition to the port number (if given).
     */
    strcpy(plabel.label, XDDP_PORT_LABEL);
    ret = setsockopt(s, SOL_XDDP, XDDP_LABEL,
                    &plabel, sizeof(plabel));
    if (ret)
        fail("setsockopt");

    /*
     * Bind the socket to the port, to setup a proxy to channel
     * traffic to/from the Linux domain. Assign that port a label,
     * so that peers may use a descriptive information to locate
     * it. For instance, the pseudo-device matching our RT
     * endpoint will appear as
     * /proc/xenomai/registry/rtipc/xddp/<XDDP_PORT_LABEL> in the
     * Linux domain, once the socket is bound.
     */
    /* saddr.sipc_port specifies the port number to use. If -1 is
     * passed, the XDDP driver will auto-select an idle port.
     */
    memset(&saddr, 0, sizeof(saddr));
    saddr.sipc_family = AF_RTIPC;
    saddr.sipc_port = -1;
    ret = bind(s, (struct sockaddr *)&saddr, sizeof(saddr));
    if (ret)
        fail("bind");

    for (;;) {
        /* Get packets relayed by the regular thread */
        ret = recvfrom(s, buf, sizeof(buf), 0, NULL, 0);
        if (ret <= 0)
            fail("recvfrom");
    }
}

```

```

        rt_printf("%s: \"%.*s\" relayed by peer\n", __FUNCTION__, ret, buf);
    }

    return NULL;
}

static void *realtime_thread2(void *arg)
{
    struct rtipc_port_label plabel;
    struct sockaddr_ipc saddr;
    int ret, s, n = 0, len;
    struct timespec ts;
    struct timeval tv;
    socklen_t addrlen;

    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_XDDP);
    if (s < 0) {
        perror("socket");
        exit(EXIT_FAILURE);
    }

    /*
     * Set the socket timeout; it will apply when attempting to
     * connect to a labeled port, and to recvfrom() calls. The
     * following setup tells the XDDP driver to wait for at most
     * one second until a socket is bound to a port using the same
     * label, or return with a timeout error.
     */
    tv.tv_sec = 1;
    tv.tv_usec = 0;
    ret = setsockopt(s, SOL_SOCKET, SO_RCVTIMEO,
                    &tv, sizeof(tv));
    if (ret)
        fail("setsockopt");

    /*
     * Set a port label. This name will be used to find the peer
     * when connecting, instead of the port number.
     */
    strcpy(plabel.label, XDDP_PORT_LABEL);
    ret = setsockopt(s, SOL_XDDP, XDDP_LABEL,
                    &plabel, sizeof(plabel));
    if (ret)
        fail("setsockopt");

    memset(&saddr, 0, sizeof(saddr));
    saddr.sipc_family = AF_RTIPC;
    saddr.sipc_port = -1; /* Tell XDDP to search by label. */
    ret = connect(s, (struct sockaddr *)&saddr, sizeof(saddr));
    if (ret)
        fail("connect");

    /*
     * We succeeded in making the port our default destination
     * address by using its label, but we don't know its actual
     * port number yet. Use getpeername() to retrieve it.
     */
    addrlen = sizeof(saddr);
    ret = getpeername(s, (struct sockaddr *)&saddr, &addrlen);
    if (ret || addrlen != sizeof(saddr))
        fail("getpeername");

    rt_printf("%s: NRT peer is reading from /dev/rtp%d\n",
              __FUNCTION__, saddr.sipc_port);

    for (;;) {
        len = strlen(msg[n]);
        /*
         * Send a datagram to the NRT endpoint via the proxy.
         * We may pass a NULL destination address, since the
         * socket was successfully assigned the proper default
         * address via connect(2).
         */
        ret = sendto(s, msg[n], len, 0, NULL, 0);
        if (ret != len)
            fail("sendto");

        rt_printf("%s: sent %d bytes, \"%.*s\"\n",
                  __FUNCTION__, ret, ret, msg[n]);

        n = (n + 1) % (sizeof(msg) / sizeof(msg[0]));
        /*
         * We run in full real-time mode (i.e. primary mode),
         * so we have to let the system breathe between two
         * iterations.
         */
    }
}

```

```

        ts.tv_sec = 0;
        ts.tv_nsec = 500000000; /* 500 ms */
        clock_nanosleep(CLOCK_REALTIME, 0, &ts, NULL);
    }

    return NULL;
}

static void *regular_thread(void *arg)
{
    char buf[128], *devname;
    int fd, ret;

    if (asprintf(&devname,
                "/proc/xenomai/registry/rtipc/xddp/%s",
                XDDP_PORT_LABEL) < 0)
        fail("asprintf");

    fd = open(devname, O_RDWR);
    free(devname);
    if (fd < 0)
        fail("open");

    for (;;) {
        /* Get the next message from realtime_thread2. */
        ret = read(fd, buf, sizeof(buf));
        if (ret <= 0)
            fail("read");

        /* Relay the message to realtime_thread1. */
        ret = write(fd, buf, ret);
        if (ret <= 0)
            fail("write");
    }

    return NULL;
}

static void cleanup_upon_sig(int sig)
{
    pthread_cancel(rt1);
    pthread_cancel(rt2);
    pthread_cancel(nrt);
    signal(sig, SIG_DFL);
    pthread_join(rt1, NULL);
    pthread_join(rt2, NULL);
    pthread_join(nrt, NULL);
}

int main(int argc, char **argv)
{
    struct sched_param rtparam = { .sched_priority = 42 };
    pthread_attr_t rtattr, regattr;
    sigset_t mask, oldmask;

    mlockall(MCL_CURRENT | MCL_FUTURE);

    sigemptyset(&mask);
    sigaddset(&mask, SIGINT);
    signal(SIGINT, cleanup_upon_sig);
    sigaddset(&mask, SIGTERM);
    signal(SIGTERM, cleanup_upon_sig);
    sigaddset(&mask, SIGHUP);
    signal(SIGHUP, cleanup_upon_sig);
    pthread_sigmask(SIG_BLOCK, &mask, &oldmask);

    /*
     * This is a real-time compatible printf() package from
     * Xenomai's RT Development Kit (RTDK), that does NOT cause
     * any transition to secondary (i.e. non real-time) mode when
     * writing output.
     */
    rt_print_auto_init(1);

    pthread_attr_init(&rtattr);
    pthread_attr_setdetachstate(&rtattr, PTHREAD_CREATE_JOINABLE);
    pthread_attr_setinheritsched(&rtattr, PTHREAD_EXPLICIT_SCHED);
    pthread_attr_setschedpolicy(&rtattr, SCHED_FIFO);
    pthread_attr_setschedparam(&rtattr, &rtparam);

    /* Both real-time threads have the same attribute set. */

    errno = pthread_create(&rt1, &rtattr, &realtime_thread1, NULL);
    if (errno)
        fail("pthread_create");

    errno = pthread_create(&rt2, &rtattr, &realtime_thread2, NULL);

```

```

    if (errno)
        fail("pthread_create");

    pthread_attr_init(&regattr);
    pthread_attr_setdetachstate(&regattr, PTHREAD_CREATE_JOINABLE);
    pthread_attr_setinheritsched(&regattr, PTHREAD_EXPLICIT_SCHED);
    pthread_attr_setschedpolicy(&regattr, SCHED_OTHER);

    errno = pthread_create(&nrt, &regattr, &regular_thread, NULL);
    if (errno)
        fail("pthread_create");

    sigsuspend(&oldmask);

    return 0;
}

```

## 8.12 xddp-stream.c

```

/*
 * XDDP-based RT/NRT threads communication demo.
 *
 * Real-time Xenomai threads and regular Linux threads may want to
 * exchange data in a way that does not require the former to leave
 * the real-time domain (i.e. secondary mode). Message pipes - as
 * implemented by the RTDM-based XDDP protocol - are provided for this
 * purpose.
 *
 * On the Linux domain side, pseudo-device files named /dev/rtp<minor>
 * give regular POSIX threads access to non real-time communication
 * endpoints, via the standard character-based I/O interface. On the
 * Xenomai domain side, sockets may be bound to XDDP ports, which act
 * as proxies to send and receive data to/from the associated
 * pseudo-device files. Ports and pseudo-device minor numbers are
 * paired, meaning that e.g. port 7 will proxy the traffic for
 * /dev/rtp7. Therefore, port numbers may range from 0 to
 * CONFIG_XENO_OPT_PIPE_NRDEV - 1.
 *
 * All data sent through a bound/connected XDDP socket via sendto(2) or
 * write(2) will be passed to the peer endpoint in the Linux domain,
 * and made available for reading via the standard read(2) system
 * call. Conversely, all data sent using write(2) through the non
 * real-time endpoint will be conveyed to the real-time socket
 * endpoint, and made available to the recvfrom(2) or read(2) system
 * calls.
 *
 * In addition to sending datagrams, real-time threads may stream data
 * in a byte-oriented mode through the proxy as well. This increases
 * the bandwidth and reduces the overhead, when a lot of data has to
 * flow down to the Linux domain, if keeping the message boundaries is
 * not required. The example code below illustrates such use.
 *
 * realtime_thread----->-----+
 * => get socket |
 * => bind socket to port 0 v
 * => write scattered traffic to NRT domain via sendto() |
 * => read traffic from NRT domain via recvfrom() <--|--+
 * | |
 * regular_thread-----+ |
 * => open /dev/rtp0 | ^
 * => read traffic from RT domain via read() | |
 * => echo traffic back to RT domain via write() +--+
 *
 * See Makefile in this directory for build directives.
 *
 * NOTE: XDDP is a replacement for the legacy RT_PIPE interface
 * available from the native skin until Xenomai 3.
 */
#include <sys/mman.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <signal.h>
#include <string.h>
#include <malloc.h>
#include <pthread.h>
#include <fcntl.h>
#include <errno.h>
#include <rtdk.h>
#include <rtdm/rtpc.h>

pthread_t rt, nrt;

```



```

#define XDDP_PORT 0      /* [0..CONFIG-XENO_OPT_PIPE_NRDEV - 1] */

static const char *msg[] = {
    "Surfing With The Alien",
    "Lords of Karma",
    "Banana Mango",
    "Psycho Monkey",
    "Luminous Flesh Giants",
    "Moroccan Sunset",
    "Satch Boogie",
    "Flying In A Blue Dream",
    "Ride",
    "Summer Song",
    "Speed Of Light",
    "Crystal Planet",
    "Raspberry Jam Delta-V",
    "Champagne?",
    "Clouds Race Across The Sky",
    "Engines Of Creation"
};

static void fail(const char *reason)
{
    perror(reason);
    exit(EXIT_FAILURE);
}

static void *realtime_thread(void *arg)
{
    struct sockaddr_ipc saddr;
    int ret, s, n = 0, len, b;
    struct timespec ts;
    size_t streamsz;
    char buf[128];

    /*
     * Get a datagram socket to bind to the RT endpoint. Each
     * endpoint is represented by a port number within the XDDP
     * protocol namespace.
     */
    s = socket(AF_RTIPC, SOCK_DGRAM, IPCPROTO_XDDP);
    if (s < 0) {
        perror("socket");
        exit(EXIT_FAILURE);
    }

    /*
     * Tell the XDDP driver that we will use the streaming
     * capabilities on this socket. To this end, we have to
     * specify the size of the streaming buffer, as a count of
     * bytes. The real-time output will be buffered up to that
     * amount, and sent as a single datagram to the NRT endpoint
     * when fully gathered, or when another source port attempts
     * to send data to the same endpoint. Passing a null size
     * would disable streaming.
     */
    streamsz = 1024; /* bytes */
    ret = setsockopt(s, SOL_XDDP, XDDP_BUFSZ,
                    &streamsz, sizeof(streamsz));
    if (ret)
        fail("setsockopt");

    /*
     * Bind the socket to the port, to setup a proxy to channel
     * traffic to/from the Linux domain.
     *
     * saddr.sipc_port specifies the port number to use.
     */
    memset(&saddr, 0, sizeof(saddr));
    saddr.sipc_family = AF_RTIPC;
    saddr.sipc_port = XDDP_PORT;
    ret = bind(s, (struct sockaddr *)&saddr, sizeof(saddr));
    if (ret)
        fail("bind");

    for (;;) {
        len = strlen(msg[n]);
        /*
         * Send a datagram to the NRT endpoint via the proxy.
         * The output is artificially scattered in separate
         * one-byte sendings, to illustrate the use of
         * MSG_MORE.
         */
        for (b = 0; b < len; b++) {
            ret = sendto(s, msg[n] + b, 1, MSG_MORE, NULL, 0);
            if (ret != 1)
                fail("sendto");
        }
    }
}

```

```

        rt_printf("%s: sent (scattered) %d-bytes message, \"%. *s\\\"n",
                __FUNCTION__, len, len, msg[n]);

        /* Read back packets echoed by the regular thread */
        ret = recvfrom(s, buf, sizeof(buf), 0, NULL, 0);
        if (ret <= 0)
            fail("recvfrom");

        rt_printf(" => \"%. *s\\\" echoed by peer\n", ret, buf);

        n = (n + 1) % (sizeof(msg) / sizeof(msg[0]));
        /*
         * We run in full real-time mode (i.e. primary mode),
         * so we have to let the system breathe between two
         * iterations.
         */
        ts.tv_sec = 0;
        ts.tv_nsec = 500000000; /* 500 ms */
        clock_nanosleep(CLOCK_REALTIME, 0, &ts, NULL);
    }

    return NULL;
}

static void *regular_thread(void *arg)
{
    char buf[128], *devname;
    int fd, ret;

    if (asprintf(&devname, "/dev/rtp%d", XDDP_PORT) < 0)
        fail("asprintf");

    fd = open(devname, O_RDWR);
    free(devname);
    if (fd < 0)
        fail("open");

    for (;;) {
        /* Get the next message from realtime_thread. */
        ret = read(fd, buf, sizeof(buf));
        if (ret <= 0)
            fail("read");

        /* Echo the message back to realtime_thread. */
        ret = write(fd, buf, ret);
        if (ret <= 0)
            fail("write");
    }

    return NULL;
}

static void cleanup_upon_sig(int sig)
{
    pthread_cancel(rt);
    pthread_cancel(nrt);
    signal(sig, SIG_DFL);
    pthread_join(rt, NULL);
    pthread_join(nrt, NULL);
}

int main(int argc, char **argv)
{
    struct sched_param rtparam = { .sched_priority = 42 };
    pthread_attr_t rtattr, regattr;
    sigset_t mask, oldmask;

    mlockall(MCL_CURRENT | MCL_FUTURE);

    sigemptyset(&mask);
    sigaddset(&mask, SIGINT);
    signal(SIGINT, cleanup_upon_sig);
    sigaddset(&mask, SIGTERM);
    signal(SIGTERM, cleanup_upon_sig);
    sigaddset(&mask, SIGHUP);
    signal(SIGHUP, cleanup_upon_sig);
    pthread_sigmask(SIG_BLOCK, &mask, &oldmask);

    /*
     * This is a real-time compatible printf() package from
     * Xenomai's RT Development Kit (RTDK), that does NOT cause
     * any transition to secondary (i.e. non real-time) mode when
     * writing output.
     */
    rt_print_auto_init(1);
}

```

```
pthread_attr_init(&rtattr);
pthread_attr_setdetachstate(&rtattr, PTHREAD_CREATE_JOINABLE);
pthread_attr_setinheritsched(&rtattr, PTHREAD_EXPLICIT_SCHED);
pthread_attr_setschedpolicy(&rtattr, SCHED_FIFO);
pthread_attr_setschedparam(&rtattr, &rtparam);

errno = pthread_create(&rt, &rtattr, &realtime_thread, NULL);
if (errno)
    fail("pthread_create");

pthread_attr_init(&regattr);
pthread_attr_setdetachstate(&regattr, PTHREAD_CREATE_JOINABLE);
pthread_attr_setinheritsched(&regattr, PTHREAD_EXPLICIT_SCHED);
pthread_attr_setschedpolicy(&regattr, SCHED_OTHER);

errno = pthread_create(&nrt, &regattr, &regular_thread, NULL);
if (errno)
    fail("pthread_create");

sigsuspend(&oldmask);

return 0;
}
```

# Index

- BUFP\_BUFSZ
  - Real-time IPC protocols, [47](#)
- BUFP\_LABEL
  - Real-time IPC protocols, [48](#)
- bind\_\_AF\_RTIPC
  - Real-time IPC protocols, [54](#)
- CAN Devices, [9](#)
  - CAN\_BITTIME\_BTR, [28](#)
  - CAN\_BITTIME\_STD, [28](#)
  - CAN\_BITTIME\_TYPE, [28](#)
  - CAN\_CTRLMODE\_3\_SAMPLES, [17](#)
  - CAN\_CTRLMODE\_LISTENONLY, [17](#)
  - CAN\_CTRLMODE\_LOOPBACK, [18](#)
  - CAN\_ERR\_LOSTARB\_UNSPEC, [18](#)
  - CAN\_MODE, [28](#)
  - CAN\_MODE\_SLEEP, [28](#)
  - CAN\_MODE\_START, [28](#)
  - CAN\_MODE\_STOP, [28](#)
  - CAN\_RAW\_ERR\_FILTER, [18](#)
  - CAN\_RAW\_FILTER, [18](#)
  - CAN\_RAW\_LOOPBACK, [19](#)
  - CAN\_RAW\_RECV\_OWN\_MSGS, [20](#)
  - CAN\_STATE, [28](#)
  - CAN\_STATE\_BUS\_OFF, [29](#)
  - CAN\_STATE\_ERROR\_ACTIVE, [29](#)
  - CAN\_STATE\_ERROR\_PASSIVE, [29](#)
  - CAN\_STATE\_ERROR\_WARNING, [29](#)
  - CAN\_STATE\_SCANNING\_BAUDRATE, [29](#)
  - CAN\_STATE\_SLEEPING, [29](#)
  - CAN\_STATE\_STOPPED, [29](#)
  - can\_filter\_t, [28](#)
  - can\_frame\_t, [28](#)
  - RTCAN\_RTIOC\_RCV\_TIMEOUT, [20](#)
  - RTCAN\_RTIOC\_SND\_TIMEOUT, [20](#)
  - RTCAN\_RTIOC\_TAKE\_TIMESTAMP, [21](#)
  - SIOCGCANBAUDRATE, [22](#)
  - SIOCGCANCTRLMODE, [22](#)
  - SIOCGCANCUSTOMBITTIME, [23](#)
  - SIOCGCANSTATE, [23](#)
  - SIOCGIFINDEX, [24](#)
  - SIOCSCANBAUDRATE, [24](#)
  - SIOCSCANCTRLMODE, [25](#)
  - SIOCSCANCUSTOMBITTIME, [26](#)
  - SIOCSCANMODE, [26](#)
  - SOL\_CAN\_RAW, [27](#)
- CAN\_BITTIME\_BTR
  - CAN Devices, [28](#)
- CAN\_BITTIME\_STD
  - CAN Devices, [28](#)
- CAN\_BITTIME\_TYPE
  - CAN Devices, [28](#)
- CAN\_CTRLMODE\_3\_SAMPLES
  - CAN Devices, [17](#)
- CAN\_CTRLMODE\_LISTENONLY
  - CAN Devices, [17](#)
- CAN\_CTRLMODE\_LOOPBACK
  - CAN Devices, [18](#)
- CAN\_ERR\_LOSTARB\_UNSPEC
  - CAN Devices, [18](#)
- CAN\_MODE
  - CAN Devices, [28](#)
- CAN\_MODE\_SLEEP
  - CAN Devices, [28](#)
- CAN\_MODE\_START
  - CAN Devices, [28](#)
- CAN\_MODE\_STOP
  - CAN Devices, [28](#)
- CAN\_RAW\_ERR\_FILTER
  - CAN Devices, [18](#)
- CAN\_RAW\_FILTER
  - CAN Devices, [18](#)
- CAN\_RAW\_LOOPBACK
  - CAN Devices, [19](#)
- CAN\_RAW\_RECV\_OWN\_MSGS
  - CAN Devices, [20](#)
- CAN\_STATE
  - CAN Devices, [28](#)
- CAN\_STATE\_BUS\_OFF
  - CAN Devices, [29](#)
- CAN\_STATE\_ERROR\_ACTIVE
  - CAN Devices, [29](#)
- CAN\_STATE\_ERROR\_PASSIVE
  - CAN Devices, [29](#)
- CAN\_STATE\_ERROR\_WARNING
  - CAN Devices, [29](#)
- CAN\_STATE\_SCANNING\_BAUDRATE
  - CAN Devices, [29](#)
- CAN\_STATE\_SLEEPING
  - CAN Devices, [29](#)
- CAN\_STATE\_STOPPED
  - CAN Devices, [29](#)
- can\_bittime, [149](#)
- can\_bittime\_btr, [149](#)
- can\_bittime\_std, [150](#)
- can\_filter, [150](#)
  - can\_id, [151](#)
  - can\_mask, [151](#)
- can\_filter\_t
  - can\_id, [151](#)
  - can\_mask, [151](#)
  - can\_filter\_t

- CAN Devices, 28
- can\_frame, 151
  - can\_id, 152
- can\_frame\_t
  - CAN Devices, 28
- can\_id
  - can\_filter, 151
  - can\_frame, 152
- can\_ifindex
  - sockaddr\_can, 159
- can\_mask
  - can\_filter, 151
- Clock Services, 92
  - rtdm\_clock\_read, 92
  - rtdm\_clock\_read\_monotonic, 92
- close\_\_AF\_RTIPC
  - Real-time IPC protocols, 55
- close\_rt
  - rtdm\_operations, 156
- connect\_\_AF\_RTIPC
  - Real-time IPC protocols, 55
- Device Profiles, 146
  - RTIOC\_DEVICE\_INFO, 147
  - RTIOC\_PURGE, 147
- Device Registration Services, 82
  - RTDM\_CLOSING, 84
  - RTDM\_CREATED\_IN\_NRT, 84
  - RTDM\_DEVICE\_TYPE\_MASK, 84
  - RTDM\_EXCLUSIVE, 84
  - RTDM\_NAMED\_DEVICE, 84
  - RTDM\_PROTOCOL\_DEVICE, 84
  - rtdm\_close\_handler\_t, 84
  - rtdm\_context\_to\_private, 89
  - rtdm\_dev\_register, 89
  - rtdm\_dev\_unregister, 90
  - rtdm\_ioctl\_handler\_t, 85
  - rtdm\_open\_handler\_t, 85
  - rtdm\_private\_to\_context, 90
  - rtdm\_read\_handler\_t, 85
  - rtdm\_recvmsg\_handler\_t, 86
  - rtdm\_select\_bind\_handler\_t, 86
  - rtdm\_sendmsg\_handler\_t, 87
  - rtdm\_socket\_handler\_t, 87
  - rtdm\_write\_handler\_t, 87
- Driver Development API, 91
- EXPORT\_SYMBOL\_GPL
  - Synchronisation Services, 112–117
- getpeername\_\_AF\_RTIPC
  - Real-time IPC protocols, 57
- getsockname\_\_AF\_RTIPC
  - Real-time IPC protocols, 57
- getsockopt\_\_AF\_RTIPC
  - Real-time IPC protocols, 57
- IDDP\_LABEL
  - Real-time IPC protocols, 48

- IDDP\_POOLSZ
  - Real-time IPC protocols, 49
- IPCPROTO\_BUF
  - Real-time IPC protocols, 54
- IPCPROTO\_IDDP
  - Real-time IPC protocols, 54
- IPCPROTO\_IPC
  - Real-time IPC protocols, 54
- IPCPROTO\_XDDP
  - Real-time IPC protocols, 54
- include/rtdm/rtdm.h, 161
- include/rtdm/rtdm\_driver.h, 169
- include/rtdm/rtdm.h, 167
- include/rtdm/rtdm\_driver.h, 169
- include/rtdm/rtdm.h, 167
- include/rtdm/rtipc.h, 174
- include/rtdm/rtserial.h, 176
- include/rtdm/rtesting.h, 180
- Inter-Driver API, 73
  - rtdm\_accept, 74
  - rtdm\_bind, 74
  - rtdm\_close, 74
  - rtdm\_connect, 74
  - rtdm\_context\_get, 74
  - rtdm\_context\_lock, 76
  - rtdm\_context\_put, 76
  - rtdm\_context\_unlock, 77
  - rtdm\_getpeername, 77
  - rtdm\_getsockname, 77
  - rtdm\_getsockopt, 78
  - rtdm\_ioctl, 78
  - rtdm\_listen, 78
  - rtdm\_open, 78
  - rtdm\_read, 78
  - rtdm\_recv, 79
  - rtdm\_recvfrom, 79
  - rtdm\_recvmsg, 79
  - rtdm\_select\_bind, 79
  - rtdm\_send, 80
  - rtdm\_sendmsg, 80
  - rtdm\_sendto, 80
  - rtdm\_setsockopt, 80
  - rtdm\_shutdown, 81
  - rtdm\_socket, 81
  - rtdm\_write, 81
- Interrupt Management Services, 129
  - rtdm\_irq\_disable, 130
  - rtdm\_irq\_enable, 131
  - rtdm\_irq\_free, 131
  - rtdm\_irq\_get\_arg, 129
  - rtdm\_irq\_handler\_t, 130
  - rtdm\_irq\_request, 132
- ksrc/skins/rtdm/core.c, 186
- ksrc/skins/rtdm/device.c, 181
- ksrc/skins/rtdm/drvlib.c, 182
- ksrc/skins/rtdm/module.c, 186
- label
  - rtipc\_port\_label, 157

- nanosecs\_abs\_t
  - Real-Time Driver Model, [31](#)
- nanosecs\_rel\_t
  - Real-Time Driver Model, [31](#)
- Non-Real-Time Signalling Services, [133](#)
  - rtm\_nrtsig\_destroy, [133](#)
  - rtm\_nrtsig\_handler\_t, [133](#)
  - rtm\_nrtsig\_init, [134](#)
  - rtm\_nrtsig\_pend, [134](#)
- open\_rt
  - rtm\_device, [154](#)
- RTCAN\_RTIOC\_RCV\_TIMEOUT
  - CAN Devices, [20](#)
- RTCAN\_RTIOC\_SND\_TIMEOUT
  - CAN Devices, [20](#)
- RTCAN\_RTIOC\_TAKE\_TIMESTAMP
  - CAN Devices, [21](#)
- RTDM\_CLOSING
  - Device Registration Services, [84](#)
- RTDM\_CREATED\_IN\_NRT
  - Device Registration Services, [84](#)
- RTDM\_DEVICE\_TYPE\_MASK
  - Device Registration Services, [84](#)
- RTDM\_EXCLUSIVE
  - Device Registration Services, [84](#)
- RTDM\_EXECUTE\_ATOMICALLY
  - Synchronisation Services, [108](#)
- RTDM\_NAMED\_DEVICE
  - Device Registration Services, [84](#)
- RTDM\_PROTOCOL\_DEVICE
  - Device Registration Services, [84](#)
- RTDM\_SELECTTYPE\_EXCEPT
  - Synchronisation Services, [112](#)
- RTDM\_SELECTTYPE\_READ
  - Synchronisation Services, [112](#)
- RTDM\_SELECTTYPE\_WRITE
  - Synchronisation Services, [112](#)
- RTDM\_TIMEOUT\_INFINITE
  - Real-Time Driver Model, [30](#)
- RTDM\_TIMEOUT\_NONE
  - Real-Time Driver Model, [30](#)
- RTDM\_TIMERMODE\_ABSOLUTE
  - Timer Services, [102](#)
- RTDM\_TIMERMODE\_REALTIME
  - Timer Services, [102](#)
- RTDM\_TIMERMODE\_RELATIVE
  - Timer Services, [102](#)
- RTIOC\_DEVICE\_INFO
  - Device Profiles, [147](#)
- RTIOC\_PURGE
  - Device Profiles, [147](#)
- RTSER\_RTIOC\_BREAK\_CTL
  - Serial Devices, [66](#)
- RTSER\_RTIOC\_GET\_CONFIG
  - Serial Devices, [67](#)
- RTSER\_RTIOC\_GET\_CONTROL
  - Serial Devices, [67](#)
- RTSER\_RTIOC\_GET\_STATUS
  - Serial Devices, [68](#)
- RTSER\_RTIOC\_SET\_CONFIG
  - Serial Devices, [68](#)
- RTSER\_RTIOC\_SET\_CONTROL
  - Serial Devices, [69](#)
- RTSER\_RTIOC\_WAIT\_EVENT
  - Serial Devices, [69](#)
- Real-Time Driver Model, [30](#)
  - nanosecs\_abs\_t, [31](#)
  - nanosecs\_rel\_t, [31](#)
  - RTDM\_TIMEOUT\_INFINITE, [30](#)
  - RTDM\_TIMEOUT\_NONE, [30](#)
- Real-time IPC protocols, [45](#)
  - BUFP\_BUFSZ, [47](#)
  - BUFP\_LABEL, [48](#)
  - bind\_\_AF\_RTIPC, [54](#)
  - close\_\_AF\_RTIPC, [55](#)
  - connect\_\_AF\_RTIPC, [55](#)
  - getpeername\_\_AF\_RTIPC, [57](#)
  - getsockname\_\_AF\_RTIPC, [57](#)
  - getsockopt\_\_AF\_RTIPC, [57](#)
  - IDDP\_LABEL, [48](#)
  - IDDP\_POOLSZ, [49](#)
  - IPCPROTO\_BUFP, [54](#)
  - IPCPROTO\_IDDP, [54](#)
  - IPCPROTO\_IPC, [54](#)
  - IPCPROTO\_XDDP, [54](#)
  - recvmsg\_\_AF\_RTIPC, [58](#)
  - SO\_RCVTIMEO, [50](#)
  - SO\_SNDTIMEO, [50](#)
  - sendmsg\_\_AF\_RTIPC, [58](#)
  - setsockopt\_\_AF\_RTIPC, [59](#)
  - socket\_\_AF\_RTIPC, [59](#)
  - XDDP\_BUFSZ, [50](#)
  - XDDP\_EVTDOWN, [51](#)
  - XDDP\_EVTIN, [51](#)
  - XDDP\_EVTNOBUF, [51](#)
  - XDDP\_EVTOUT, [51](#)
  - XDDP\_LABEL, [51](#)
  - XDDP\_MONITOR, [52](#)
  - XDDP\_POOLSZ, [53](#)
- recvmsg\_\_AF\_RTIPC
  - Real-time IPC protocols, [58](#)
- rt\_dev\_accept
  - User API, [33](#)
- rt\_dev\_bind
  - User API, [33](#)
- rt\_dev\_close
  - User API, [33](#)
- rt\_dev\_connect
  - User API, [35](#)
- rt\_dev\_getpeername
  - User API, [35](#)
- rt\_dev\_getsockname
  - User API, [36](#)
- rt\_dev\_getsockopt
  - User API, [36](#)

- rt\_dev\_ioctl
  - User API, [37](#)
- rt\_dev\_listen
  - User API, [37](#)
- rt\_dev\_open
  - User API, [37](#)
- rt\_dev\_read
  - User API, [39](#)
- rt\_dev\_recv
  - User API, [39](#)
- rt\_dev\_recvfrom
  - User API, [40](#)
- rt\_dev\_recvmsg
  - User API, [40](#)
- rt\_dev\_send
  - User API, [41](#)
- rt\_dev\_sendmsg
  - User API, [41](#)
- rt\_dev\_sendto
  - User API, [42](#)
- rt\_dev\_setsockopt
  - User API, [42](#)
- rt\_dev\_shutdown
  - User API, [43](#)
- rt\_dev\_socket
  - User API, [43](#)
- rt\_dev\_write
  - User API, [43](#)
- rtdm\_accept
  - Inter-Driver API, [74](#)
- rtdm\_bind
  - Inter-Driver API, [74](#)
- rtdm\_clock\_read
  - Clock Services, [92](#)
- rtdm\_clock\_read\_monotonic
  - Clock Services, [92](#)
- rtdm\_close
  - Inter-Driver API, [74](#)
- rtdm\_close\_handler\_t
  - Device Registration Services, [84](#)
- rtdm\_connect
  - Inter-Driver API, [74](#)
- rtdm\_context\_get
  - Inter-Driver API, [74](#)
- rtdm\_context\_lock
  - Inter-Driver API, [76](#)
- rtdm\_context\_put
  - Inter-Driver API, [76](#)
- rtdm\_context\_to\_private
  - Device Registration Services, [89](#)
- rtdm\_context\_unlock
  - Inter-Driver API, [77](#)
- rtdm\_copy\_from\_user
  - Utility Services, [135](#)
- rtdm\_copy\_to\_user
  - Utility Services, [136](#)
- rtdm\_dev\_context, [152](#)
- rtdm\_dev\_register
  - Device Registration Services, [89](#)
- rtdm\_dev\_unregister
  - Device Registration Services, [90](#)
- rtdm\_device, [153](#)
  - open\_rt, [154](#)
  - socket\_rt, [154](#)
- rtdm\_device\_info, [155](#)
- rtdm\_event\_clear
  - Synchronisation Services, [117](#)
- rtdm\_event\_destroy
  - Synchronisation Services, [118](#)
- rtdm\_event\_init
  - Synchronisation Services, [118](#)
- rtdm\_event\_pulse
  - Synchronisation Services, [118](#)
- rtdm\_event\_select\_bind
  - Synchronisation Services, [119](#)
- rtdm\_event\_signal
  - Synchronisation Services, [119](#)
- rtdm\_event\_timedwait
  - Synchronisation Services, [120](#)
- rtdm\_event\_wait
  - Synchronisation Services, [120](#)
- rtdm\_free
  - Utility Services, [137](#)
- rtdm\_getpeername
  - Inter-Driver API, [77](#)
- rtdm\_getsockname
  - Inter-Driver API, [77](#)
- rtdm\_getsockopt
  - Inter-Driver API, [78](#)
- rtdm\_in\_rt\_context
  - Utility Services, [137](#)
- rtdm\_ioctl
  - Inter-Driver API, [78](#)
- rtdm\_ioctl\_handler\_t
  - Device Registration Services, [85](#)
- rtdm\_iomap\_to\_user
  - Utility Services, [137](#)
- rtdm\_irq\_disable
  - Interrupt Management Services, [130](#)
- rtdm\_irq\_enable
  - Interrupt Management Services, [131](#)
- rtdm\_irq\_free
  - Interrupt Management Services, [131](#)
- rtdm\_irq\_get\_arg
  - Interrupt Management Services, [129](#)
- rtdm\_irq\_handler\_t
  - Interrupt Management Services, [130](#)
- rtdm\_irq\_request
  - Interrupt Management Services, [132](#)
- rtdm\_listen
  - Inter-Driver API, [78](#)
- rtdm\_lock\_get
  - Synchronisation Services, [109](#)
- rtdm\_lock\_get\_irqsave
  - Synchronisation Services, [109](#)
- rtdm\_lock\_init

- Synchronisation Services, [110](#)
- rtm\_lock\_irqrestore
  - Synchronisation Services, [110](#)
- rtm\_lock\_irqsave
  - Synchronisation Services, [110](#)
- rtm\_lock\_put
  - Synchronisation Services, [111](#)
- rtm\_lock\_put\_irqrestore
  - Synchronisation Services, [111](#)
- rtm\_malloc
  - Utility Services, [138](#)
- rtm\_mmap\_to\_user
  - Utility Services, [139](#)
- rtm\_munmap
  - Utility Services, [140](#)
- rtm\_mutex\_destroy
  - Synchronisation Services, [122](#)
- rtm\_mutex\_init
  - Synchronisation Services, [122](#)
- rtm\_mutex\_lock
  - Synchronisation Services, [123](#)
- rtm\_mutex\_timedlock
  - Synchronisation Services, [123](#)
- rtm\_mutex\_unlock
  - Synchronisation Services, [124](#)
- rtm\_nrtsig\_destroy
  - Non-Real-Time Signalling Services, [133](#)
- rtm\_nrtsig\_handler\_t
  - Non-Real-Time Signalling Services, [133](#)
- rtm\_nrtsig\_init
  - Non-Real-Time Signalling Services, [134](#)
- rtm\_nrtsig\_pend
  - Non-Real-Time Signalling Services, [134](#)
- rtm\_open
  - Inter-Driver API, [78](#)
- rtm\_open\_handler\_t
  - Device Registration Services, [85](#)
- rtm\_operations, [155](#)
  - close\_rt, [156](#)
- rtm\_printk
  - Utility Services, [140](#)
- rtm\_printk\_ratelimited
  - Utility Services, [141](#)
- rtm\_private\_to\_context
  - Device Registration Services, [90](#)
- rtm\_ratelimit
  - Utility Services, [141](#)
- rtm\_read
  - Inter-Driver API, [78](#)
- rtm\_read\_handler\_t
  - Device Registration Services, [85](#)
- rtm\_read\_user\_ok
  - Utility Services, [142](#)
- rtm\_rcv
  - Inter-Driver API, [79](#)
- rtm\_rcvfrom
  - Inter-Driver API, [79](#)
- rtm\_rcvmsg
  - Inter-Driver API, [79](#)
- rtm\_rcvmsg\_handler\_t
  - Device Registration Services, [86](#)
- rtm\_rt\_capable
  - Utility Services, [142](#)
- rtm\_rw\_user\_ok
  - Utility Services, [143](#)
- rtm\_safe\_copy\_from\_user
  - Utility Services, [143](#)
- rtm\_safe\_copy\_to\_user
  - Utility Services, [144](#)
- rtm\_select\_bind
  - Inter-Driver API, [79](#)
  - Synchronisation Services, [124](#)
- rtm\_select\_bind\_handler\_t
  - Device Registration Services, [86](#)
- rtm\_selecttype
  - Synchronisation Services, [112](#)
- rtm\_sem\_destroy
  - Synchronisation Services, [125](#)
- rtm\_sem\_down
  - Synchronisation Services, [125](#)
- rtm\_sem\_init
  - Synchronisation Services, [125](#)
- rtm\_sem\_select\_bind
  - Synchronisation Services, [126](#)
- rtm\_sem\_timeddown
  - Synchronisation Services, [126](#)
- rtm\_sem\_up
  - Synchronisation Services, [127](#)
- rtm\_send
  - Inter-Driver API, [80](#)
- rtm\_sendmsg
  - Inter-Driver API, [80](#)
- rtm\_sendmsg\_handler\_t
  - Device Registration Services, [87](#)
- rtm\_sendto
  - Inter-Driver API, [80](#)
- rtm\_setsockopt
  - Inter-Driver API, [80](#)
- rtm\_shutdown
  - Inter-Driver API, [81](#)
- rtm\_socket
  - Inter-Driver API, [81](#)
- rtm\_socket\_handler\_t
  - Device Registration Services, [87](#)
- rtm\_strncpy\_from\_user
  - Utility Services, [145](#)
- rtm\_task\_busy\_sleep
  - Task Services, [95](#)
- rtm\_task\_current
  - Task Services, [95](#)
- rtm\_task\_destroy
  - Task Services, [95](#)
- rtm\_task\_init
  - Task Services, [97](#)
- rtm\_task\_join\_nrt
  - Task Services, [97](#)



- rtdm\_task\_proc\_t
  - Task Services, [95](#)
- rtdm\_task\_set\_period
  - Task Services, [98](#)
- rtdm\_task\_set\_priority
  - Task Services, [98](#)
- rtdm\_task\_sleep
  - Task Services, [99](#)
- rtdm\_task\_sleep\_abs
  - Task Services, [99](#)
- rtdm\_task\_sleep\_until
  - Task Services, [99](#)
- rtdm\_task\_unlock
  - Task Services, [100](#)
- rtdm\_task\_wait\_period
  - Task Services, [100](#)
- rtdm\_timer\_destroy
  - Timer Services, [103](#)
- rtdm\_timer\_handler\_t
  - Timer Services, [102](#)
- rtdm\_timer\_init
  - Timer Services, [103](#)
- rtdm\_timer\_mode
  - Timer Services, [102](#)
- rtdm\_timer\_start
  - Timer Services, [103](#)
- rtdm\_timer\_start\_in\_handler
  - Timer Services, [104](#)
- rtdm\_timer\_stop
  - Timer Services, [104](#)
- rtdm\_timer\_stop\_in\_handler
  - Timer Services, [105](#)
- rtdm\_toseq\_init
  - Synchronisation Services, [127](#)
- rtdm\_write
  - Inter-Driver API, [81](#)
- rtdm\_write\_handler\_t
  - Device Registration Services, [87](#)
- rtipc\_port\_label, [156](#)
  - label, [157](#)
- rtser\_config, [157](#)
- rtser\_event, [158](#)
- rtser\_status, [158](#)
  
- SIOCGCANBAUDRATE
  - CAN Devices, [22](#)
- SIOCGCANCTRLMODE
  - CAN Devices, [22](#)
- SIOCGCANCUSTOMBITTIME
  - CAN Devices, [23](#)
- SIOCGCANSTATE
  - CAN Devices, [23](#)
- SIOCGIFINDEX
  - CAN Devices, [24](#)
- SIOCSCANBAUDRATE
  - CAN Devices, [24](#)
- SIOCSCANCTRLMODE
  - CAN Devices, [25](#)
- SIOCSCANCUSTOMBITTIME
  - CAN Devices, [26](#)
- SIOCSCANMODE
  - CAN Devices, [26](#)
- SO\_RCVTIMEO
  - Real-time IPC protocols, [50](#)
- SO\_SNDTIMEO
  - Real-time IPC protocols, [50](#)
- SOL\_CAN\_RAW
  - CAN Devices, [27](#)
- sendmsg\_\_AF\_RTIPC
  - Real-time IPC protocols, [58](#)
- Serial Devices, [61](#)
  - RTSER\_RTIOC\_BREAK\_CTL, [66](#)
  - RTSER\_RTIOC\_GET\_CONFIG, [67](#)
  - RTSER\_RTIOC\_GET\_CONTROL, [67](#)
  - RTSER\_RTIOC\_GET\_STATUS, [68](#)
  - RTSER\_RTIOC\_SET\_CONFIG, [68](#)
  - RTSER\_RTIOC\_SET\_CONTROL, [69](#)
  - RTSER\_RTIOC\_WAIT\_EVENT, [69](#)
- setsockopt\_\_AF\_RTIPC
  - Real-time IPC protocols, [59](#)
- sipc\_port
  - sockaddr\_ipc, [160](#)
- sockaddr\_can, [159](#)
  - can\_ifindex, [159](#)
- sockaddr\_ipc, [159](#)
  - sipc\_port, [160](#)
- socket\_\_AF\_RTIPC
  - Real-time IPC protocols, [59](#)
- socket\_rt
  - rtdm\_device, [154](#)
- Synchronisation Services, [106](#)
  - EXPORT\_SYMBOL\_GPL, [112–117](#)
  - RTDM\_EXECUTE\_ATOMICALLY, [108](#)
  - RTDM\_SELECTTYPE\_EXCEPT, [112](#)
  - RTDM\_SELECTTYPE\_READ, [112](#)
  - RTDM\_SELECTTYPE\_WRITE, [112](#)
  - rtdm\_event\_clear, [117](#)
  - rtdm\_event\_destroy, [118](#)
  - rtdm\_event\_init, [118](#)
  - rtdm\_event\_pulse, [118](#)
  - rtdm\_event\_select\_bind, [119](#)
  - rtdm\_event\_signal, [119](#)
  - rtdm\_event\_timedwait, [120](#)
  - rtdm\_event\_wait, [120](#)
  - rtdm\_lock\_get, [109](#)
  - rtdm\_lock\_get\_irqsave, [109](#)
  - rtdm\_lock\_init, [110](#)
  - rtdm\_lock\_irqrestore, [110](#)
  - rtdm\_lock\_irqsave, [110](#)
  - rtdm\_lock\_put, [111](#)
  - rtdm\_lock\_put\_irqrestore, [111](#)
  - rtdm\_mutex\_destroy, [122](#)
  - rtdm\_mutex\_init, [122](#)
  - rtdm\_mutex\_lock, [123](#)
  - rtdm\_mutex\_timedlock, [123](#)
  - rtdm\_mutex\_unlock, [124](#)
  - rtdm\_select\_bind, [124](#)

- rt dm\_selecttype, 112
- rt dm\_sem\_destroy, 125
- rt dm\_sem\_down, 125
- rt dm\_sem\_init, 125
- rt dm\_sem\_select\_bind, 126
- rt dm\_sem\_timeddown, 126
- rt dm\_sem\_up, 127
- rt dm\_toseq\_init, 127

#### Task Services, 94

- rt dm\_task\_busy\_sleep, 95
- rt dm\_task\_current, 95
- rt dm\_task\_destroy, 95
- rt dm\_task\_init, 97
- rt dm\_task\_join\_nrt, 97
- rt dm\_task\_proc\_t, 95
- rt dm\_task\_set\_period, 98
- rt dm\_task\_set\_priority, 98
- rt dm\_task\_sleep, 99
- rt dm\_task\_sleep\_abs, 99
- rt dm\_task\_sleep\_until, 99
- rt dm\_task\_unblock, 100
- rt dm\_task\_wait\_period, 100

#### Testing Devices, 71

#### Timer Services, 102

- RTDM\_TIMERMODE\_ABSOLUTE, 102
- RTDM\_TIMERMODE\_REALTIME, 102
- RTDM\_TIMERMODE\_RELATIVE, 102
- rt dm\_timer\_destroy, 103
- rt dm\_timer\_handler\_t, 102
- rt dm\_timer\_init, 103
- rt dm\_timer\_mode, 102
- rt dm\_timer\_start, 103
- rt dm\_timer\_start\_in\_handler, 104
- rt dm\_timer\_stop, 104
- rt dm\_timer\_stop\_in\_handler, 105

#### User API, 32

- rt\_dev\_accept, 33
- rt\_dev\_bind, 33
- rt\_dev\_close, 33
- rt\_dev\_connect, 35
- rt\_dev\_getpeername, 35
- rt\_dev\_getsockname, 36
- rt\_dev\_getsockopt, 36
- rt\_dev\_ioctl, 37
- rt\_dev\_listen, 37
- rt\_dev\_open, 37
- rt\_dev\_read, 39
- rt\_dev\_recv, 39
- rt\_dev\_recvfrom, 40
- rt\_dev\_recvmsg, 40
- rt\_dev\_send, 41
- rt\_dev\_sendmsg, 41
- rt\_dev\_sendto, 42
- rt\_dev\_setsockopt, 42
- rt\_dev\_shutdown, 43
- rt\_dev\_socket, 43
- rt\_dev\_write, 43

#### Utility Services, 135

- rt dm\_copy\_from\_user, 135
- rt dm\_copy\_to\_user, 136
- rt dm\_free, 137
- rt dm\_in\_rt\_context, 137
- rt dm\_iomap\_to\_user, 137
- rt dm\_malloc, 138
- rt dm\_mmap\_to\_user, 139
- rt dm\_munmap, 140
- rt dm\_printk, 140
- rt dm\_printk\_ratelimited, 141
- rt dm\_ratelimit, 141
- rt dm\_read\_user\_ok, 142
- rt dm\_rt\_capable, 142
- rt dm\_rw\_user\_ok, 143
- rt dm\_safe\_copy\_from\_user, 143
- rt dm\_safe\_copy\_to\_user, 144
- rt dm\_strncpy\_from\_user, 145

#### XDDP\_BUFSZ

- Real-time IPC protocols, 50

#### XDDP\_EVTDOWN

- Real-time IPC protocols, 51

#### XDDP\_EVTIN

- Real-time IPC protocols, 51

#### XDDP\_EVTNOBUF

- Real-time IPC protocols, 51

#### XDDP\_EVTOUT

- Real-time IPC protocols, 51

#### XDDP\_LABEL

- Real-time IPC protocols, 51

#### XDDP\_MONITOR

- Real-time IPC protocols, 52

#### XDDP\_POOLSZ

- Real-time IPC protocols, 53